Architecture and Access Protocol for a Wavelength–Selective Single–Hop Packet Switched Metropolitan Area Network

vorgelegt von
Diplom–Ingenieur
Martin Maier

Von der Fakultät IV – Elektrotechnik und Informatik
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Ingenieurwissenschaften
– Dr.–Ing. –

Promotionsausschuß:

Vorsitzender: Prof. Dr.–Ing. Gerhard Mönich
Berichter: Prof. Dr.–Ing. Adam Wolisz
Berichter: Prof. Dr.rer.nat. Ulrich Killat
Berichter: Prof. Dr.–Ing. Klaus Petermann

Tag der wissenschaftlichen Aussprache: 6. Juni 2003

Berlin 2003

D 83

Abstract

Current SONET/SDH metro ring networks create the so–called metro gap which prevents high–speed clients from tapping into the vast amounts of bandwidth available in the backbone. To bridge this gap we propose and investigate by means of analysis and simulation a novel arrayed–waveguide grating (AWG) based single–hop wavelength division multiplexing (WDM) network with a physical star topology. Star networks offer a better optical power budget than ring and bus networks and are easy to install, configure, manage, and troubleshoot. Unlike their multi-hop counterparts single–hop networks provide a minimum mean hop distance (unity), inherent transparency, future–proofness, easy upgradability, simplified management, and an improved throughput–delay performance since no bandwidth is wasted due to packet forwarding. Owing to spatial wavelength reuse the AWG keeps the wavelength pool small which enables the deployment of fast tunable transceivers with a negligible tuning time. The proposed network consists of an AWG with wavelength–insensitive combiners (splitters) attached to each AWG input (output) port, where the splitters are used for realizing optical multicasting. Due to its completely passive nature the network is cost–effective and reliable. Each node at the network periphery is equipped with one single tunable transceiver and one low–cost broadband light source which is spectrally sliced for broadcasting control information. Direct sequence spread spectrum techniques are deployed to improve the network security and enable simultaneous transmission of data and control within the same wavelength channel without requiring any additional control wavelength and receiver. Each node has access to all wavelengths resulting in efficient multicasting and load balancing. Wavelengths are on–demand allocated by means of a reservation medium access control (MAC) protocol which provides both packet and circuit switching. By not fixed assigning the reservation slots and deploying performance enhancing reservation ALOHA (R–ALOHA) and code division multiple access (CDMA) the network is made scalable. All nodes have global knowledge and schedule variable–size packets on a deterministic first–come–first–served and first–fit basis without requiring explicit acknowledgements guaranteeing fairness, decreased latency, and quality of service (QoS) and completely avoiding both channel and receiver collisions of data packets. The network efficiency is significantly increased by spatially reusing all wavelengths at each AWG port, exploiting multiple free spectral ranges (FSRs) of the AWG, and wormhole scheduling. For unicast packet switched traffic our AWG based network achieves a mean wavelength utilization of more than 100%, a mean channel utilization of approximately 53%, and a mean aggregate throughput that is about 70% larger than the maximum aggregate throughput of the DT–WDMA access protocol that runs on a PSC based single–hop metro WDM network. Partitioning in conjunction with spatial wavelength reuse and exploiting the reservation phases of the MAC protocol allow for very efficient multicasting. Furthermore, the presented single–hop network significantly reduces the complexity of the protocol stack in that routing is replaced with simple wavelength tuning and the data link layer can be omitted. By using computationally efficient multiobjective optimization tech-

niques the proposed network and protocol are Pareto–optimized in order to efficiently enable dynamic multiservice convergence. We develop and evaluate the new concept of heterogeneous protection of the single point of failure which clearly outperforms its conventional counterpart in terms of throughput and delay and is generally applicable in all single–hop WDM networks. Our findings aim at contributing to bridge the bandwidth abyss at the metro level, enable new applications benefitting from the huge amounts of bandwidth available in the backbone, stimulate revenue growth, and possibly offer a turnaround in the telecom sector.

Preface

After completing my master's thesis on passive optical WDM networks under his guidance, Dr. Andreas Gladisch from Deutsche Telekom offered me the opportunity to become a Ph.D. candidate. In June of 1999 Deutsche Telekom has granted me a two–year scholarship for conducting research in the area of optical wavelength division multiplexing (WDM) networks. The target of my future work was described quite vaguely. Within my Ph.D. research studies I was supposed to find and investigate possible deployments of the periodic wavelength routing characteristics of an arrayed–waveguide grating (AWG) in passive IP–over–WDM networks. Professor Adam Wolisz from the Technical University Berlin was willing to accept me in his Telecommunication Networks Group (TKN) and to supervise my dissertation since then.

After countless fruitful discussions with my colleague Hagen Woesner I have decided to focus on single–hop WDM networks while he was investigating AWG based multihop WDM networks. Up to 1999 most single–hop WDM networks were based on a passive star coupler (PSC). To verify the potential of AWG based single–hop networks I started comparing them with their PSC based counterparts. It turned out that the performance of AWG based single–hop networks was quite promising due to spatial wavelength reuse. To better understand AWG based WDM networks Hagen and I have examined AWG based single–hop and multihop networks in greater detail. Through this comparison we have gained valuable insight in their respective merits and drawbacks. These architectural investigations formed the initial phase of my dissertation. Next, it was time to fix ideas about a novel cost–effective architecture and medium access control (MAC) protocol that was able to transport variable–size packets efficiently. During the design and first performance evaluations of my proposed AWG based single–hop network I got in touch with Professor Martin Reisslein. Martin has shown interest in my work and has been contributing to the work in a very positive way since then. After specifying the network and providing preliminary evaluation results we payed attention to more detailed features of the MAC protocol. During this time period, Professor Wolisz announced the possibility of collaborating with Professor Michael Scheutzow who as a mathematician was interested in analyzing the performance of communication networks. Very soon Michael and I got in touch and after some time we came up with an analytical model of more sophisticated aspects of our MAC protocol. In early 2001 I visited Martin at Arizona State University. During my visit I was happy to meet Hyo–Sik Yang and Chun Fan who were just beginning to pursue the Ph.D. degree. Martin and I introduced them to our work and formulated open questions which remained to be addressed. Since then Hyo–Sik and Chun have been members of our group and we have received significant help from them. After the US visit I concentrated on examining the feasibility and transmission limits of our proposed network. Together with Martin Herzog we examined additional features of the protocol by means of simulation which have been analytically intractable. Within his master's thesis Martin Herzog made a significant contribution to determining the transmission limits of our AWG based single–hop network.

The results generated and presented in this dissertation represent a very stout body of work. Parts of the dissertation have led to several publications in journals and conference proceedings, as listed in Appendix A. In my opinion, the results form an extraordinary collection of accomplishments and will hopefully make a significant contribution to the metropolitan network infrastructure of the next generation Internet and to its consumers.

Finally, I want to thank my advisor Professor Adam Wolisz. He generously supported me and provided valuable criticism. He gave me the independence to follow my own research interests. I also want to thank Professor Ulrich Killat from the Technical University Hamburg–Harburg and Professor Klaus Petermann from the Technical University Berlin for acting as supervisors. I want to thank Deutsche Telekom and the Federal Ministry of Education and Research for funding our work within the national TransiNet project. I am particularly grateful to my mentors and peers Professor Martin Reisslein, Professor Michael Scheutzow, Hyo–Sik Yang, Chun Fan, Martin Herzog, and Hagen Woesner for helpful comments and important contributions. I also owe thanks to all my colleagues of the TKN group. Most importantly, I want to thank my family, friends, and Constanze for supporting me during sometimes quite troublesome and tedious periods of my dissertation.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The ultimate goal of the Internet and communication networks in general is to provide access to information *when we need it, where we need it, and in whatever format we need it* [Muk00]. To achieve this goal wireless and optical technologies play a key role in future communication networks. Wireless and optical networks can be thought of as quite complementary. Optical fiber does not go everywhere, but where it does go, it provides a huge amount of available bandwidth. Wireless networks, on the other hand, potentially go almost everywhere, but provide a highly bandwidth–constrained transmission channel, susceptible to a variety of impairments [Ram02]. As opposed to the wireless channel, the fiber exhibits a number of advantageous transmission properties such as low attenuation, large bandwidth, and immunity from electromagnetic interference.

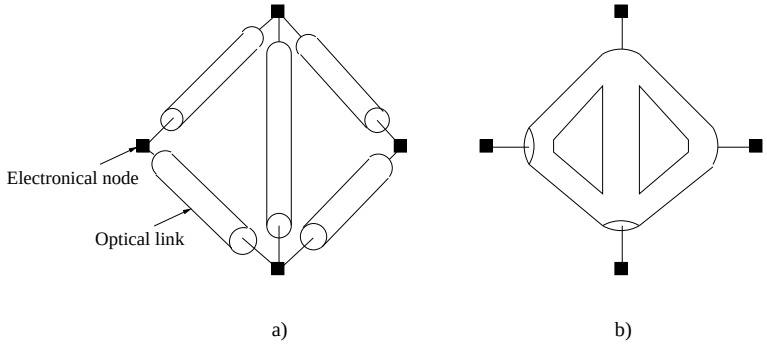

Figure 1.1: Optical networks: a) First generation, b) second generation.

Optical networks are the medium of choice to provide enough bandwidth to the ever increasing number of users and bandwidth–hungry applications, e.g., video conferences, distributed games, visualization, supercomputer interconnection, or medical imaging which does not trust image–compressing techniques [RS92]. There are two generations of optical networks. As illustrated in Fig. 1.1 a), in the first–generation optical networks copper links are replaced with fiber links while the nodes at either end of the fiber remain electronical. In such *opaque* optical networks optical–electronic–optical (OEO) conversions of the signal take place at each node [Gre93]. Initially, each fiber carried only one single wavelength such as in FDDI and IEEE 802.6 DQDB (for the definition of these and all following acronyms please refer to Appendix B). To cope with the exponentially increasing amount of data traffic in an economic fashion and to fully exploit the gain bandwidth of the optical Erbium doped fiber amplifier (EDFA) *wavelength*

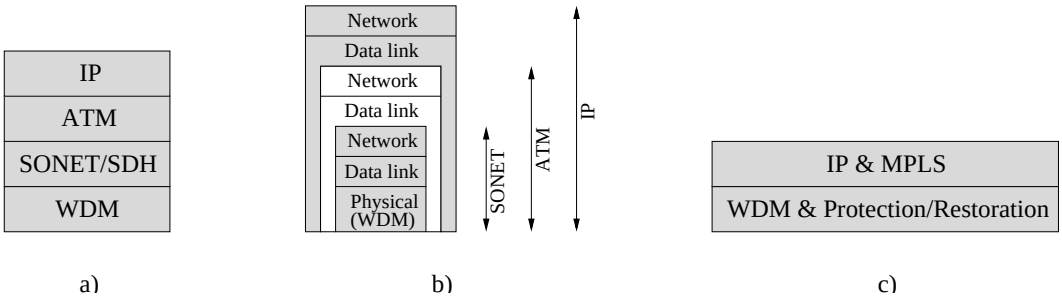

Figure 1.2: Protocol stacks: a) IP/ATM/SONET(SDH)/WDM, b) detailed layer architecture of IP/ATM/SONET/WDM, c) slim protocol stack IP/WDM.

division multiplexing (WDM) was introduced in the 90's. With WDM each fiber link carries multiple wavelengths, each operating at any arbitrary line rate, e.g., electronic peak rate. After providing these huge pipes attention turned from optical transmission to optical networking [Ram02]. In the second–generation optical networks OEO conversions occur only at the source and destination nodes while all intermediate nodes are *optically bypassed*, as illustrated in Fig. 1.1 b). By optically bypassing nodes the electro–optic bottleneck is alleviated and the number of electronic port cards can be reduced at each node resulting in significantly reduced network costs, which is one of the most important drivers for optics [Mok00]. Furthermore, the resultant end–to–end all–optical light paths are able to provide *transparent* channels to end users which are free to choose bit rate, modulation format, and protocol. This transparency easily enables the support of various legacy as well as future services.

Traffic in future optical networks is widely expected to be predominantly IP based. Typically, IP packet transmission in optical WDM networks is done in a mix–and–match fashion. Fig. 1.2 a) depicts the ubiquitous IP/ATM/SONET(SDH)/WDM protocol stack which today's networks deploy for transmitting IP packets. Variable–size IP packets are segmented into fixed–size ATM cells which in turn are transmitted in SONET/SDH frames over optical WDM links. This protocol stack requires a number of mapping operations between the various protocols which not only increase the network costs and complexity but also tend to cause a computational bottleneck in high–speed networks. Moreover, as Fig. 1.2 b) illustrates the protocol stack is rather inefficient since the same aspects concerning the network and data link layer are addressed in each protocol leading to redundant functionalities and rather complex layer interworking schemes [RS98][JFW+00]. For instance, in case of a link and/or node failure the affected traffic can either be re–routed by updating the routing tables of the IP routers or saved by activating protection mechanisms at the SONET cross–connects. Those parallel activities have to be harmonized in order to enable a consistent and efficient network operation. Another example for redundant functionalities is traffic engineering, which can be done in both the IP and ATM layers. IP traffic has some distinctive properties such as burstiness, asymmetry, and server–based congestions (hot spots) [LES00]. Time–division–multiplexing (TDM) based SONET/SDH can carry this type of traffic only very inefficiently since it is designed for synchronous and symmetric traffic. ATM also suffers from inefficiencies which stem from a large ATM cell overhead, which is also known as cell tax. To avoid these inefficiencies and to simplify network operation the complex IP/ATM/SONET(SDH)/WDM layer structure is likely to be replaced with a significantly less complex IP/WDM protocol stack [Gre01]. The ATM function of traffic engineering (e.g., quality of service (QoS)) will likely be absorbed into the IP layer by means of multiprotocol label switching (MPLS), and the transport capabilities of SONET/SDH

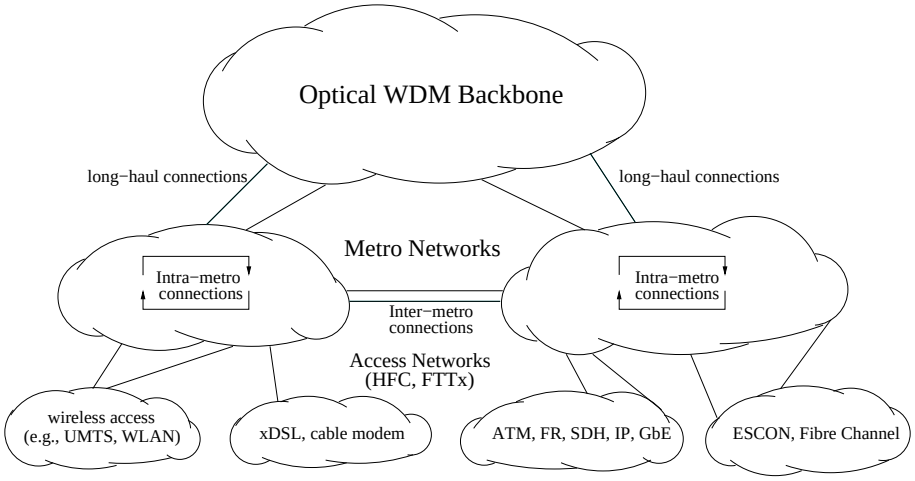

Figure 1.3: Network hierarchy (acronyms are defined in Appendix B).

(e.g., protection and restoration) will be absorbed by the optical WDM layer. Consequently, future optical WDM networks are expected to have a slim IP/WDM protocol stack, depicted in Fig. 1.2 c).

Optical IP/WDM networks hold great promise to efficiently provide a large amount of bandwidth at low network complexity, even though the current photonic technology pose some constraints on their feasibility and cost–effectiveness [Gre92]. Note, however, that current complexity and cost issues in optical WDM networks might be overcome in the future. The classic example of cost reduction of photonic technology is the short–wavelength laser diode used in compact disc players. Who would have thought that lasers would eventually cost U$5 a piece and be more plentiful than phonograph needles [Gre91]?

1.1 Motivation and scope

As illustrated in Fig. 1.3, the hierarchy of communication networks can be viewed as consisting of backbone, metro, and access networks where the latter ones collect (distribute) data from (to) different clients such as wireless stations and LANs. Let us first consider the backbone wide area networks (WANs). Future optical WDM backbone networks seem to be converging to a two–layer infrastructure in which the data transport is by means of interconnected *all–optical islands of transparency* while the remainder of the communication layers are based on IP [Gre01]. The islands of transparency contain very high–speed WDM links. Recently, it was shown that transmission of 80 wavelengths each operating at 42.7 Gb/s is feasible over a distance of 5200 km resulting in a bandwidth–distance product of 16.6 petabit·km/s for a nominal 40 Gb/s transmission [ZLG$^+$02]. These links are stitched together by all–optical cross connects (OXCs) and add–drop multiplexers (OADMs). (OXCs are basically space switches which connect input and output ports on a per–wavelength basis. OADMs allow for locally dropping and adding one or more wavelengths from or to an incoming or outgoing fiber link, respectively.) The OEO boundaries between islands will be retained due to management, jurisdiction, billing, and/or signal regeneration issues [RPS$^+$00]. While early OADMs and OXCs provide only frozen paths, micro–electro–mechanical–systems (MEMS) can now be used to realize reconfigurable OADMs and OXCs [NR01][YLG01]. Those devices make the network more flexible and robust

against link and/or node failures. Recently, much research has focused on controlling such reconfigurable network elements [LYWK02]. MPLS routing and signalling protocols have been extended and modified in order to enable optical label switching (OLS) (using wavelengths as labels leads to a specific kind of OLS termed MPλS) and to enrich optical WDM networks with point–and–click provisioning [AR01], protection [FV00], restoration [DY01][YDA00][MCG+02], traffic engineering [WLL+01], and future services, e.g., rent–a–wavelength. Since MEMS have a switching time of about 10 ms only circuit switching can be realized. Those circuits correspond to wavelength light paths and can either be permanent or reconfigured, for example twice a day, in order to provide sufficient bandwidth to offices during the day and to residential areas in the evening. Future OXCs and OADMs may exhibit a significantly improved switching time of a few nanoseconds by deploying semiconductor optical amplifiers (SOAs) or lithium–niobate based components.

Meanwhile, optical burst switching (OBS) appears to be a viable intermediate step between current lambda switching and future packet switching [VCR00]. OBS aggregates multiple packets, e.g., IP packets, into bursts at the optical network edge and makes a one–way reservation by sending a control packet on one wavelength prior to transmitting the corresponding data burst on another wavelength. OBS does not require buffering at intermediate nodes. Moreover, as opposed to lambda switching OBS allows for statistical multiplexing, resulting in a higher channel utilization compared to circuit switching. OBS is also able to provide differentiated services by controlling the offset time between control packet and data burst and/or by using fiber delay lines (FDLs) at intermediate nodes [YQD00][YQD01][Qia00]. Other approaches to improve the switching granularity include optical flow routing (OFR) [HS01] and photonic slot routing (PSR) [CEFS99][EFW01]. To further improve the switching granularity, optical packet switching (OPS) has been receiving much attention both in backbone and metro WDM networks [YMD00][YYMD01][OSHT01][JCD+01]. Several OPS node architectures have been proposed [HNC+99][HA00][CJZ+00]. However, due to the lack of optical RAM buffers and economical optical processing of packet headers the current OPS state–of–the–art provides only limited performance [EBS02].

Now, let us take a look at the network periphery, depicted in Fig. 1.3. Current Gigabit Ethernet (GbE) LANs together with the IEEE 802.3ae 10 GbE standard completed in 2002 are expected to provide sufficient bandwidth for at least the next few years. Phone companies typically deploy some form of digital subscriber loop (DSL) and cable companies deploy cable modems. The bottleneck in the first mile is likely to be further alleviated by deploying the IEEE 802.3ah Ethernet in the first mile (EFM) standard which is anticipated by September 2003. These broadband access technologies in conjunction with next–generation wireless services, e.g., UMTS and wireless LANs (WLANs), and high–speed protocols such as ATM, Frame Relay (FR), IP, ESCON, and Fibre Channel will require a huge amount of bandwidth and quality of service (QoS) support from the networks higher up in the hierarchy in Fig. 1.3.

Between those high–speed clients and the huge pipes in the backbone lie the access and metro networks. Access networks were initially hybrid fiber coax (HFC) systems where only the feeding part between the central office and the remote node of the network is optical while the distribution network between the remote node and the subscribers is still electrical. Both the telcos and cable providers are steadily moving the fiber–to–copper discontinuity point out toward the network periphery [Gre01]. As a consequence, so–called FTTx access networks have received a great deal of attention. FTTx networks, e.g., fiber to the curb (FTTC) or fiber to the home (FTTH), are completely optical, i.e., the signal is transmitted via fiber from the central office close or all the way to the subscribers. Typically, for cost reasons such all–optical

access networks are unpowered which are accordingly called passive optical networks (PONs). PONs have been considered for the access network since the mid 90's, well before the Internet spurred bandwidth demand [Fri97]. Recently evolving Ethernet PONs are promising candidates for providing enough bandwidth for efficient transport of data traffic [KP02].

Today's metropolitan area networks (MANs) are mostly SONET/SDH ring networks which suffer from a number of drawbacks:

- Circuit provisioning in SONET/SDH networks takes too much time, typically from 6 weeks to 6 months. Thus, quick service provisioning is impossible.
- SONET/SDH equipment is very expensive and significantly decreases the margins in the cost–sensitive metro market where costs are shared by fewer subscribers than in the backbone. The high costs prevent new companies from entering the metro market.
- Upgrading a SONET/SDH ring affects all nodes, not only the corresponding source and destination nodes wishing to communicate at higher data rates.
- SONET/SDH's automatic protection switching (APS) mechanism (1+1 protection) is highly inefficient in terms of bandwidth since both working and protection paths carry the same traffic.
- SONET/SDH is designed for symmetric traffic. As a consequence, asymmetric IP traffic is transported only very inefficiently.
- The voice–centric TDM operation is unable to efficiently support bursty data traffic resulting in wasted bandwidth.

The above mentioned inefficiencies of SONET/SDH ring networks create a severe bandwidth bottleneck at the metro level. This so–called *metro gap* prevents the high–speed clients (and also service providers) from tapping into the vast amounts of bandwidth available in the backbone [Ali][HM02]. This bottleneck is expected to get even more severe due to the fact that an increasing part of IP traffic will be local by placing more proxy cache servers in the metro networks in order to achieve network latency reduction, server load balancing, and higher content availability [BO00]. The increase in the usage of cellular phones and handheld devices for Internet services will increase the amount of locally stored content that is accessed and has to be updated regularly, especially as home appliances, cars, and other electronic devices begin to utilize the metro network [KWSR01]. In addition, Napster being simply a precursor of the controversial file sharing, future peer–to–peer applications where each attached user will also operate as a server will dramatically increase the amount of intra–MAN traffic. *To bridge this bandwidth abyss between high–speed clients and the backbone novel metro architectures and protocols have to be developed.*

Recently, research has begun to pay attention to alleviating the metro gap [KSW+01]. The importance of the metro gap is also reflected by the large number of recently initiated standardization activities and industry forums such as IETF WG IPoRPR, IEEE 802.17 RPRWG, the Metro Ethernet Forum (MEF) [MEF], and the resilient packet ring (RPR) alliance [RPR] which comprises more than 70 companies.

In this work, we focus on packet switched metro WDM networks where packets are stored in electronic RAMs rather than optical fiber delay lines. With the architecture and protocol presented and examined in this work we aim at enabling high–speed clients and service providers to bridge the metro gap and fully benefit from the abundant bandwidth in the backbone in an efficient, cost–effective, and future–proof way.

1.2 Approach

In our approach we make use of the respective strengths of the optical and electronic domains while avoiding their respective shortcomings: Transmission is done in the optical domain while buffering and logical operations are done in the electronic domain. Due to the lack of optical buffers (RAMs) we consider *bufferless* optical networks. We propose a *switchless* network architecture which is completely passive. Passive networks are not only quite reliable but also naturally move the intelligence towards the network periphery resulting in reduced network costs and a simplified network operation, administration, and maintenance (OAM). The network under consideration is *wavelength selective*. In such a static wavelength–selective network each source node is able to reach different destination nodes by simply changing the transmitting wavelength. In doing so, the conventional store–and–forward packet switching paradigm is replaced with *wavelength tuning* on a *per–packet* basis at the network periphery. In wavelength–selective networks each wavelength is routed only to a small part of the network. In the remaining parts of the network the same wavelength can be reused. The resultant *spatial wavelength reuse* not only increases the *degree of concurrency* but also keeps the number of wavelengths small. A small wavelength pool requires wavelength–agile transceivers to be tunable only over a small tuning range. This in turn allows for deploying tunable transceivers whose tuning time is *several orders of magnitude* smaller compared to transceivers with a relatively large tuning range. Our proposed network has a physical *star* topology. Star topologies exhibit a better optical power budget than bus networks. Star and bus networks suffer from splitting and tapping loss, respectively. While the overall tapping loss (in dB) in bus networks grows linearly with the number of network nodes, the splitting loss (in dB) in star networks grows only logarithmically. Moreover, physical star configurations are easy to install, configure, manage, and troubleshoot. On top of the physical star network we embed a logical *single–hop* network. In single–hop networks any pair of source and destination nodes communicate directly without any forwarding intermediate nodes. As opposed to their multihop counterparts, single–hop networks have a number of advantageous properties: (*i*) The mean hop distance is minimum (unity), (*ii*) no bandwidth is wasted due to forwarding since each node reaches any arbitrary destination in one single hop, (*iii*) each node has to process only those packets which are addressed to itself resulting in reduced nodal processing requirements, (*iv*) transparency is inherently provided, and (*v*) upgrading a given source–destination pair involves only these two nodes, i.e., no intermediate nodes have to be upgraded as opposed to multihop networks. Note that single–hop networks also provide a significantly *decreased protocol stack complexity*: Due to the single–hop communication routing/forwarding is not required. As a consequence, the network layer is effectively absent. In addition, packets traverse one single passive optical hop between source and destination, resulting in very miniscule bit error probabilities. Thus, error detection and correction at the data link layer can be omitted and residual transmission errors can be handled by the transport layer.

1.2.1 Architecture

Our proposed single–hop WDM network architecture is based on a wavelength–selective arrayed–waveguide grating (AWG). A wavelength–insensitive combiner (splitter) is attached to each AWG input (output) port. Nodes are attached to the combiners and splitters via a pair of fibers, one for transmission and the other one for reception. The AWG allows for extensive *spatial wavelength reuse* by simultaneously deploying all wavelengths at each AWG input port without channel collisions at the AWG output ports. Furthermore, by using *multiple free spectral*

ranges (FSRs) of the AWG the network capacity can be easily upgraded. By spatially reusing all wavelengths and using multiple FSRs of the underlying AWG the network degree of concurrency is significantly increased. The wavelength–insensitive splitters allow for optical *multicasting* by distributing an incoming signal to all attached nodes. By equipping each node with one single tunable transceiver and one off–the–shelf broadband light source the node structure is rather simple and economic.

1.2.2 Protocol

In the proposed network each node has access to all wavelengths. Wavelengths are assigned *dynamically on–demand* on a *per–packet* basis by means of a distributed reservation medium access control (MAC) protocol. The degree of concurrency is improved by sending reservation requests *simultaneously* with data by means of spreading techniques. Reservation slots are not fixed assigned to nodes. Thus, without any change of the MAC protocol new nodes can join the reservation process making the network easily *scalable*. Owing to a cyclic reservation timing structure all nodes are able to acquire *global knowledge* after one end–to–end propagation delay, i.e., half the round–trip time, without requiring explicit acknowledgements. By exploiting the global knowledge, data packet transmissions are scheduled in a distributed way such that data packets are transmitted *completely collisionfree*. Hence, bandwidth is not wasted due to data packet collisions, resulting in an improved channel utilization. The protocol supports both *packet and circuit switching* where the latter one can be used for providing *guaranteed QoS*.

1.3 Methodology and thesis outline

We first give an overview of metro WDM network architectures and MAC protocols found in the literature and summarize learnt lessons. After discussing the pros and cons of the various network candidates we decide to concentrate on physical star topologies. We then compare different possible node structures by means of analysis in order to better understand the performance of AWG based logical multihop and single–hop WDM networks. Based on the obtained results we decide to focus on AWG based single–hop networks. To assess their potential we conduct an analytical performance comparison between AWG and passive star coupler (PSC) based single–hop networks where the latter one has been the most commonly studied type of single–hop network at the time of writing.

Incorporating the aforementioned learnt lessons we develop a novel network architecture and MAC protocol. Through analysis and simulation we investigate the performance of the network and the impact of the various architecture and protocol parameters on the network performance and their interplay are examined. After gaining some insight into the dynamics of the proposed network we optimize the parameter settings in terms of throughput–delay performance. We then proceed to study the feasibility and transmission limits of the developed network by means of analysis and simulation. In addition, the network performance is examined by using packet header trace files as traffic load in the simulation. Finally, to remove the single point of failure and improve the survivability of our network we present a novel protection scheme whose throughput–delay performance is evaluated analytically together with verifying simulations.

The outline of the thesis is as follows:

In Chapter 2, we introduce and discuss various optical components which will be used in our architectural investigations in Chapter 4 and as building blocks in our proposed network

described in Chapter 5. Special attention is thereby paid to the AWG and its transmission characteristics and possible applications. Different types of transmitters and receivers are presented. We discuss the tuning properties of the transceivers, i.e., tuning range and tuning time, which play a major role in our work. We also address several transmission impairments encountered in optical fiber networks which come into play when examining the feasibility of our network in Chapter 8.

Chapter 3 describes the distinctive properties of metro networks. We provide a review of the related work on metro WDM networks and testbeds which typically have either a physical ring or star topology. The discussed star metro WDM networks belong to the family of single–hop networks since communications between any arbitrary pair of nodes are done in one single hop. We extensively discuss the state–of–the–art of single–hop WDM network architectures and MAC protocols and introduce a taxonomy for them. Learnt lessons are summarized and guidelines for the design of MAC protocols in single–hop WDM networks are formulated. On the basis of these guidelines we develop our MAC protocol in Chapter 5.

In Chapter 4 we investigate metro WDM network architectures in greater detail. We conduct analytical comparisons of different metro WDM network and node architecture candidates and discuss their merits and drawbacks. Due to their numerous advantages we put particular emphasis on physical star topologies. First, we compare AWG based star networks where nodes are equipped with either fixed–tuned or tunable transceivers. Depending on the node structure we obtain either logical multihop or logical single–hop networks which are embedded on the physical AWG based star topology. By means of analysis we examine both logical topologies in terms of mean hop distance and capacity, i.e., maximum achievable throughput. Secondly, we analyze the throughput–delay performance and packet loss of AWG and PSC based single–hop WDM networks and demonstrate the potential of an AWG based single–hop WDM network by comparing it to its PSC based counterpart.

Based on the comparison results of Chapter 4 we introduce our proposed AWG based single–hop metro WDM network in Chapter 5. After discussing underlying principles we describe in detail the network and node architecture which makes use of the components presented in Chapter 2. In addition, we specify the MAC protocol. The key parameters of both architecture and MAC protocol are identified and summarized.

Chapter 6 investigates the impact of the architecture and protocol parameters on the network performance by means of analysis and simulation. We proceed in two steps. First, we consider different parameters in isolation to better understand their impact. Secondly, we examine several parameters concurrently for different network configurations in order to gain some insight into the interplay of the parameters. Different performance metrics are considered. In extensive simulations we relax the assumptions made in the foregoing analyses for tractability reasons and investigate additional aspects of our network. We also provide a benchmark comparison of our proposed AWG based network with a PSC based network.

In Chapter 7 we address the dimensioning and reconfiguration of our proposed network. We introduce and develop an optimization approach which finds the optimal setting of the architecture (hardware) and protocol (software) parameters for efficiently supporting various traffic types with different throughput–delay requirements, thereby allowing for efficient multiservice convergence.

In Chapter 8 we study the feasibility of our proposed network architecture. We investigate the impact of various phyiscal transmission impairments discussed in Chapter 2 on the transmission limitations of our network. In addition, we use packet header trace files to evaluate the network performance.

Similar to other single–hop networks, our AWG based metro WDM network suffers from a single point of failure. In Chapter 9 we tackle this problem. We propose the novel concept of heterogeneous protection where the working and protection devices are not identical. More specifically, our proposed heterogeneous protection scheme aims at combining the strengths of AWG and PSC which have been compared previously in Chapter 4 such that the network efficiency and thoughput–delay performance are significantly improved.

Chapter 10 concludes the thesis. Apart from summarizing the main contributions of our work we also outline possible future research avenues which are worth being examined in detail. Furthermore, we review the current standardization activities on metro WDM networks and discuss how our work fits into this context.

Chapter 2

Basics

In this chapter, we introduce several components which are used in our architectural comparisons and as building blocks for our proposed network. In Section 2.1, we consider these components in isolation and describe their respective properties. In Chapter 5 we will discuss how the components can be connected in order to form our WDM network. Section 2.2 deals with several transmission impairments encountered in optical networks. Their impact on our proposed network will be investigated when discussing the network feasibility in Chapter 8. For an in–depth discussion and detailed information on the following and other components and transmission impairments the interested reader is referred to [Muk97][RS98][MS00].

2.1 Components

The following components serve as fundamental building blocks for the design of WDM networks. In our description we concentrate on components which are of importance for the remainder of this work.

2.1.1 Combiners and splitters

Coupler is a general term that covers all devices that combine light into and/or split light out of a fiber. Combiners are devices that combine light from different fibers. Splitters, as the name suggests, separate light into several fibers. Both combiners and splitters are passive devices. The most common splitter is a 1×2 splitter as shown in Fig. 2.1 a). The ratio of the output powers is called *splitting ratio* α and can be controlled. A fraction α of the input power is distributed to one output and the remaining fraction $(1 - \alpha)$ to the other output. Expressing the splitting ratio in dB provides the *splitting loss*. For a two–port splitter a splitting ratio of $50 : 50$ is very popular which results in a splitting loss of 3 dB at each output port. Couplers are also used to tap off a small portion of the power from a light stream for monitoring purposes or other reasons. Such couplers are called taps and are designed with values of α close to 1, typically $0.90 - 0.95$.

When turned around, a splitter can be used as a combiner, as depicted in Fig. 2.1 b). An input signal to the 2×1 combiner experiences a power loss of 3 dB, which is sometimes referred to as *combining loss*. By deploying a combiner together with a splitter, couplers can be realized. Fig. 2.1 c) shows a 2×2 coupler consisting of a 2×1 combiner, immediately followed by a 1×2 splitter, which has the effect of broadcasting the signals from two input fibers onto two output fibers. For a coupling ratio of $50 : 50$ we obtain the so–called 3–dB coupler where each

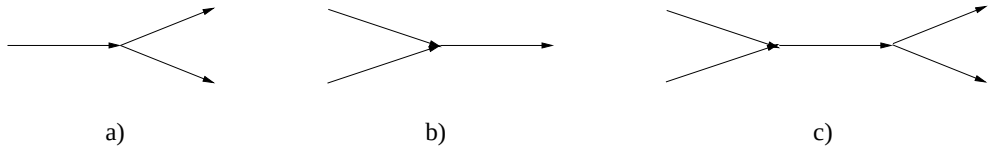

Figure 2.1: a) Splitter, b) combiner, and c) coupler.

input signal is equally distributed to both output ports. In addition to the 50 : 50 power split incurred in a coupler, a signal experiences insertion loss and also return and excess loss. The insertion loss is the fraction of power (usually expressed in dB) that is lost between the input and output ports of the device (coupler in this case). If the signal enters an input of the coupler a small amount of power is reflected in the opposite direction and is directed back to the inputs of the coupler (return loss). Excess loss is caused by manufacture imperfections due to the very small dimensions. A coupler can be made either wavelength independent or wavelength selective, meaning its coupling behavior depends on the wavelength.

2.1.2 Passive Star Coupler (PSC)

An $N \times N$ passive star coupler (PSC) is a natural generalization of the 3–dB 2×2 coupler [Dra89][DHKK89][OTS$^+$91][OOOK92]. As illustrated in Fig. 2.2, it is an N–input, N–output device with the property that the power from each input P_{in} is divided equally among all the outputs. Hence, the optical power that each output receives P_{out} equals

$$P_{out} = \frac{P_{in}}{N}, \tag{2.1}$$

which translates into a splitting ratio $\alpha = 1/N$ or equivalently a splitting loss of $10 \log_{10} N$ dB. One way to implement the PSC is to suitably interconnect a number of 3–dB couplers. This approach requires $(N/2 \cdot \log_2 N)$ 3–dB couplers to construct an $N \times N$ PSC.

2.1.3 Arrayed–Waveguide Grating (AWG)

In this section, we discuss the arrayed–waveguide grating (AWG), which is also known as phased array (PHASAR) or waveguide grating router (WGR), in greater detail since this component plays a major role in this work. We briefly explain the underlying physical concepts. Formulas are given for the free spectral range (FSR), channel spacing, and full width at half maximum (FWHM). Clearly, the channel spacing denotes the spectral distance between two adjacent wavelength channels. The FSR is the spectral range between two successive passbands of the

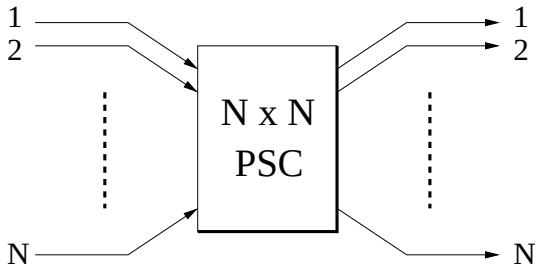

Figure 2.2: An $N \times N$ passive star coupler (PSC).

Parameter	Meaning
c	Light velocity in vacuum
n_g	Group index of the waveguides
ΔL	Path length difference of arrayed waveguide
d	Pitch of arrayed waveguide
θ_i	Diffraction angle in input slab waveguide
θ_o	Diffraction angle in output slab waveguide
λ	Optical wavelength
Δf	Frequency channel spacing
Δx	Spacing of input/output waveguides
L_f	Focal length of focusing slab waveguide
m	Diffraction order
n_s	Effective refractive index of slab waveguide
n_c	Effective refractive index of channel waveguide
ω_o	Spot size (half width at $1/e^2$ of maximum power)

Table 2.1: Definition of AWG parameters.

AWG. The FWHM is a measure of the full width of a passband between the points where the transfer function is half of its maximum. These quantities determine the number of available wavelengths and the channel capacity. Moreover, we discuss AWG based applications which form the building blocks of WDM networks.

Transmission characteristics

The definitions of parameters used to describe the transmission characteristics of an AWG are listed in Table 2.1 [TOTI95].

An $N \times N$ AWG is schematically shown in Fig. 2.3, where $N \geq 2$. It consists of N input–output waveguides, two focusing slab waveguides (free propagation regions) and an arrayed waveguide grating, where the lengths of adjacent waveguides differ by a constant value. The waveplate at the symmetry line of the device eliminates the polarization dependence. Thus, polarization independent AWGs can be realized [IIH+01]. The excess loss is about 0.4 dB. Both slab waveguides work as identical $N \times M$ star couplers, where $M \gg N$, so that all the light power diffracted in the slab can be collected. If $M \gg N$ the crosstalk near the center of a passband is reduced compared with $M = N$. The signal from any of the N input ports is distributed over the M outputs of the slab waveguide to the array inputs. Each input light is diffracted in the input slab, passed through the arrayed waveguides, focused in the output slab, and coupled into the output waveguides.

The arrayed waveguides introduce wavelength–dependent phase delays such that only frequencies with a phase difference of integer times 2π interfere constructively in the output slab waveguide. Thus, each output port carries periodic pass frequencies. The spacing of these periodic pass frequencies is called free spectral range (FSR) and is approximated by [TOTI95]

$$\mathrm{FSR} = \frac{c}{n_g \left(\Delta L + d \sin\theta_i + d \sin\theta_o\right)} \quad [\mathrm{Hz}], \tag{2.2}$$

or simply [Zir98]

$$\mathrm{FSR} = \frac{\lambda^2}{n_g \Delta L} \quad [\mathrm{m}]. \tag{2.3}$$

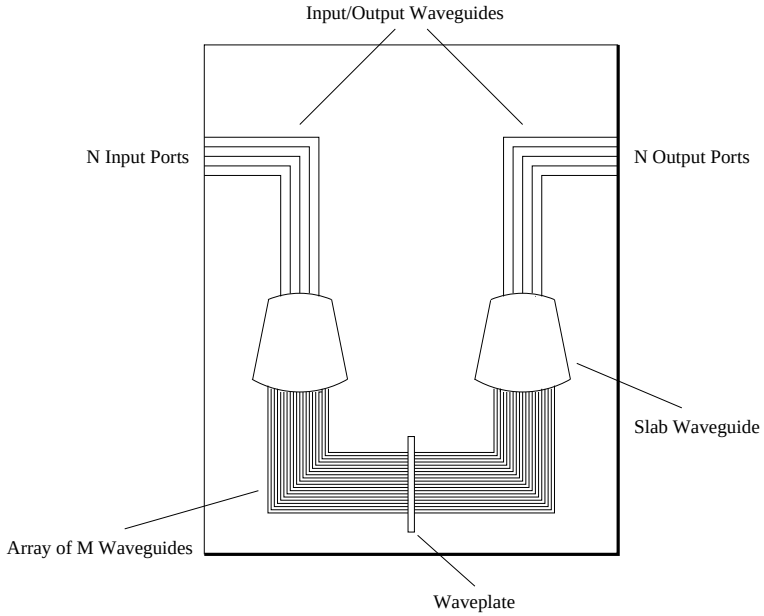

Figure 2.3: Schematic layout of an $N \times N$ AWG.

In each FSR an $N \times N$ AWG accepts a total of N wavelengths from each input port and it transmits each wavelength to a particular output port. Each output port receives N wavelengths, one from each input port. There exist cyclic wavelength permutations at the output waveguides if different input waveguides are used. In Fig. 2.4 the routing connectivity of an 8×8 AWG is illustrated. Each optical frequency (alternatively, we could also say each wavelength) gives routing instructions that are independent of the input port. Thus, f_k's routing information is to exit the output port that is $(k-1)$ ports below the corresponding input port, i.e., f_1 goes from input port 1 to exit port 1 and from input port 5 to exit port 5. Similarly, f_3 incident on input port 1 is directed to output port 3, whereas if f_3 were incident on port 5, it would be directed to output port 7. Due to the periodicity property of the AWG the optical frequency f_9 (i.e., one free spectral range higher than f_1) entering port 1 exits at port 1 like f_{17} and other frequencies separated by an integral number of free spectral ranges [Fri97]. But there are also AWGs with different channel routing patterns, e.g., channel f_i entering at input j exits at output k, where $k = (8 - i + j)_{mod8} + 1$ with $i \in \mathbb{N}$ and $j, k \in \{1, 2, \ldots, 8\}$ [SB98]. The wavelength routing pattern of an AWG can be formally described by means of a wavelength transfer matrix [Ogu96]. Owing to the wavelength permutations at its output ports the AWG belongs to the family of permutation wavelength routers [SB99] (permutation wavelength routers are a special case of so–called Latin routers [BH93]). An $N \times N$ AWG provides full $N \times N$ interconnection. Using one FSR a total of N^2 simultaneous connections are possible. Note that an $N \times N$ PSC is capable of simultaneously carrying only N channels.

The following transmission characteristics of an $N \times N$ AWG play a major role in WDM networks [TOTI95]. AWGs exhibit a low insertion loss of about $3 - 5$ dB. Also, AWGs with uniform loss across all channels can be realized [OHI+97]. The frequency channel spacing of the

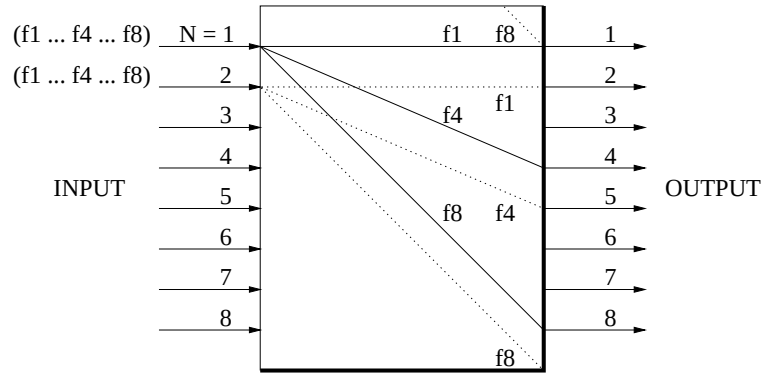

Figure 2.4: Connectivity of an 8×8 AWG.

multiplexer is given by

$$\Delta f = \frac{\Delta x}{L_f}\left(\frac{d\theta}{df}\right)^{-1} = \frac{\Delta x}{L_f}\left(\frac{m\lambda^2}{n_s dc}\frac{n_g}{n_c}\right)^{-1} \quad \text{[Hz]}. \tag{2.4}$$

The channel spacing is typically 100 or 200 GHz. A 100 GHz channel spacing in the low loss region at 1.55 μm corresponds to a 0.8 nm channel spacing, thus resulting in dense wavelength division multiplexing (DWDM). A 64 × 64 AWG with 0.4 nm (50 GHz) channel spacing was reported in [OMS95]. Recently, a 400–channel 25–GHz spacing AWG was reported in [HHK$^+$01].

The frequency response of the AWG is important for applications. It is the same as that of a Gaussian bandpass filter. The full width at half maximum (FWHM) is given by

$$\text{FWHM} = \frac{2\sqrt{\ln 2}\omega_o \Delta f}{\Delta x} \quad \text{[Hz]}. \tag{2.5}$$

Generally, the FWHM is about 30% of the channel spacing. The shape of Gaussian bandpass filters places tight restrictions on the wavelength tolerance of laser diodes and requires accurate temperature control for both AWGs and laser diodes. Frequency fluctuations in the light source result in power penalties in order to achieve the same BER at the receiver. Moreover, the passband width of cascaded AWGs becomes much narrower than that of the single–stage AWG filter. Recently, AWG multiplexers with a flat spectral response were reported [OS96][TBB$^+$97][KS01]. The 3dB bandwidth could be increased up to 124 GHz for a 200 GHz channel spacing and the interchannel crosstalk from a given wavelength channel to neighbouring channels is less than −27 dB. (Crosstalk is the general term given to the effect of other signals on the desired signal. Two forms of crosstalk arise in WDM systems: Interchannel and intrachannel crosstalk. Interchannel crosstalk occurs if the crosstalk signal is at a different wavelength than the desired signal's wavelength. Intrachannel takes place if the crosstalk signal is at the same wavelength as that of the desired signal.) But this comes at the expense of an about 3 dB higher insertion loss. AWGs with flat and broadened frequency response relax the above mentioned restrictions. In general, the crosstalk is about −30 dB. Hence, channel demultiplexing is possible with negligible power penalty, i.e., the crosstalk is small enough not to increase the BER significantly.

The pass frequency shift due to fluctuations of the multiplexer temperature has also to be taken into account. The slope equals approximately −1.5 GHz/°C [TOTI95]. The power penalty can be avoided by applying a thermo–electric cooler (TEC). Thus, an AWG can operate over

Insertion loss	$\sim 3-5$ dB
Excess loss	~ 0.4 dB
Interchannel crosstalk	~ -30 dB
Channel spacing	$\{25, 50, 100, 200\}$ GHz
FWHM	$\sim 30\%$ of channel spacing
Number of ports	$2 \ldots 400$

Table 2.2: Key parameters of AWG.

the temperature range from $0-85°$C [Ton98]. Alternatively, the temperature control can be avoided by deploying athermal all–polymer AWGs [KYZ+01]. An AWG can be made more robust, compact, and economical if the entire component is integrated onto a single chip. The chip size can be as small as 1×1 cm^2 [DEK91][ZDJ92][dBDH+02]. The above mentioned key transmission parameters of an AWG and typical parameter values or ranges are summarized in Table 2.2.

Table 2.3 summarizes the properties of the AWG and contrasts them to the PSC. Note that AWGs can support broad– and multicasting if a broadband light source is used, e.g., a light emitting diode (LED). This input signal is spectrally sliced by the AWG such that a portion of the broadband input signal is transmitted to each output port, as we will see shortly in Section 2.1.4. Privacy means that wavelengths are routed only to output ports they are destined to as opposed to the broadcast–and–select PSC where every output can receive all wavelengths at the same time. As a consequence, if the same wavelength is fed into two or more PSC input ports simultaneously channel collisions occur at all PSC output ports. In other words, the PSC does not allow for spatial reuse of wavelengths at multiple input ports. Conversely, the wavelength–routing AWG allows for spatial wavelength reuse without resulting channel collisions at the AWG output ports. Therefore, an $N \times N$ PSC is able to support at most N transmissions simultaneously, each using a different wavelength. Whereas an $N \times N$ AWG does not impose any restrictions on input signals allowing for up to N^2 simultaneous transmissions without channel collisions at the AWG output ports.

	PSC	AWG
broadcasting	yes	no
multicasting	yes	no
wavelength routing	no	yes
spatial wavelength reuse	no	yes
periodicity	no	yes
splitting loss	yes	no
privacy	no	yes
channel collision	yes	no
number of simultaneous transmissions	N	N^2

Table 2.3: Property comparison of PSC and AWG.

Next, we highlight the main applications which try to exploit the properties of an AWG in order to create building blocks for the realization of WDM networks.

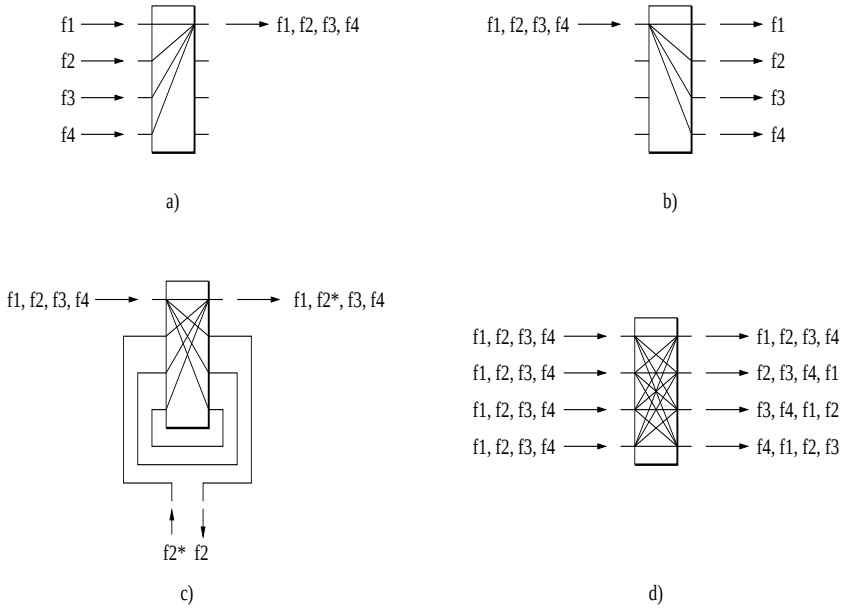

Figure 2.5: Basic usage scenarios of $N \times N$ AWG: a) Multiplexer, b) demultiplexer, c) add–drop multiplexer, d) $N \times N$ full–interconnect wavelength router.

Applications

The $N \times N$ AWG provides four basic usage scenarios [TOTI95]. It can operate as a simple $N \times 1$ multiplexer and a $1 \times N$ demultiplexer. As an add/drop multiplexer (ADM) the AWG carries out both multiplexing and demultiplexing simultaneously. The AWG can also be used as a full–interconnection wavelength router. All four functions are summarized in Fig. 2.5 for $N = 4$. Full–interconnection wavelength router can be employed in wavelength–addressed star topologies whereas ADMs are well suited for applications in wavelength–addressed ring or bus networks.

The conventional structure of an ADM is based on one $1 \times N$ demultiplexer and one $N \times 1$ multiplexer. Both demultiplexer and multiplexer must be carefully adjusted so as to have identical wavelength responses. In [TIK+93] an ADM with 15 add–drop channels was proposed. This ADM consists of a single 16×16 AWG with loop–back optical paths connecting each output port with its corresponding input port. One input port and its corresponding output port are reserved as common input and output ports for the transmission line. A desired wavelength λ_i can be dropped and added by opening one of the loop–back paths corresponding to λ_i. Wavelength channel selectors were accomplished by inserting laser diode amplifiers (LDA) into the loop–back fiber paths. Any wavelength can be selected by switching the LDAs on/off [TIIN96]. By using a single AWG troublesome tuning of wavelength responses can be avoided. We also note that an ADM can be applied as a comb filter for suppressing the optical fiber amplifier noise in DWDM transmission lines.

Some other applications of the AWG were reported in [GKW94]. As a discretely tunable receiver, an AWG is employed as a $1 \times N$ demultiplexer to separate a WDM signal into individual channels, which in turn are coupled to a photodiode array. A fast transistor switch array sends the intensity–modulated (IM) data from one or more selected photodiode elements to a common

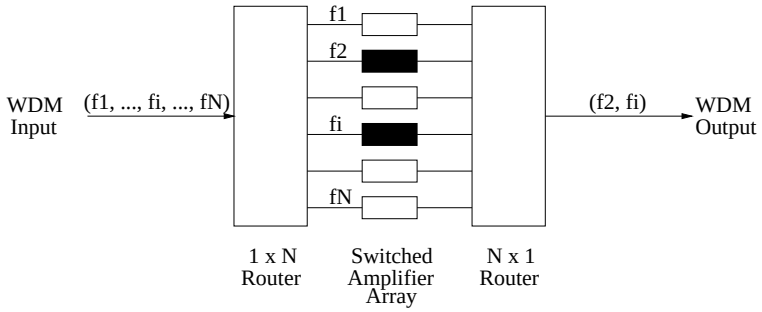

Figure 2.6: Discretely tunable filter/equalizer.

amplifier and receiver output. The complete device can be built on a single semiconductor chip. Switching times of ∼1 ns are feasible. Similarly, an AWG can be deployed as a discretely tunable filter/equalizer, a discretely tunable single– or multifrequency laser, or as a wideband tunable laser/filter. In all cases, the incoming WDM signal is demultiplexed and coupled to a switched amplifier array which selects the desired wavelengths. Fig. 2.6 illustrates a discretely tunable filter/equalizer.

Optical cross–connects (OXC) improve the flexibility and survivability of a system. They provide network restoration and can reconfigure the network to accomodate load changes and to compensate for link and/or device failures. An OXC performs this function entirely in the optical domain. An OXC is an $N \times N \times M$ component with N input fibers, N output fibers and M wavelength channels on each fiber, as shown in Fig. 2.7. A demultiplexer is attached to each input fiber. Each output from a single demultiplexer goes into a unique λ layer. Each λ layer has a space–division switch that directs each channel to a selected multiplexer. Each multiplexer collects light from M space–division switches and multiplexes the wavelengths onto a single output fiber. Ideally, the OXC is bidirectional as well as transparent to transmission format.

The conventional structure requires $2N$ multiplexers/demultiplexers and M space–division switches. Such an OXC can be constructed with N AWGs acting as demultiplexers and N AWGs acting as multiplexers. Costs can be reduced by using several input/output ports of a single AWG simultaneously. A $2 \times 2 \times 8$ OXC using only two AWGs was reported in [McG98]. Two

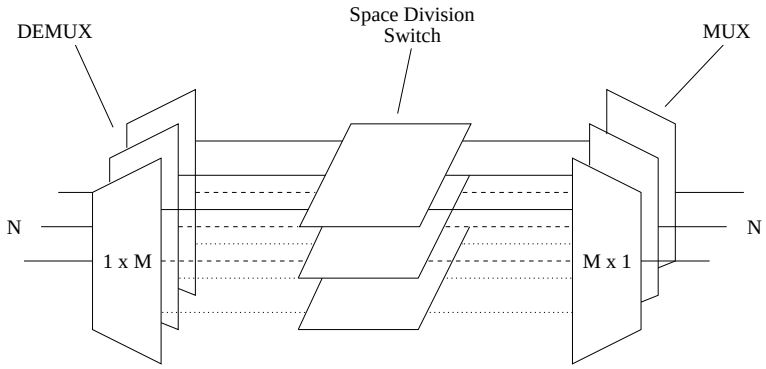

Figure 2.7: Schematic of an OXC.

fibers were attached to two ports of each AWG. In addition, a $2 \times 2 \times 4$ OXC with only one AWG in loopback configuration was mentioned. In this configuration, each wavelength passes through the AWG twice in the same direction. In the first pass, the AWG acts as a demultiplexer, in the second pass, it acts as a multiplexer. Wavelength misalignment is avoided since a single AWG is used, thus preventing power and crosstalk penalty. Similarly, an optical add–drop multiplexer (OADM) can be realized with a single AWG in foldback configuration.

Several other applications make use of the wavelength multiplexing/demultiplexing nature of the AWG. For example, much attention has been paid to placing an AWG at the remote node of passive optical networks (PONs) in order to create a switched WDM overlay [FIM+94][ZJS+95] [HHSW96][RIB+00][DFS+01][MMPS00]. The AWG can also be used to build dispersion (slope) compensators [TTH+00][HYM+00][TGI01], DWDM wavelength interleavers [HCL01], and multiwavelength lasers and receivers [Zir98][OYA01]. In [FHA98] the AWG is applied as an optical packet synchronizer which is capable of compensating for different packet delays. Recently, optical switches with an AWG based passive core and attached tunable transmitters (44 nm tuning range, nanosecond tuning time) or tunable wavelength converters were reported [GBC+00][CHA+01].

2.1.4 Transmitters and receivers

Besides the above mentioned components we need transmitters and receivers for building a WDM communication network. A transmitter comprises a light source, a modulator, and supporting electronics. A receiver is composed of an optical filter, a photodetector, a demodulator, and supporting electronics. In the following we discuss different types of transmitters and receivers which are considered in our architectural comparisons and our proposed node architecture in the remainder of this work. Specifically, we pay attention to the tuning range and tuning time of the different light sources and optical filters since the tuning characteristics of transceivers play a major role in our work.

Broadband light sources

The light output of a broadband light source has a broad spectrum in the range of $10 - 100$ nm. Light emitting diodes (LEDs) are a very common and cost–effective example of a broadband light source. Due to their relatively small bandwidth–distance product LEDs are mainly applied where the data rates are low and/or distances are short. Typical output powers are of the order of -10 dBm. However, superluminescent diodes with an output power in single–mode fiber of 18.0 dBm and a 3dB–bandwidth of 35 nm are also commercially available [opt].

In low–speed, low–budget wavelength–sensitive systems, LED slicing provides a cheap alternative to an array of expensive lasers [RHZ+88][WLK+88][WKR+88][IFD95][LRI00]. Fig. 2.8 schematically illustrates the spectral slicing of a broadband LED signal by a wavelength–routing AWG. At each AWG output port a different slice of the original broadband input signal is received by nodes, each slice carrying the same information. Thus, one LED can be shared by a number of receivers.

Lasers

To achieve a significantly increased bandwidth–distance product lasers are deployed. Essentially, a laser is an optical amplifier enclosed within a reflective cavity that causes the light to oscillate

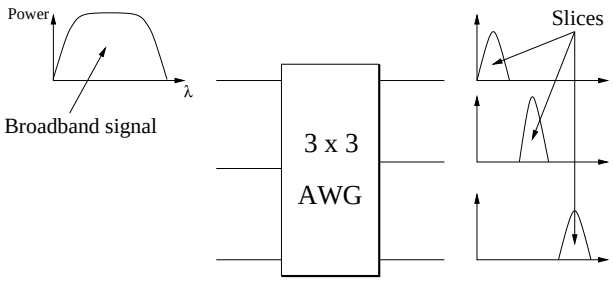

Figure 2.8: Spectral slicing of a broadband signal.

via positive feedback. Lasers are capable to achieve high output powers, typically between 0 and 10 dBm.

Lasers can either be fixed tuned to a nominal wavelength (though it can drift with temperature and age) or tunable, where tunable lasers can either be continuously or discretely tunable. Since only wavelengths which match the period and refractive index of the laser will be constructively reinforced, a laser can be tuned by controlling the cavity length and/or the refractive index of the lasing medium. Common examples are mechanically, acousto–optically, electro–optically, and injection–current tunable lasers. Most mechanically tuned lasers use an external Fabry–Perot cavity whose length is physically adjusted. Mechanically tunable lasers exhibit a relatively wide tuning range of up to 500 nm but a relatively slow tuning time of $1 - 10$ ms. In an acousto–optic or electro–optic laser the refractive index in the external cavity is changed by using either sound waves or electrical current, respectively. An acousto–optic laser combines a moderate tuning range of ~ 100 nm with a moderate tuning time of ~ 10 μs. Electro–optical lasers can be tuned over a range of $10 - 15$ nm within a few nanoseconds. Injection–current–tuned lasers form a family of light sources which allow wavelength selection via a diffraction grating, e.g., distributed feedback (DFB) and distributed Bragg reflector (DBR) lasers. Tuning is achieved by changing the injected current density and thereby the refractive index. This type of laser typically consists of multiple sections in order to allow for independent control of output power and wavelength of the laser [KVG$^+$90]. Recently, fast tunable multisection transmitters which can be tuned to adjacent wavelengths within 4 ns [FSA$^+$00][LRB00] and over a wide range of ~ 30 nm within 15 ns [SWR$^+$01] were reported. In particular, so–called SG–DBR lasers hold promise for use as fast tunable transmitters with a wide tuning range and high output power [Mas00][WRRW00].

The tuning ranges and tuning times of the different transmitter types are summarized in Table 2.4. Note that instead of tunable lasers one might use an array of fixed–tuned lasers, each operating at a different wavelength, or multifrequency lasers [Zir98].

Transmitter Type	*Tuning Range*	*Tuning Time*
Mechanically tunable	500 nm	1–10 ms
Acousto–optic	∼100 nm	∼10 μs
Electro–optic	10–15 nm	1–10 ns
Injection current	∼30 nm	15 ns

Table 2.4: Transmitters: Tuning ranges and tuning times.

Optical filters

Optical filters are used to select a slice of a broadband signal or one wavelength out of the WDM comb. The selected wavelength is subsequently opto–electrically converted by a photodetector. Optical filters are either fixed tuned or tunable, whereby tunable filters can either be continuously or discretely tunable. Examples for fixed–tuned filters are diffraction gratings, dielectric thin–film filters, and fiber Bragg gratings (FBGs) [Muk97]. Tunable optical filters encompass mechanically, thermally, acousto–optically, electro–optically tuned filters, and liquid–crystal Fabry–Perot filters [SB98]. In the following, we describe the different types of tunable optical filters and their properties in greater detail.

Mechanically tunable filters consist of one (or more) cavity formed by two parallel mirrors (facets). By mechanically adjusting the distance between the mirrors different wavelengths can be selected. This type of filter has a tuning range of ~ 500 nm and a tuning time in the range of $1 - 10$ ms.

The Mach–Zehnder interferometer (MZI) is an example for a thermally controlled optical filter. In an MZI a splitter splits the incoming light into two waveguides and a combiner recombines the signals at the outputs of the waveguides. A thermally adjustable delay element controls the optical path length in one of the waveguides. Due to the resulting phase difference a single desired wavelength can be selected by means of constructive interference. An MZI can be tuned over > 10 nm within a few milliseconds.

In acousto–optic tunable filters (AOTFs) a sound wave periodically changes the refractive index of the filtering medium which enables the medium to act as a grating. By changing the frequency of the sound wave a single optical wavelength can be chosen to pass through while the remainder of the wavelengths interfere destructively. If more than one sound wave is applied more than one wavelength can be filtered out. One drawback of AOTFs is that they are unable to filter out crosstalk from adjacent channels if the channels are closely spaced, thus limiting the number of channels. AOTFs can be tuned over a range of ~ 100 nm within ~ 10 μs.

Electro–optic tunable filters (EOTFs) use electrodes which rest in the filtering medium. Currents are applied to change the refractive index of the filtering medium which allows a desired wavelength to pass through while others interfere destructively. The tuning time is limited only by the speed of the electronics. Hence, EOTFs can be tuned on the order of $1 - 10$ ns. However, EOTFs provide a relatively small tuning range of ~ 15 nm.

Liquid–crystal (LC) Fabry–Perot filters appear to be an inexpensive filter technology with low power requirements. The design of a LC filter is similar to the design of a Fabry–Perot filter, but the cavity consists of a liquid crystal. The refractive index of the LC is controlled by an electrical current to filter out the corresponding wavelength. The tuning time is on the order of $0.5 - 10$ μs and the tuning range is $30 - 40$ nm.

The tuning ranges and tuning times of the different receiver types are summarized in Table 2.5. Note that alternatively to tunable optical filters, arrays of fixed–tuned receivers or multiwavelength receivers can be deployed [Ton98][OYA01].

2.2 Transmission impairments

To build communication systems the above described components are connected by fibers. In such a system, a light signal which propagates from the transmitter to the receiver undergoes a number of impairments. These transmission impairments are discussed next and their impact on our proposed network is investigated in Chapter 8 when we discuss the network's feasibility

Receiver Type	*Tuning Range*	*Tuning Time*
Mechanically tunable	500 nm	1–10 ms
Thermally tunable	> 10 nm	1–10 ms
Acousto–optic	$\sim$100 nm	$\sim$10 μs
Electro–optic	10–15 nm	1–10 ns
Liquid crystal	30–40 nm	0.5–10 μs

Table 2.5: Receivers: Tuning ranges and tuning times.

issues.

2.2.1 Attenuation

Beside the optical power loss caused by the components, the fiber further reduces the signal power. Fig. 2.9 shows the attenuation loss of a fiber as a function of wavelength [BSdMN01]. The peak in loss in the 1400 nm region is due to hydroxyl ion (OH^-) impurities in the fiber. However, in Lucent's AllWave fiber this peak is reduced significantly. In today's optical communications systems three wavelength bands are used: 0.85 μm, 1.3 μm, and 1.55 μm, where the latter band provides the smallest attenuation of ~ 0.25 dB/km.

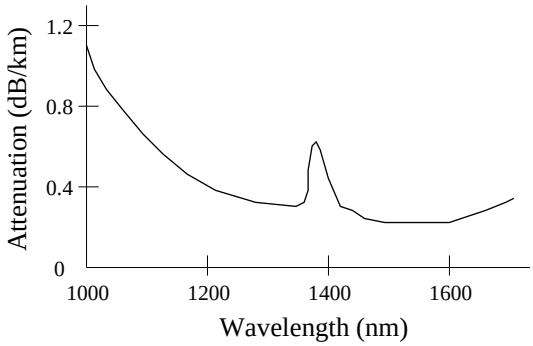

Figure 2.9: Attenuation of an optical fiber.

2.2.2 Dispersion

Dispersion is the name given to any effect wherein different components of the transmitted signal travel at different velocities in the fiber, arriving at different times at the receiver. As a result, the pulse widens and causes intersymbol interference (ISI). Thus, dispersion limits the minimum bit spacing, i.e., the maximum transmission rate. The amount of accumulated dispersion depends on the length of the link. The important forms of dispersion are modal dispersion, chromatic (material) dispersion, waveguide dispersion, and polarization–mode dispersion (PMD).

Modal dispersion

Modal dispersion arises only in multimode fiber where different modes travel at different velocities. Clearly, in single–mode fibers modal dispersion is not a problem.

Waveguide dispersion

Waveguide dispersion is caused because the propagation of different wavelengths depends on waveguide characteristics such as indices and shape of the fiber core and cladding. After entering a single–mode fiber, an information carrying light pulse is distributed between the core and the cladding. Its major portion travels within the core, the rest within the cladding. Both portions propagate at different velocities since the core and the cladding have different refractive indices.

Chromatic dispersion

Chromatic or material dispersion arises because different frequency components of a pulse (and also signals at different wavelengths) travel at different velocities due to the fact that the refractive index of the fiber is a function of the wavelength. It is typically measured in units of ps/nm·km, where ps refers to the time spread of the pulse, nm is the spectral width of the pulse, and km corresponds to the link length. Typically, standard single–mode fibers (SMFs) have a chromatic dispersion of 17 ps/nm·km at 1550 nm.

Recently, so–called nonzero dispersion shifted fibers (NZ–DSFs) are installed more often. By controlling the waveguide dispersion accordingly, NZ–DSFs have a chromatic dispersion between 1 and 8 ps/nm·km, or between -1 and -8 ps/nm·km at 1550 nm. For example, Alcatel's TeraLight Metro Fiber has a dispersion of 8 ps/nm·km. Such a low–dispersion fiber is targeted towards 10 Gb/s operation over $80-200$ km without requiring dispersion compensation. Another example is Corning's MetroCor Fiber. Its low negative dispersion enables the use of low–cost directly modulated DFB lasers. Both fibers are mainly devised for metro WDM networks in order to reduce network complexity and costs.

Polarization–mode dispersion

PMD arises because the fiber core is not perfectly circular, particularly in older installations. Thus, different polarizations of the signal travel at different velocities. PMD is proving to be a serious impediment in very high–speed systems operating at 10 Gb/s and beyond.

2.2.3 Nonlinearities

As long as the optical power within an optical fiber is small, the fiber can be treated as a linear medium, i.e., the loss and refractive index of the fiber are independent of the signal power. However, when the power levels get fairly high in the system the nonlinearities can place significant limitations on high–speed systems as well as WDM systems. Nonlinearities can be classified into two categories. The first set of effects occurs owing to the dependence of refractive index on the optical power. This category includes self–phase modulation (SPM), cross–phase modulation (CPM or XPM), and four–wave mixing (FWM). The second set of effects occurs owing to scattering effects in the fiber medium due to the interaction of light waves with phonons (molecular vibrations) in the silica medium. The two main effects in this category are stimulated Raman scattering (SRS) and stimulated Brillouin scattering (SBS).

Self–phase modulation

SPM is caused by variations in the power of an optical signal and results in variations in the phase of the signal. SPM leads to the spectral broadening of pulses. Instantaneous variations in a signal's phase caused by changes in the signal's intensity will result in instantaneous variations

of the frequency around the signal's central frequency. For very short pulses, the additional frequency components generated by SPM combined with the effects of material dispersion also lead to spreading or compression of the pulse in the time domain, affecting the maximum bit rate and the bit error rate (BER).

Cross–phase modulation

XPM is a shift in the phase of a signal caused by the change in intensity of a signal propagating at a different wavelength. XPM can lead to asymmetric spectral broadening, and combined with SPM and dispersion, may also affect the pulse shape in the time domain.

Four–wave mixing

FWM occurs when two wavelengths, operating at frequencies f_1 and f_2, respectively, mix to cause signals at frequencies such as $2f_1 - f_2$ and $2f_2 - f_1$. These extra signals can cause interference if they overlap with frequencies used for data transmission. Similarly, mixing can occur between combinations of three and more wavelengths.

Stimulated Raman scattering

SRS is caused by the interaction of light with molecular vibrations. Light incident on the molecules creates scattered light at a longer wavelength than that of the incident light. A portion of the light traveling at each frequency is downshifted across a region of lower frequencies. The light generated at the lower frequencies is called the Stokes wave. The fraction of power transferred to the Stokes wave grows rapidly as the power of the input signal is increased. In multiwavelength systems, the shorter–wavelength channels will lose some power to the higher–wavelength channels. To reduce the amount of loss, the power on each channel needs to be below a certain level.

Stimulated Brillouin scattering

SBS is similar to SRS, except that the frequency shift is caused by sound waves rather than molecular vibrations. Other characteristics of SBS are that the Stokes wave propagates in the opposite direction of the input light. The intensity of the scattered light is much greater in SBS than in SRS, but the frequency range of SBS is much lower than that of SRS. To counter the effects of SBS, one must ensure that the input power is below a certain threshold. Also, in multiwavelength systems, SBS may induce crosstalk between channels. Crosstalk occurs when two counterpropagating channels differ in frequency by the Brillouin shift, which is around 11 GHz for wavelengths at 1550 nm.

2.2.4 Crosstalk

Crosstalk decreases the signal–to–noise ratio (SNR) leading to an increased BER. Crosstalk may either be caused by signals on different wavelengths (interchannel crosstalk) or by signals on the same wavelength on another fiber (intrachannel crosstalk) due to imperfect transmission characteristics of components, e.g., AWG. Interchannel crosstalk must be considered when determining channel spacing. In some cases, interchannel crosstalk may be removed through the use of appropriate narrowband filters. Intrachannel crosstalk usually occurs in switching/routing nodes where multiple signals on the same wavelength are being switched/routed from different inputs

to different outputs. This form of crosstalk is more of a concern than interchannel crosstalk because intrachannel crosstalk can not be removed through filtering.

2.2.5 Noise

The SNR is deteriorated by different noise terms. In particular, we consider amplified spontaneous emission (ASE) of optical Erbium–doped fiber amplifiers (EDFAs), shot noise of photodetectors, and thermal noise of electrical amplifiers.

Amplified spontaneous emission

An optical EDFA amplifies an incoming light signal by means of stimulated emission. Besides stimulated emission also spontaneous emission takes place which has a deleterious effect on the system. The amplifier treats spontaneous emission radiation as another input signal and the spontaneous emission is amplified in addition to the incident light signal. The resulting ASE appears as noise at the output of the EDFA.

Shot noise

A photodetector converts the optical signal into an electrical photocurrent. The main complication in recovering the transmitted bit is that in addition to the photocurrent there is a shot noise current. Shot noise current occurs due to the random distribution of the electrons generated by the photodetection process even when the input light intensity is constant. (Note that the shot noise current is not added to the generated photocurrent but is merely a convenient representation of the variability in the generated photocurrent as a separate component.)

Thermal noise

Since the photocurrent is rather small it is subsequently amplified by an electrical amplifier. This electrical amplifier introduces an additional thermal noise current due to the random motion of electrons that is always present at typical temperatures.

Chapter 3

Related Work

Metropolitan are networks (MANs) are located between access and backbone networks, as depicted in Fig. 1.3 of Chapter 1. MANs have a number of *distinctive properties* which have to be taken into account in the design of metro network architectures and access protocols:

- The geographical coverage of MANs is limited. Typically, MANs have a diameter of 50 to 200 km.
- The number of nodes in a MAN is rather modest in the range from 10 to 200 nodes.
- Compared to backbone networks, MANs have to be more cost–effective due to the smaller number of subscribers and the traffic in MANs is more bursty.
- While the nature of data traffic in LANs [LTWW94] and WANs [PF95] has been investigated, defining traffic models for MANs is an open question at the time of writing.
- While mesh networks are quite common in the backbone, metro networks typically have a ring, bus, or star topology.

Metro WDM networks belong to the family of *multichannel* networks. Specifically, in WDM networks multiple parallel channels are created by transmitting/receiving data on different wavelengths where each wavelength forms a separate channel. In such multichannel WDM networks two types of data collisions can occur: (*i*) On each wavelength a *channel collision* takes place when two or more nodes send data on the same wavelength channel simultaneously, and (*ii*) a *receiver collision* occurs when the receiver of the intended destination node is not tuned to the wavelength on which the corresponding source node is transmitting. Receiver collisions are also known as destination conflicts. As we will see shortly, there is a large number of approaches to avoid or mitigate channel and/or receiver collisions, either on the architecture or access protocol level. For instance, channel collisions can be completely prevented on the architecture level by assigning each node its own *home channel* (wavelength) for transmission. Similarly, receiver collisions are completely avoided by equipping each node with one fixed–tuned receiver, each operating on a separate home channel (wavelength). In the home channel architecture the number of required wavelengths is identical to the number of nodes. Hence, such architectures do not scale well. Furthermore, assigning each node a *dedicated* wavelength results in a poor channel and transceiver utilization for bursty data traffic. The resource utilization can be improved by *sharing* each wavelength among multiple nodes. Nodes access the shared wavelength channels by deploying a *medium access control* (MAC) protocol which aims at either avoiding collisions or mitigating their impact on the network performance.

A large number of MAC protocols for metro WDM networks have been proposed and investigated in the literature. Many of those MAC protocols make use of well known single–channel media access techniques and adopt them to the high–speed multichannel WDM environment. For instance, slotted ALOHA and carrier sense multiple access (CSMA) are frequently used to arbitrate the wavelength access. For an extensive survey on slotted ALOHA, CSMA, and many other media access techniques the interested reader is referred to [vA94]. Most media access techniques used in MAC protocols for metro WDM networks belong to one of the following categories:

- *Preallocation*: In preallocation protocols resources (time slots, wavelengths) are statically assigned to sources and/or destinations, i.e., resources are dedicated to a given node and can not be used by other nodes, even though the resources might be idle.

- *Random access*: In random access protocols resources are not statically allocated. Nodes using this mechanism start transmitting data but a successful completion can not be guaranteed. Owing to the distributed environment, nodes transmit uncoordinated (ALOHA) or at best with limited coordination (CSMA). Therefore, packet transmission of different nodes may collide and be destroyed.

- *Reservation*: In reservation protocols resources are assigned dynamically on demand via pretransmission coordination. Each node ready to send a data packet disseminates control information prior to transmitting the corresponding data packet. Control packets are typically sent over a separate wavelength and are received by the intended destination node or other nodes as well. Through this signalling the source node informs the destination node (and other nodes) about the time and wavelength of the data transmission. More precisely, in so–called *tell–and–go* protocols a given node sends the data packet immediately after the control packet irrespective of the success of the control packet. On the other hand, in so–called *attempt–and–defer* protocols a given node sends the data packet only after having learnt that the corresponding control packet was successful. Here successful means that the control packet was sent collisionfree and that there are enough free resources for transmitting the corresponding data packet. The corresponding data packet is then sent on the reserved wavelength during the reserved time interval.

By applying the concept of home channel or preallocation *circuit switched* metro WDM networks can be realized, where channel resources are dedicated to a given node. On the other hand, random access and reservation MAC protocols are used for dynamically allocating channel resources to a given node on demand. Doing this on a per–packet basis leads to *packet switched* metro WDM networks.

In this chapter we review different metro WDM networks which have been presented in the literature to date. Typically, metro WDM networks have either a ring or star topology. In Section 3.1 we first describe several ring metro WDM networks. In Section 3.2 we discuss star metro WDM networks which are based on either a PSC or an AWG. In those star networks any pair of nodes is able to communicate in one *single hop*, i.e., no processing, storing, and forwarding at intermediate nodes is required. Such logical single–hop WDM networks embedded on a physical star topology play a major role in the subsequent chapters and they are therefore discussed at length. The design and evaluation of MAC protocols for single–hop WDM networks have attracted much attention. Section 3.3 provides a comprehensive survey on various MAC protocols for this type of network. Most of them are geared towards packet switching, where packets are either of fixed or variable size. As we will see shortly, in this context the term packet

is used in a more general sense and is not restricted to IP packets. In our survey we concentrate on the main ideas. Detailed information, e.g., assumptions on traffic generation and traffic matrices, can be found in the provided references. We highlight the merits and shortcomings of the various MAC protocols. Finally, we summarize the learnt lessons and formulate resulting guidelines for the design of MAC protocols for single–hop WDM networks. These guidelines will help us develop our MAC protocol in Chapter 5.

3.1 Ring metro WDM networks

Most of the ring metro WDM networks described below operate at a line rate of 2.5 Gb/s. For practical reasons most of these networks deploy fixed–tuned rather than tunable receivers.

3.1.1 KomNet

The KomNet metro WDM field trial network consists of three optical add–drop multiplexers (OADMs) interconnected in a bidirectional fiber ring topology [RBF+01][SLBR01]. The structure of an OADM is shown in detail in Fig. 3.1. On each fiber different wavelengths can be dropped by deploying tunable fiber Bragg gratings (FBGs). By using wavelength–insensitive combiners wavelengths can be added to each fiber. Each FBG has a relatively small insertion loss equal to 0.1 dB. The FBGs can be mechanically tuned within the millisecond range. Therefore, KomNet is well suited for (lambda) circuit switching, but is inefficient for packet switching due to the relatively large tuning time of each FBG.

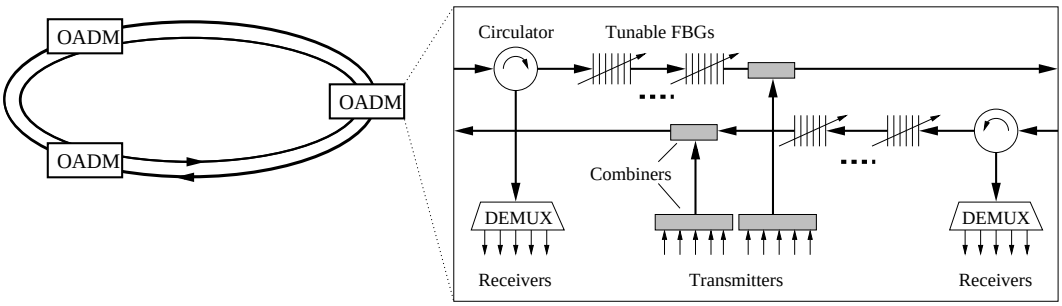

Figure 3.1: KomNet metro WDM network.

3.1.2 RINGO

The packet switched RINGO metro network is a unidirectional fiber ring network [GCF+01]. It comprises N nodes where N equals the number of wavelengths. Each node is equipped with an array of fixed–tuned transmitters and one fixed–tuned receiver operating on a given wavelength that identifies the node. Node j drops wavelength λ_j from the ring. Thus, in order to communicate with node j, a given node i has to transmit data by using the laser operating on wavelength λ_j, as illustrated in Fig. 3.2. All wavelengths are slotted with the slot length equal to the transmission time of a fixed–size data packet plus guard time. Each node checks the state of the wavelength occupation (λ–monitoring) on a slot–by–slot basis avoiding collisions by means of a multichannel generalization of the empty–slot approach (in the empty–slot approach one bit at the beginning of each slot indicates the state of the corresponding slot, i.e., whether

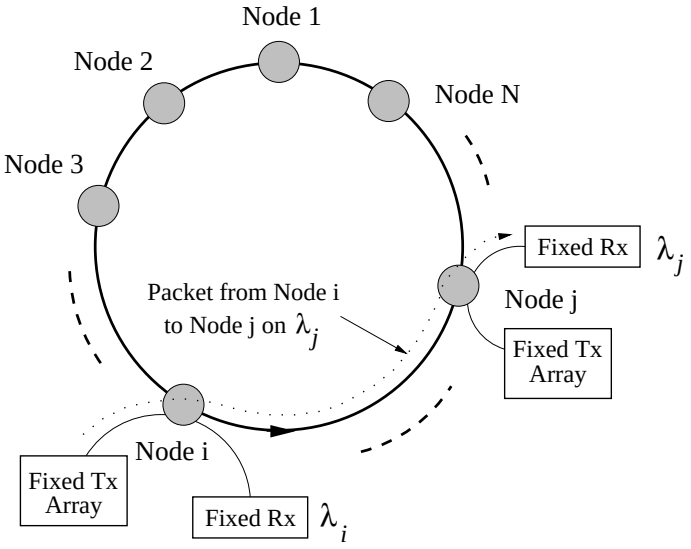

Figure 3.2: RINGO metro WDM network.

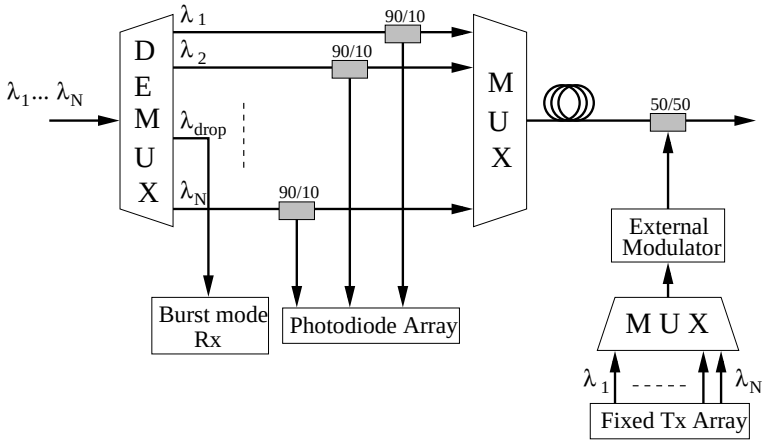

Figure 3.3: RINGO node structure.

the slot is free (empty) or occupied). This access mechanism gives priority to in–transit traffic by allowing a monitoring node to use only empty slots.

Fig. 3.3 depicts the node structure in greater detail. At each node all wavelengths are demultiplexed. The drop wavelength is routed to a burst mode receiver while the status of the remaining wavelengths is monitored by using 90/10 taps and an array of photodiodes. Subsequently, the wavelengths are multiplexed on the outgoing ring fiber. With a 50/50 combiner and an external modulator the corresponding node is able to send data packets by activating one or more fixed–tuned transmitters.

3.1.3 HORNET

HORNET is a unidirectional WDM ring network [GWW+99][SWW+00]. All wavelengths are slotted with the slot length equal to the transmission time of a *fixed–size* packet (plus guard time). Each wavelength is shared by several nodes for data reception. Every node is equipped

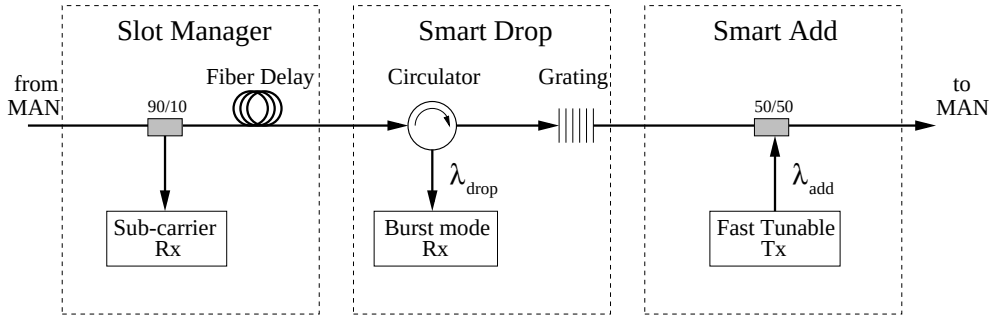

Figure 3.4: HORNET node structure.

with one fast tunable transmitter [FSA+00][SWR+01] and one fixed–tuned burst mode receiver [WWG+99]. The node structure consists of a slot manager, a smart drop, and a smart add module, as shown in Fig. 3.4.

Access to all wavelengths is governed by means of a carrier sense multiple access with collision avoidance (CSMA/CA) MAC protocol [WFS+00][WRS+00]. When a node transmits a packet it multiplexes a sub–carrier tone onto the packet at a unique sub–carrier frequency that corresponds to the wavelength the packet is sent on. Thus, all packets on the ring are carrying with them a sub–carrier multiplexed tone that denotes the wavelength which they are occupying. For carrier sensing the slot manager (see Fig. 3.4) simply taps off a small amount of optical power and detects it with one photodiode. As illustrated in Fig. 3.5, the data on all wavelengths collide at baseband while leaving the (ASK or FSK modulated) sub–carrier frequencies intact. The absence of a sub–carrier tone indicates the absence of a packet on the corresponding wavelength. This allows the node to determine whether a wavelength is free or not. If the wavelength of the corresponding destination node is idle the sensing node transmits the packet by deploying its smart add module.

Each module uses its smart drop module (see Fig. 3.4) to receive on its fixed assigned wavelength. The corresponding sub–carrier frequency is FSK modulated and carries the destination address of the respective packet. If the packet destination address does not match the node's address the node forwards the packet by using the smart add module. The CSMA/CA MAC protocol can be extended to support variable–size IP packets [SSW+00]. By adding a counterdirec-

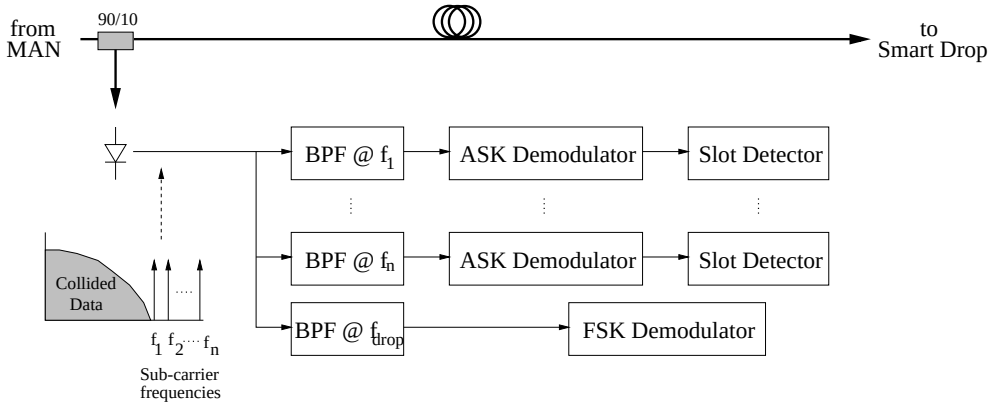

Figure 3.5: Structure of the slot manager.

tional fiber ring HORNET can be made tolerant against fiber/node failures [WSR+00][WRH+02].

3.1.4 IEEE 802.17 RPR

At the time of writing IEEE 802.17 and IETF WG IPoRPR are working on a new standard for metro ring networks. The resulting resilient packet ring (RPR) standard is scheduled to be completed by the end of 2003. In Section 10.2 we highlight its main principles when discussing how the work presented in this thesis can provide an evolutionary upgrade of RPR.

3.2 Star metro WDM networks

Star metro WDM networks are based on either a PSC or an AWG. In the following star networks communication between any arbitrary pair of nodes — either circuit or packet switched — takes place in one single hop, i.e., transmitted data does not have to be processed and forwarded by intermediate nodes.

3.2.1 RAINBOW

RAINBOW is an IBM initiated metro WDM network based on a PSC [DGL+90][JBM96]. As illustrated in Fig. 3.6, the network accomodates 32 nodes. Each node is equipped with one transmitter fixed tuned to a separate wavelength and one tunable receiver. Every receiver uses a tunable Fabry–Perot filter with a millisecond tuning speed. RAINBOW is geared toward full–duplex circuit switching. Circuits between nodes are set up and torn down by deploying a circular search protocol. In order to set up a connection a node continuously broadcasts a connection request message on its assigned wavelength. The intended destination, if idle, scans its tunable filter across all wavelengths looking for such a request and locks to a wavelength if it sees such a request. It then sends back a connection accept message that the originator looks for while itself scanning across all wavelengths. In RAINBOW I each node is able to send data at 300 Mb/s [JRS92][JRS93]. In RAINBOW II, on the other hand, the transmitting rate of each node is equal to 1 Gb/s [HKR+96].

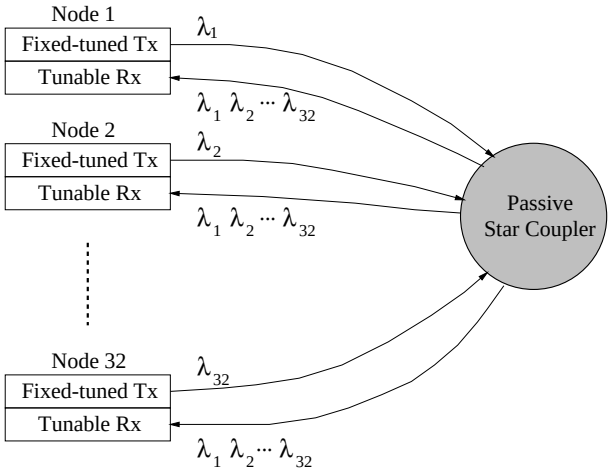

Figure 3.6: IBM's RAINBOW star metro network.

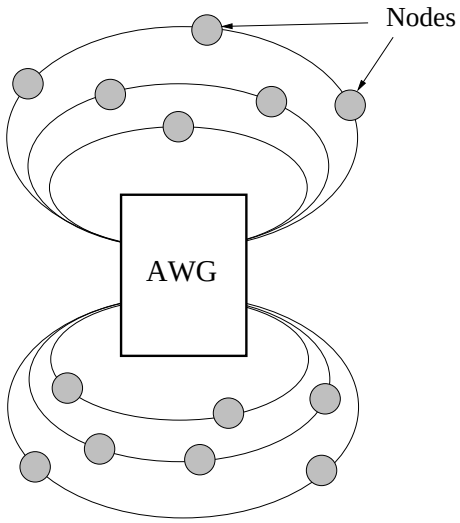

Figure 3.7: Telstra's AWG based metro network interconnecting multiple rings.

3.2.2 Telstra

Telstra deploys a central AWG (without any attached combiners and splitters) as a passive wavelength router for interconnecting WDM ring networks in a star topology, as illustrated in Fig. 3.7 [RA01]. Each node deploys fixed–tuned transceivers. By activating different transceivers, each node is able to send data to different ring networks through the wavelength–routing AWG. The main rationale behind this architecture is the fact that any two rings are directly connected by the central AWG. Thus, traffic does not have to pass multiple intermediate ring networks resulting in a reduced traffic load on each ring and an increased bandwidth efficiency.

3.2.3 NTT

NTT presented an AWG based star metro WDM network which interconnects 32 nodes [KOS+00]. As shown in Fig. 3.8, each node has 32 fixed–tuned transceivers. Each transceiver operates on a different wavelength such that any pair of nodes can communicate in one single hop at 10 Gb/s. At a network diameter of 20 km no optical amplifiers are required. The network can be extended up to 96 nodes leading to a network capacity of $96 \times 96 \times 10$ Gb/s = 92 Tb/s.

A modified network where each node is equipped with only two fixed–tuned transmitters (and n receivers, where n denotes the number of nodes) was described in [OSS+01]. Each node transmits data packets on the same 1.55 μm wavelength and the corresponding packet headers on the same 1.3 μm wavelength. To enable single–hop communication between any arbitrary pair of nodes, wavelengths have to be converted at the central AWG. Fig. 3.9 depicts a header and packet arriving from a given node to the central AWG. A WDM coupler routes the header to the optical header analyzer which determines the target wavelength on which the packet has to be transmitted in order to reach the corresponding destination node. The data packet is amplified and forwarded to the wavelength converter which consists of multiple light sources each operating at a different wavelength. By using semiconductor optical amplifiers (SOAs) gates and cross–gain modulation (XGM) the incoming packet is translated onto the target wavelength. After passing the multiplexer the AWG routes the packet according to the wavelength to the corresponding destination node.

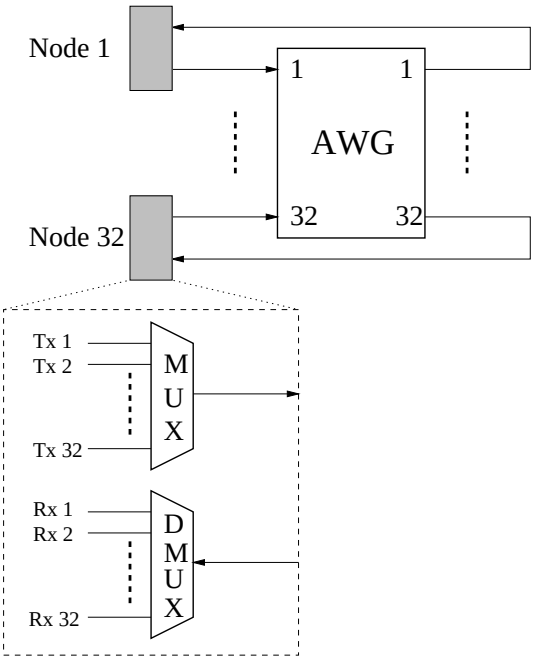

Figure 3.8: NTT's AWG based star metro WDM network.

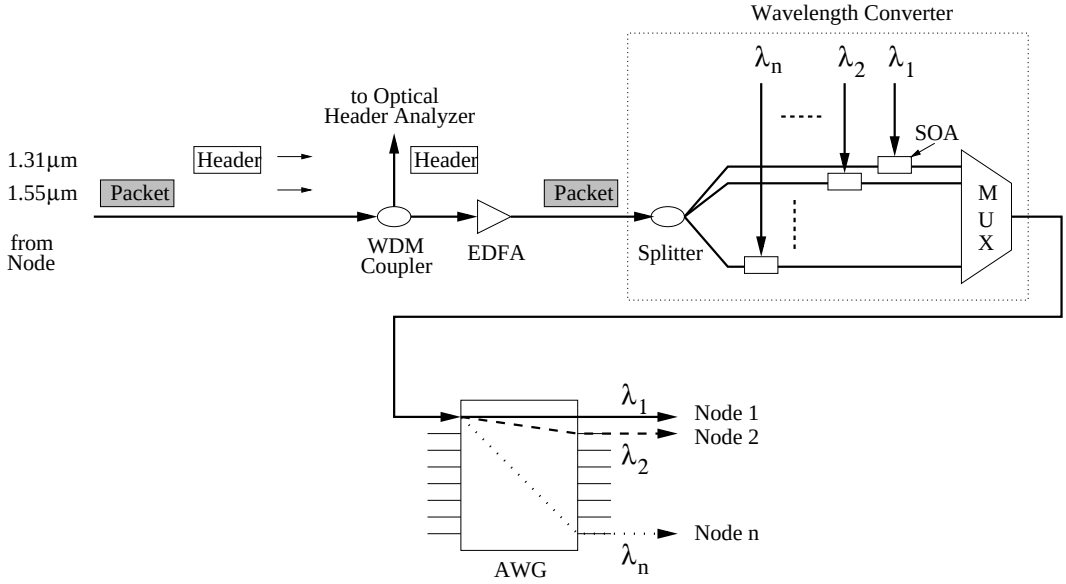

Figure 3.9: Structure of NTT's central AWG with wavelength converter.

3.3 Single–hop WDM networks

The concept of single–hop communication is not restricted to metro WDM networks. It has also been applied in WDM wide area networks (WANs), local area networks (LANs), switches, and high–speed interconnections of multiple processors and memories. Several single–hop WDM networks have been realized as testbeds. SONATA is a national–scale single–hop network based on a free–space AWG with attached arrays of wavelength converters and a centralized resource (time slots, wavelengths) controller [CHNS00][BLMN00a][BLMN00b][Hil00]. Examples for single–hop WDM LANs are Bellcore's LAMBDANET [GKL86][KBG+87][VBG+88] [GKV+90], Fairnet [BM93a], STARNET of Stanford University [KP93][CAM+96], SYMFONET [Kir90][Wes91], and [WC92]. Examples for single–hop WDM–based switches are provided by AT&T Bell Labs' photonic knockout switch [YHA87][EHY87][Eng88], Bellcore's FOX [ACG+88], HYPASS [AGKV88] and BHYPASS [Goo89], and a photonic ATM switch [CGMJ+93]. LIGHTNING represents an example for single–hop WDM multiprocessor communications systems [DPC+96].

For cost reasons each node in single–hop WDM networks deploys a rather small number of transceivers which is typically smaller than the number of wavelengths available for data transmission/reception. To increase the network efficiency all wavelengths should be used at any given time. This can be achieved if each wavelength is used by a different subset of nodes. However, if each node's transceiver(s) is (are) fixed–tuned each node can not listen to all wavelengths due to the limited number of transceivers per node. As a result, single–hop networks with fixed–tuned transceivers at each node are not able to provide full connectivity. This problem can be solved by either equipping each node with one fixed–tuned transceiver per wavelength or making each node forward packets towards the destination node resulting in multihop networks. Alternatively, single–hop networks with any–to–any connectivity can be realized if each node's transmitter(s) and/or receiver(s) are tunable. In doing so, each node has access to all wavelengths and is able to send and/or receive packets on any wavelength (within the transceiver tuning range). According to the different node structures, single–hop WDM networks can be categorized as follows:

- Fixed transmitter(s) and fixed receiver(s) (FT–FR)
- Fixed transmitter(s) and tunable receiver(s) (FT–TR)
- Tunable transmitter(s) and fixed receiver(s) (TT–FR)
- Tunable transmitter(s) and tunable receiver(s) (TT–TR)

Adopting the notation given in [Muk92], single–hop WDM networks can be described as FT^i–TT^j–FR^m–TR^n systems, where $i, j, m, n \in \mathbb{N}_0$ denote the number of the respective device(s) at each node.

Almost all single–hop WDM networks reported in the literature deploy a node structure with tunable transmitter(s) and/or receiver(s) rather than an array of fixed–tuned transceivers (one for each wavelength). In such networks packets can suffer from channel and receiver collisions. A channel collision occurs if two or more nodes transmit simultaneously on the same wavelength. A receiver collision occurs if a packet is sent collisionfree on a given wavelength to a destination node whose receiver is tuned to a different wavelength. Thus, the packet can not be received by the destination node, resulting in a receiver collision (also known as destination conflict). To mitigate or completely avoid channel and receiver collisions, the wavelength access has to be arbitrated by a MAC protocol. As shown in Fig. 3.10, MAC protocols for single–hop

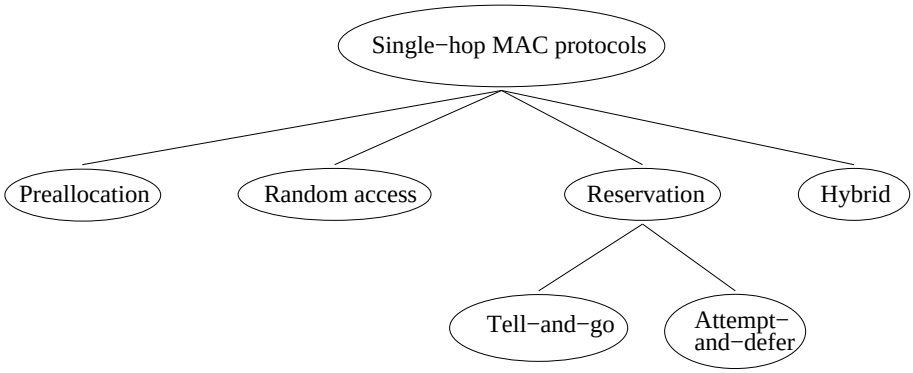

Figure 3.10: Classification of single–hop WDM network MAC protocols.

WDM networks can be classified into preallocation, random access, and reservation protocols, where the group of reservation protocols is further subdivided into so–called tell–and–go and attempt–and–defer protocols. In addition, there are MAC protocols which are hybrids of the aforementioned protocols. In the following sections we describe the basic operation and discuss examples of these different classes of MAC protocols.

The remainder of this chapter is organized as follows. First, we will highlight previous work on MAC protocols for single–hop WDM networks. For each MAC protocol category we distinguish between protocols with and protocols without receiver collisions. More detailed and extensive reviews of single–hop WDM network MAC protocols can be found in [Muk92][vA94][BCC98][MRW02]. In Section 3.3.5, we then summarize the learnt lessons and provide some resulting guidelines for the design of MAC protocols for single–hop WDM networks. In the following, let N denote the number of nodes and W denote the number of wavelengths in the network.

3.3.1 Preallocation protocols

Protocols with receiver collision

The source allocation protocol in [CG88b] is targeted towards bus and star networks. Each node is equipped with one transmitter and one receiver where either both (TT–TR) or one of them (TT–FR or FT–TR) must be tunable. This access protocol assigns transmission permits to source nodes in a fixed round–robin TDMA fashion where each cycle consists of $\lceil \frac{N}{W} \rceil$ slots of equal length. In each slot, W different source nodes are allowed to transmit to any destination, each source node sending on a separate wavelength. This protocol avoids channel collisions but receiver collisions can occur when more than one of the W nodes simultaneously send packets to the same destination. This situation arises more likely at medium to high traffic loads and for $W \approx N$. For $W \ll N$ and/or low to medium traffic loads receiver collisions happen less frequently since transmitted packets are unlikely to be addressed to the same destination. Note that due to the fact that each node can use its allocated slot for transmission to any one of the $(N-1)$ possible destination nodes wavelengths are utilized more efficiently. However, for medium to high loads of uniform and nonbursty traffic it is reasonable to assign each source node $(N-1)$ slots, one for each destination node, as discussed next.

Protocols without receiver collision

A modification of the above mentioned access protocol is the source/destination allocation protocol [CG88b]. In this protocol each cycle comprises $\left\lceil \frac{N(N-1)}{W} \right\rceil$ slots of equal size, i.e., the number of slots per cycle is increased by a factor of $(N-1)$ compared to the source allocation protocol. By allocating a separate slot to each source–destination pair in each cycle both channel and receiver collisions are prevented at the expense of a decreased channel utilization and a higher delay for bursty nonuniform traffic and/or low to medium loads.

A similar collisionfree round–robin TDMA protocol was presented in [BSD93]. The underlying network is based on a PSC and each node has a TT–FR structure. Wavelengths are fixed assigned to one or more different receivers. Thus, receivers can not be tuned to any other wavelength avoiding receiver collisions. Since receivers do not have to be tunable the implementational and operational complexity of the network is reduced.

An interesting approach to deploy fast tunable transceivers with a negligible tuning time and yet allowing for a large total number of available wavelengths was presented in [LK93]. The network is based on a PSC and each node has a TT–FR structure. Specifically, each node receives on multiple FSRs of its receiver, i.e., the multiwavelength–pass filter characteristics of each node's receiver is used, resulting in a larger number of available wavelengths. However, each node's transmitter has to be tunable over only one receiver FSR in order to enable full connectivity. Thus, this approach increases the number of channels while capitalizing on negligible transmitter tuning times. To avoid receiver collisions, wavelengths are preassigned in such a way that each node is allowed to receive data on only one FSR at any given time.

The problem of finding an optimal round–robin time–wavelength schedule for a given traffic matrix while taking the nonzero transceiver tuning latency into account was investigated in [GG92a][GG94][GG92b]. The network is based on a PSC with a TT^i–TR^j node structure, where $1 \leq i, j \leq W$. The proposed scheduling algorithms avoid both channel and receiver collisions. The obtained schedules consist of frames which are repeated periodically. Each frame comprises packet transmission periods and transceiver tuning periods. It has been shown that the problem of minimizing both periods is computationally intractable and therefore has to be split into two subproblems. The scheduling heuristic described in [GG92a] provides minimal transmission periods while trying to minimize the tuning periods. Another time–wavelength scheduling heuristic presented in [GG94] minimizes the tuning periods while aiming at providing minimal transmission periods. The latter heuristic yields better results than the former one if the transceiver tuning time is larger than the packet transmission time, and vice versa. The scheduling algorithm in [GG94] was also applied in a different environment where each node has a TT^2–TR^2 structure and is attached to two PSCs [GG92b].

A static round–robin TDMA/WDMA assignment scheme for real–time messages was developed in [THWH96]. The network is based on a PSC and the node structure is FT–TR, TT–FR, or TT–TR. The TDMA/WDMA schedule is collisionfree and provides deterministic timing guarantees for messages with given delivery deadlines while minimizing the number of required wavelengths.

The throughput–delay performance of fixed round–robin allocation schemes can be improved as shown in [CG90]. In this approach each node is fixed assigned one slot per cycle for every destination. Instead of using only the slot that allows for direct transmission, a given source node may use any other idle slot in which it is allowed to transmit packets. In doing so, a given source node sends a packet to another (intermediate) node which subsequently forwards the packet to the final destination node such that the overall delay is reduced, provided the

intermediate node has enough free resources. This approach can improve the throughput–delay performance, eliminates the large delay of preallocation protocols at low to medium traffic loads, and makes the network adaptive to varying traffic demands.

3.3.2 Random access protocols

Protocols with receiver collision

A completely allocation free protocol was studied in [CG88b]. The network is based on a PSC or bus. The node structure is TT–TR. All wavelengths are slotted with a slot length equal to the packet transmission time. Any node with a packet to send is allowed to transmit the packet on a randomly chosen wavelength at the beginning of each slot. Both channel and receiver collisions can occur, especially at medium to high loads.

Protocols without receiver collision

The destination allocation protocol in [CG88b] allows all nodes to send a packet at the beginning of each slot. The network is based on a PSC or bus and each node has a TT–FR or a TT–TR structure. While channel collisions can happen, receiver collisions are eliminated by assigning every wavelength to a different receiver in each slot. A cycle comprises several slots and is repeated periodically. In each cycle all nodes have the opportunity to send packets to any destination.

A similar random access protocol was described in [BSD93]. The network is based on a PSC. For implementational and operational simplicity each node is allocated a fixed home channel (wavelength) for reception while the transmitter is tunable (TT–FR), i.e., receiver collisions are avoided. All wavelengths are equally slotted and each node can transmit a packet at the beginning of each slot, possibly resulting in channel collisions.

The previous two protocols are identical to slotted ALOHA extended to a multichannel environment. In [Dow91] two variations of this multichannel slotted ALOHA were investigated which differ in the synchronization boundaries. The first protocol is slotted on minislot boundaries whereas in the second protocol each slot is longer and comprises L minislots which equals the packet length, where $L \geq 1$. As expected, the second random access protocol shows a higher throughput than the first protocol due to the reduced vulnerable period and thus decreased channel collision probability. More importantly, it was also shown that for a small number of wavelengths it is beneficial to use all wavelengths for data transmission than to put one wavelength aside for control and reservation as is typically done in reservation protocols (which are discussed in Section 3.3.3). The additional data channel becomes less beneficial as the number of wavelengths increases.

The performance of multichannel slotted ALOHA was analyzed for another PSC based single–hop network under nonhomogeneous traffic in [GK91]. Each node is equipped with one transmitter which can be tuned only to a subset of the W wavelengths and m receivers each fixed tuned to a different wavelength, where $1 \leq m \leq W$. The performance of multichannel slotted ALOHA is compared to that of another synchronous access protocol called random TDMA which randomly determines a collisionfree transmission schedule in a distributed fashion at the beginning of each slot. Specifically, each node deploys the same random number generator with the same seed. Until all wavelengths are assigned to source nodes, each node proceeds as follows. First, a given node, say node i, randomly selects one of the unallocated wavelengths, say λ_k. Second, node i randomly chooses one source node, say node j, among the nodes which are

able to transmit on wavelength λ_k. Node j randomly selects one packet among those stored in its buffer which are addressed to one of the destination nodes which are able to receive on wavelength λ_k. Last, source node j sends the packet (if any) to the corresponding destination node at the beginning of the next slot. For low system loads, multichannel slotted ALOHA outperforms random TDMA in terms of throughput and delay, and vice versa for medium to high traffic loads since random TDMA is free of channel collisions. As expected, the performance of both protocols is best when each node's transmitter can be tuned over all wavelengths and the number of receivers per node equals the number of wavelengths.

A single–hop network deploying two PSCs in parallel was presented in [Gla92][KG93a][KG94]. Each node is equipped with one tunable transmitter and one receiver fixed tuned to its dedicated home channel (TT–FR). Before transmitting a packet a given source node probes the home channel of the corresponding destination by sending a small bit burst. Only if this burst does not collide with other bursts and packets currently transmitted to the same destination, the source node gains access to the corresponding wavelength and starts sending the packet. Otherwise, the source node has to retransmit the burst at a later time.

A carrier sensing access protocol was reported in [CG88a]. The network has a unidirectional folded bus topology. Each node has a TT–FR structure. A given node's home channel for reception may be dedicated or shared by other nodes. Each node is allowed to transmit not only single packets but also packet trains (a packet train is created by sending multiple packets back to back) once the node has gained access to the corresponding home channel of the destination. Time is divided into periodically recurring cycles. A given node with a packet ready to send senses the corresponding home channel of the destination once per cycle. If the wavelength is idle the node starts transmitting the packet or packet train. If the wavelength is busy the node does not start transmission and keeps sensing the wavelength in the next cycles until the wavelength becomes idle.

3.3.3 Reservation protocols

Protocols with receiver collision

The single–hop star network considered in [HKS87] is based on a PSC. Each node has a TT–TR structure. Wavelengths are not fixed assigned. One wavelength is used for pretransmission coordination. Idle nodes tune their receivers to this control channel to collect control packets each consisting of the source address, destination address, and a randomly selected wavelength on which the respective data packet is intended to be sent. The paper examines several combinations of (pure and slotted) ALOHA and Carrier Sense Multiple Access (CSMA) for controlling the access to the control channel and the data wavelengths. A given node with a data packet in its buffer transmits the data packet immediately after sending the corresponding control packet, i.e., the protocol belongs to the tell–and–go category. Clearly, in this random reservation protocol both control and data packets can collide. Moreover, busy nodes are not able to monitor the control channel and might be tuned to another wavelength resulting in receiver collisions.

The protocols in [HKS87] provide a relatively poor throughput–delay performance due to the fact that data packets are sent irrespective of the success of the corresponding control packets. It was shown in [Meh90] and [SK91] that the throughput–delay performance can be improved by sending data packets only if the corresponding control packets were transmitted collisionfree. In other words, the throughput–delay performance can be improved by replacing the tell–and–go protocol with an attempt–and–defer one. However, receiver collisions can still occur and were not taken into account in the analysis.

Note that the aforementioned attempt–and–defer protocol suffers not only from receiver collisions but also from channel collisions of data packets. To see this, consider two subsequently transmitted control packets without channel collision which try to make a reservation for the same wavelength. Both control packets are successful and the corresponding data packets are transmitted accordingly. If data packets are longer than one control packet (which is typically the case) the two data packets collide resulting in wasted bandwidth and a decreased throughput. This problem was solved in [Lee91] by avoiding such channel collisions of data packets. Specifically, a given data packet is sent if (i) the corresponding control packet is transmitted without collision, and (ii) no other collisionfree control packet has made a reservation for the same wavelength during the last $(L-1)$ slots, where $L \geq 1$ denotes the data packet length and one slot equals the transmission time of a control packet.

The throughput–delay performance of the slotted ALOHA protocol combinations presented in [HKS87] and [Meh90] was improved by subdividing the time into periodically recurring cycles and allowing in each slot on the control channel reservations on a different data wavelength [SGK91]. In doing so, in each slot only nodes wishing to send a data packet on the same wavelength make their reservation resulting in a decreased number of collided control packets. This in turn translates into an improved throughput–delay performance due to the smaller number of collisions and retransmissions of control packets. Furthermore, the cyclic timing structure gives rise to Reservation ALOHA (R–ALOHA) [SGK91]. In R–ALOHA a given node which has successfully made a reservation in a slot is fixed assigned this slot until the given node has no more data packets to send. R–ALOHA yields high throughput and small delay values for medium to high traffic loads.

The slotted ALOHA protocols in [SGK91] can be further improved by sending control packets on multiple data wavelengths instead of one single control channel [SKG91]. In this protocol, there is no separate control channel and nodes are divided into groups. Each group is assigned a different wavelength. Time is divided into cycles which are repeated periodically. Each cycle is composed of a control phase and a data phase. An idle node tunes its receiver to the wavelength of the group to which it belongs. A source node wishing to transmit a data packet sends a control packet on the wavelength of the corresponding destination node during the control phase of a cycle by deploying slotted ALOHA. It was shown in [SKG91] that this protocol further improves the throughput–delay performance of the network. This is because control packets are distributed over multiple wavelengths (rather than one single control channel) and all wavelengths can be used for data transmission during the data phase resulting in a higher degree of concurrency and a decreased number of collisions and retransmissions of control packets.

In [HRS93] each node is assigned a dedicated control channel for receiving control packets. Thus, the network requires N control wavelengths in addition to the data wavelengths, where N denotes the number of nodes. Since each control channel is not shared by multiple nodes the number of collisions and retransmissions is decreased. This approach helps improve the throughput–delay performance but requires significantly more wavelengths than the reservation protocol in [SKG91].

The impact of receiver collisions on the throughput–delay performance was analytically investigated in [JM92a]. The authors considered the PSC based network with a TT–TR node structure and the reservation protocol in [Meh90] which was described above. It was shown that unless the population is large receiver collisions lead to a degeneration of the network throughput–delay behavior. This is because in networks with a small number of nodes two simultaneously transmitted data packets are likely to be addressed to the same destination node. With each node having one single receiver the given destination node is able to pick up only one

data packet while missing the other one. As a result, the control packet of the receiver collided data packet has to be retransmitted leading to a decreased throughput and an increased delay. The authors also found that control channel based reservation protocols with random access to both control and data wavelengths exhibit a bimodal and nonmonotonic delay for increasing traffic load if the number of data wavelengths is relatively small. To see this, note that at light loads almost all packets on the control and data channels succeed. As the load is increased channel collisions on the data wavelengths start to dominate and the throughput starts to drop and the corresponding control packets have to be retransmitted. As the load is further increased the control channel throughput drops as well, thereby reducing the load on the data wavelengths. Consequently, the data channel throughput again rises and fewer control packets have to be retransmitted. However, beyond a certain traffic load the collisions on the control channel occur too frequently such that the data channel throughput decreases again.

The so–called *dynamic time–wavelength division multiaccess* (DT–WDMA) reservation protocol that completely avoids channel collisions of both control and data packets was discussed in [CDR90]. The network is a PSC based star network. Each node has one transceiver fixed tuned to the control channel, one fixed–tuned transmitter, and one tunable receiver for data. Channel collisions on the control channel are avoided by means of round–robin TDMA. Data packets are sent collisionfree since each node transmits on a different dedicated wavelength. After sending the control packet a given node transmits the corresponding data packet on its assigned home channel in the next cycle without waiting for the outcome of the reservation. A control packet consists of destination address, age of the corresponding control packet, and a mode field. The mode is used to allow for both packet and circuit switching. In packet switching mode, a given control packet makes a reservation for one single data packet. In circuit switching mode a given control packet tries to set up a circuit. Provided the reservation is successful, the reserved wavelength and destination receiver are reserved as long as the given source node repeats sending the corresponding control packet in its allocated reservation slot. As a result, sessions are not interrupted and can benefit from guaranteed QoS. The age information in each control packet is used to give priority to data packets which are queued at the source node for a longer time period. Among multiple control packets destined to the same destination node that one with the largest age value succeeds in the reservation, thus ensuring fairness. The remaining control packets have to be retransmitted due to the resulting receiver collisions of the corresponding data packets. (We will get back to this reservation protocol in the benchmark comparison of our proposed access protocol in Section 6.4.4.)

Protocols without receiver collision

A reservation protocol which completely avoids both channel and receiver collisions of control and data packets was reported in [SGK92]. The network is based on a PSC and each node has a TT–FR structure with each receiver fixed tuned to its own home channel. Time is divided into recurring cycles. Each cycle is further subdivided into N control slots, M information slots, and M data slots, where $M \geq 1$. Each control slot is preassigned to a different node. When node i has to send a data packet to node j, node i transmits a control packet in the i–th control slot on node j's home channel. After receiving the control packets node j selects one or more source nodes according to a given policy. Using the information slots node j sends permits to the selected source node(s). Since the permit contains the transmission schedule there are no channel collisions of data packets.

Another completely collisionfree reservation protocol which is able to achieve a normalized

throughput of up to 100% was presented in [CY91]. Again, the network is based on a PSC. The node structure is FT^2–TR^2–FR, which is rather complex. Each node deploys one transceiver fixed tuned to a common control channel for sending control packets. For transmitting data packets each node uses its own home channel. By using two alternately operating tunable transceivers tuning penalties can be avoided. Time is divided into cycles which are repeated periodically. On the control channel each cycle comprises N reservation slots each preallocated to a different node. Every node disseminates its backlog information by broadcasting a control packet in its assigned reservation slot. The corresponding data packet is scheduled in a distributed fashion according to a given algorithm. Successfully scheduled data packets are then transmitted collisionfree on the respective home channel of the source node. A modification of this system with a centralized scheduler was described in [CY94] where the scheduler aims at minimizing the packet delay.

A contention based reservation protocol without preassigned reservation slots was investigated in [JU90]. The network is based on a PSC. Each node has a TT–FT–TR–FR structure where one transceiver is fixed tuned to a common control channel and the other one is tunable and is used for data transmission/reception. All wavelengths are slotted with a slot length equal to the transmission time of a (fixed–size) data packet. The slot on the control channel is subdivided into several reservation slots and one end–to–end propagation delay. The reservation slots are not fixed assigned. Nodes send control packets by using one of the reservation slots randomly (slotted ALOHA). At the end of each slot, i.e., after one end–to–end propagation delay, uncollided control packets take part in the distributed first–come–first–served (FCFS) scheduling procedure. If the reservation succeeds the corresponding data packets are transmitted during the next slot. Unsuccessful control packets have to be retransmitted. Here unsuccessful means that a control packet is collided on the slotted ALOHA channel and/or does not find enough free resources. By increasing the scheduling horizon control packets which have been sent collisionfree on the slotted ALOHA control channel are more likely to find free resources. Hence, fewer control packets have to be retransmitted translating into a decreased delay [JU92].

A more efficient modification of this protocol was proposed in [JU95] where the reservation slots are extended over the entire slot, i.e., there is no more idle end–to–end propagation delay phase on the control channel. In addition, the scheduling window is enlarged from one to several slots which allows data packets to be larger than one slot. Due to the larger scheduling window control packets are more likely to find free resources resulting in fewer retransmissions.

Note that in the previous protocol each node has to maintain a relatively large amount of status information. The involved processing requirements can be significantly reduced if nodes do not have to maintain any status tables [LU96]. This is achieved by replacing the slotted ALOHA with slotted reservation ALOHA (R–ALOHA). In doing so, a node sending a successful control packet in a given slot can exclusively use this slot for repeating control packets until the end of its data transmission. Thus, multiple updating control packets replace the need for large status tables at each node.

The basic approach in [JU90] can be extended by placing an additional priority field into each control packet [KSLU95]. This enables service differentiation, e.g., real and nonreal time traffic.

A reservation protocol avoiding both channel and receiver collisions of data packets and requiring only one transceiver per node was discussed in [JM93b]. The network is based on a PSC and each node has a TT–TR structure. The protocol takes nonzero transceiver tuning time and nonzero propagation delay into account. Time is divided into data slots with a slot length equal to the transmission time of a (fixed–size) data packet. The data slot on the common

control channel is further subdivided into W control slots each comprising several minislots. A node i, $1 \leq i \leq N$, wishing to send a data packet on wavelength j, $1 \leq j \leq W$, randomly selects a minislot in control slot j for broadcasting a control packet. By using slotted ALOHA rather than TDMA on the control channel new nodes can easily join the reservation process making the network scalable. All nodes which have participated in the reservation know its outcome after one end–to–end propagation delay. If there was no other uncollided control packet addressed to the same destination or reserving wavelength j before node i's uncollided control packet node i sends the corresponding data packet on wavelength j in the next control slot. Otherwise, node i has to retransmit the control packet. It was shown that bandwidth can be used more efficiently by overlapping one node's transceiver tuning time with another node's transmission time. A modified version of this protocol which incorporates nonuniform node distances from the central PSC was examined in [JM92b]. Another protocol modification replaces the W control slots with one [JM93a]. Wavelengths are allocated to successful control packets in an increasing order.

Another approach to avoid receiver collision is the use of switched fiber delay lines (FDLs) [CF91][CF94a][CF94b]. If two or more nodes simultaneously send a data packet to the same destination node data packets are put into the destination node's switched FDL and received sequentially. This contention resolving method smoothes the traffic burstiness and improves the throughput–delay performance of the network.

An adaptive control channel based protocol which reduces the number of receiver collisions was analyzed in [PM96]. The network is based on a PSC. Each node has a FT^2–TR–FR structure and a separate home channel for data transmission, thereby avoiding channel collisions. Each node stores its backlogged data packets in different buffers one for each separate destination. (This buffer structure is known as virtual output queueing (VOQ). Note that head of line (HOL) blocking can be avoided by equipping each node with multiple destination queues one for each destination [MGLA96]. Thus, a data packet with an occupied destination receiver does not prevent another data packet whose corresponding receiver is idle from being sent immediately resulting in an improved throughput–delay performance.) Every node randomly selects one of these packets according to a given probability distribution P. The destination address of the chosen data packet is broadcast by each node sending a control packet over the common control channel and subsequently transmitting the chosen data packet on its home channel. After one end–to–end propagation delay each node learns from the control traffic whether its data packet receiver collided or not. In case of receiver collisions the probability distribution P is changed such that receiver collided data packets are chosen with a smaller probability in the next reservation attempt. It was shown that this adaptive random transmission strategy provides a better network throughput–delay performance than static random transmission and first–in–first–out (FIFO) strategies. A similar adaptive approach also works for avoiding channel collisions in a PSC based network with TT–TR node structure [PM95]. In this case the status of all wavelengths is monitored by each node and is used to update P in that packets are more likely to be transmitted on collisionfree or idle wavelengths. Conversely, if on a given wavelength collisions occur packets are less likely to be sent on this wavelength.

Two reservation protocols with varying signalling complexity which avoid both channel and receiver collisions of data packets were reported in [CZA92][CZA93]. It is a PSC based FT^2–TR–FR system. Each node has one transceiver which is fixed tuned to the common control channel, one fixed–tuned transmitter for data transmission, and one tunable receiver for data reception. In the first proposed protocol every node executes an identical arbitration algorithm using a random number generator with the same seed. Thus, all nodes will arrive at the same conclusion. A transmitter i is randomly selected among all the transmitters. Among all nonempty receiver

queues (one for each destination) at transmitter i one queue r is randomly chosen. In the upcoming slot transmitter i sends a packet to receiver r. If all receiver queues are empty the slot remains unused. On the control channel the receiver queue status of all nodes is permanently broadcast using a fixed TDM scheme. Thus, each node has global knowledge for executing the common distributed arbitration algorithm. The algorithm procedure is repeated every data slot until all transmitters are served. At any step, transmitters and receivers which have already been scheduled are excluded from the arbitration algorithm. In doing so, fairness is provided while higher priorities are given to the queues with larger arrival rates by selecting only nonempty receiver queues. Note that for large populations the signalling traffic can become relatively high. To reduce the signalling overhead the second protocol is a combination of TDM and the first one, i.e., it sustains both preassigned and dynamic slot assignment. We will return to this protocol when discussing hybrid MAC protocols in Section 3.3.4.

The protocol described in [BD91][BD92][DB92] prevents both channel and receiver collisions of data packets by using channel and receiver status tables at each node. The network is based on a PSC and each node has a TT–TR–FR structure. On the common control channel each node is assigned one control slot in a static cyclic fashion. Control packets are composed of four fields: Source address, destination address, data wavelength, and packet size (packets can be of variable size). Each node maintains two status tables. The channel status table keeps track of the status of the wavelengths and is used to prevent channel collisions. The node status table at each node avoids receiver collisions by keeping track of the status of the tunable receivers at all nodes. The table entries indicate the number of slots wavelengths and receivers will be busy and are updated by using the control information.

A reservation protocol that requires no separate control channel was investigated in [SD95]. All wavelengths are divided into a control phase and a data phase which are repeated periodically. During the control phase all nodes use all wavelengths in a preassigned TDM fashion for making reservations. The corresponding data packets are transmitted in the data phase after one end–to–end propagation delay. Nodes have either a TT–FR or a TT–TR structure. In the first case the TT is realized by an array of fixed–tuned transmitters one for each wavelength. Broadcasting of control packets is achieved by activating all fixed–tuned transmitters simultaneously, thus using all wavelengths. Bandwidth can be used more efficiently by deploying a TR instead of a FR at each node [DS94]. In doing so, a given source node needs to send the control packet only on one wavelength to which all other nodes' receivers are tuned. Moreover, due to the higher flexibility of the TT–TR system compared to the TT–FR one, the wavelength utilization is increased while the mean delay is decreased [SD96]. This holds in networks where $N > W$, i.e., each wavelength is shared by several nodes. To see this, note that in the TT–FR system it might happen that a given node's home channel is busy while other wavelengths are not used. Since that node's receiver is fixed tuned it can not be tuned to one of the free wavelengths. As a result, the corresponding data packet can not be sent and the free wavelength remains unused. In contrast, in the TT–TR system both source and destination nodes can tune their transmitter and receiver to one of the free wavelengths and start transmitting the corresponding data packet. Especially for nonuniform traffic it was shown in [Sim98] that the TT–TR node structure is superior to its TT–FR counterpart in terms of throughput and delay. While in the TT–FR system home channels are underutilized for nonuniform traffic, the TT–TR system allows for load balancing over all wavelengths resulting in a higher channel utilization and an improved throughput–delay performance of the network. Furthermore, by allowing each node to make a reservation for more than one wavelength the network performance can be improved [SW96].

By deploying acoustooptic receivers with a relatively large tuning range more wavelengths can

be accessed by each node resulting in an improved throughput–delay performance of the network [KG93b]. Moreover, exploiting the multiwavelength filtering capability of acoustooptic receivers further improves the network throughput–delay performance due to the increased degree of concurrency [KGB93].

The amount of control traffic can be reduced by allowing each node to make a reservation for variable–size data packets rather than sending multiple control packets one for each segment of the original data packet [JMI95]. In this proposal each node has global knowledge about all nodes' reservations and invokes a distributed scheduling algorithm which aims at avoiding unnecessary transceiver tuning operations. Both control and data wavelengths are slotted where each control slot is preallocated to a different node. Since slots on the control channel and data wavelengths do not necessarily have to be synchronized control slots can be added or removed, thereby making the network scalable. Furthermore, the transceiver tuning time can be masked by tuning idle transmitters to the wavelength on which the corresponding destination node's receiver is currently operating and starting to send data as soon as the destination receiver becomes available [JMIO94]. Consequently, the destination receiver does not only have to be tuned to another wavelength but can continue receiving data without any interruption leading to an increased channel utilization.

A reservation protocol that not only utilizes WDMA and TDMA but also code division multiple access (CDMA) was analyzed in [MA95][AM95]. Several nodes share a wavelength by deploying different codes. Hence, the number of required wavelengths is reduced. This in turn allows for using transceivers with a small tuning range and a negligible tuning time resulting in a smaller tuning penalty, higher channel utilization, and an improved throughput–delay performance of the network.

A sampling and probing reservation protocol with in–band signalling was reported in [LGK95]. The network is based on a PSC and each has a TT–FR–TR structure. Each node has its own home channel for receiving data packets. A given source node monitors the home channel of the corresponding destination node by deploying its TR (sampling). If the home channel is found idle the source node sends a reservation request on the home channel (probing). If this request is successful the source node starts transmitting the respective data packet on the destination node's home channel.

The throughput–delay performance of reservation protocols can be improved by sequencing the data packets at each source node prior to broadcasting the control packets [HMH97][HMH99]. Sequencing means that data packets are processed at the source node such that they are not only sorted by their destination address but also by their length. The resulting data aggregates (each typically comprising several data packets) are then subsequently processed according to their length. Different sequencing policies are possible. Among others, a given source node starts making a reservation for the longest data aggregate and continues doing so for data aggregates in a decreasing length order. In this approach one control packet makes a reservation for several data packets improving the bandwidth efficiency and the throughput–delay performance of the network.

Service differentiation can be provided by placing the arriving packets into different queues according to their QoS requirements prior to making reservations [LQ98]. The queues contain either real–time or nonreal–time packets. Each source node tries to make reservations for each data packet with higher priority given to data packets stored in the real–time data buffer.

A reservation protocol which also supports real–time services was described in [YGK96]. It is a PSC based network using one wavelength as broadcast control channel. The control channel access is not done by TDMA or slotted ALOHA but by means of token passing. This access

technique does not require the control channel to be slotted and the nodes to be synchronized.

In [GYZ94] the authors have shown that transceiver tuning overhead in reservation protocols for single–hop WDM networks can be significantly reduced if the scheduling of data packets is not done on a per–packet basis but on the basis of a varied number of data packets called worm. This so–called *wormhole scheduling* decreases the number of required tuning operations translating into a smaller tuning penalty and an improved throughput–delay performance, in particular for large tuning times, e.g., mechanically tunable transceivers. (We will get back to wormhole scheduling when discussing feasibility issues of our proposed network in Chapter 8.)

Against the common conjecture that transceiver tuning times have to be very small relative to the packet transmission time the authors in [ABM96] have demonstrated that in bandwidth–limited networks, i.e., networks with $N > W$, there is no penalty in the traffic demand matrix clearance time through optimal scheduling as long as the transceiver tuning time is no larger than the transmission time of the data packet. As a consequence, through optimal scheduling one can eliminate the need for very rapidly tunable transceivers in packet switched single–hop WDM networks.

For any all–to–all (full connectivity) transmission schedules in single–hop PSC based networks the lower and upper bounds of their minimum length are provided in [PS94].

3.3.4 Hybrid protocols

The so–called Hybrid TDM (HTDM) MAC protocol combines TDM and on–demand reservation [CZA92][CZA93]. The control channel and all data wavelengths are slotted. One part of the slots are fixed allocated while the remaining slots are dynamically assigned by broadcasting reservation packets on the control channel. The protocol can be considered a tradeoff between flexibility and signalling overhead. The fixed assigned slots do not require any control traffic but suffer from underutilization for bursty nonuniform traffic. Conversely, the remaining slots which are dynamically assigned require signalling but are better suited for bursty nonuniform traffic.

A hybrid protocol which uses round–robin TDM for unicast traffic and reservation for high multicast traffic loads in a PSC based network with a FT^2–TR–FR node structure and a separate control channel was reported in [TK00]. More precisely, as long as the multicast session length and/or the multicast group size are small the corresponding multicast packet is sent as multiple unicast replicas using the round–robin TDM scheme. However, if the multicast session length and/or the multicast group size exceed a certain threshold the corresponding multicast packet is transmitted only once and it is received by all intended destination nodes. To achieve this, the corresponding source node broadcasts a control packet in order to reserve the corresponding receivers. In doing so, the transmitted multicast packet possibly preempts unicast transmissions since some multicast destination nodes can not listen to unicast traffic while receiving the multicast packet.

A hybrid MAC protocol which is adaptive not only to the traffic type but also to the traffic load is the so–called Hybrid Dynamic Reservation Protocol (HDRP) [CZ95]. This protocol is a combination of tell–and–go and attempt–and–defer reservation protocols. More specifically, while isochronous traffic is always transmitted in attempt–and–defer reservation mode the transmission mode of non–isochronous traffic depends on the load: At low load nodes deploy the tell–and–go reservation protocol whereas at high loads each node uses the attempt–and–defer protocol. The transmission mode is selected by each node according to the traffic load where the buffer occupancy serves as a good indication. The advantage of this hybrid MAC

protocol is that it avoids round–trip propagation delays at low loads and circumvents receiver collisions at high loads (note that receiver collisions are more likely to happen at high traffic loads). Furthermore, it was shown that the attempt–and–defer reservation protocol outperforms fixed–allocation schemes such as round–robin TDMA for nonuniform traffic, e.g., client–server traffic.

In the preceding sections we have reviewed previous work on the design and performance evaluation of MAC protocols for single–hop WDM networks. This review tried to highlight the most relevant contributions and insights provided in the literature. For the sake of completeness we would like to note that there are several other papers which address some more detailed issues of single–hop WDM MAC protocols. For instance, efficient scheduling algorithms accomodating arbitrary transceiver tuning times and traffic patterns were presented in [RA95a][RA95b][BM96][RS97][DS99]. Simple scheduling algorithms for an AWG based single–hop network were presented in [BJM99]. A single–hop network based on two PSCs was investigated in [HW94]. This architecture allows for spatial wavelength reuse resulting in a larger number of simultaneous transmissions. A multiwavelength (logical) single–hop network based on a (physical) ring network which is able to provide fairness control was discussed in [MBL+96].

3.3.5 Resulting guidelines

Let us first summarize the well known pros and cons of preallocation, random access, and reservation protocols in the following:

- Preallocation protocols completely avoid both channel and receiver collisions and do not require any signalling. They provide guaranteed bandwidth, delay, and jitter performance and are suitable for QoS requiring traffic. However, unless the traffic is uniform and nonbursty and the traffic load is medium to high the channel utilization is quite low. Due to the fixed slot allocation new nodes can be added to networks with preallocated resources only with service interruption. Moreover, all nodes have to be synchronized and can experience a certain access delay even though the wavelengths are idle.

- Random access protocols also work without signalling and are suitable for nonreal–time traffic at low to medium loads. As opposed to preallocation protocols nodes can access idle wavelengths immediately without any access delay. However, at medium to high traffic loads the throughput decreases and the mean delay is not bounded due to an increasing number of collisions.

- Reservation protocols are well suited for bursty traffic in highly flexible networks, e.g., networks where both transmitter and receiver are tunable. Resources are assigned on demand by means of reservation signalling. In tell–and–go reservation protocols data packets can suffer from both channel and receiver collisions but nodes do not have to wait for the outcome of the reservation. Tell–and–go protocols are well suited for networks with a large number of nodes working at low to medium traffic loads. Attempt–and–defer reservation protocols exhibit a minimum delay which is equal to the end–to–end propagation delay. However, successfully scheduled packets do not experience any channel and receiver collisions. This type of reservation protocol is well suited for QoS requiring traffic with a relatively large session length such that the reservation overhead is amortized.

Next, we draw several conclusions for the design of MAC protocols in single–hop WDM networks from the review provided in the previous sections. We note that despite the fact that

most of the aforementioned MAC protocols were devised for PSC based networks most of the following guidelines are valid for single–hop WDM networks in general. We will use the following guidelines when developing our MAC protocol for an AWG based single–hop WDM network in Chapter 5.

Guidelines for the design of MAC protocols in single–hop WDM networks:

- A TT–TR system is superior to its TT–FR counterpart since the TT–TR system (*i*) allows for efficient broadcasting (and multicasting) in that all receivers can be tuned to the source node's transmitting wavelength such that the broadcast (multicast) packet has to be sent only once [DS94], (*ii*) provides a higher flexibility in that each receiver can be tuned to any free wavelength resulting in an improved channel utilization and throughput–delay performance [SD96], and, especially for nonuniform traffic, (*iii*) allows for load balancing in that traffic between any source–destination pair can be distributed over all wavelengths, again improving the channel utilization and throughput–delay performance of the network [Sim98].

- The channel underutilization in preallocation protocols at low to medium traffic loads can be mitigated through a partially fixed wavelength assignment, resulting in an improved throughput–delay performance of the network [CG88b].

- Using a random access rather than fixed assignment protocol on the control channel makes reservation based networks scalable. Network propagation time independent random access protocols reduce the channel collision probability and thereby achieve larger throughputs [Lee91]. Due to the relatively large ratio of propagation delay and packet transmission time slotted ALOHA is superior to CSMA in terms of throughput–delay performance [KT75]. Reservation ALOHA (R–ALHOHA) achieves a larger throughput than conventional slotted ALOHA [SGK91]. In carrier sensing protocols the sensing entities can be centralized leading to a smaller end–to–end propagation delay and an increased throughput [KG94].

- Higher transmission concurrency increases the network efficiency [GG94]. This can be achieved by wavelength reuse in multilevel configurations [GG92b] or by exploiting the multiwavelength selectivity of acoustooptic transceivers [KGB93]. In random access networks a larger transmitter tuning range and more receivers at each node reduce the contention, resulting in an improved throughput–delay performance [GK91].

- There are several approaches to reduce or avoid the transceiver tuning latency. The tuning penalty can be decreased (*i*) by minimizing the number of tuning operations [GG92a][JMIO94], (*ii*) overlapping data transmission and tuning operation [SGK91], (*iii*) deploying two alternating transceivers at each node [CY91], or (*iv*) using multiple FSRs of a fixed–tuned Fabry Perot receiver where each transmitter is assigned a separate receiver FSR, thereby requiring only a small transmitter tuning range [LK93]. The number of required wavelengths can be decreased by multiplexing techniques such as CDMA [AM95][MA95] (alternatively subcarrier multiplexing (SCM) or polarization division multiplexing (PDM) are other candidates), which in turn allows for transceivers with a small tuning range and small tuning time.

- In single–hop networks where both transmitter and receiver are tunable it is reasonable to deploy reservation protocols [HKS87], even though this does not necessarily yield better performance [Dow91]. By sending control packets over more than one wavelength the

throughput–delay performance can be improved [SKG91]. Misdimensioned control channel based networks can exhibit bimodal throughput and nonmonotonic average packet delay [JM92a]. To avoid data channel and receiver collisions a node does not have to monitor the control channel continuously [JM92b][JM93b][JM93a]. A common distributed scheduling algorithm reduces the signalling overhead (no explicit ACKs are needed) and avoids collisions of data packets. Adaptive scheduling algorithms achieve a better throughput–delay performance, especially under varying traffic conditions [PM95][PM96]. Hybrid fixed and dynamic slot assignment is a good trade–off between flexibility and signalling overhead [CZA93]. Implicit wavelength allocation by monitoring the number of earlier successful control packets reduces the pretransmission coordination overhead [JU90]. Control slot preallocation and on demand data slot assignment achieve latency reduction at low loads and stable operation at high loads [SD95][SW96].

- Bandwidth can be saved and the network performance can be improved if data packets are sent only after successfully transmitted control packets [Meh90][SK91].
- Variable–size packets can be announced by placing an additional field into the control packet [BD91][BD92][DB92]. Variable–size packets without requiring status tables at each node can be supported by transmitting multiple control packets rather than one [LU96]. Supporting variable packet lengths reduces the signalling overhead [JMI95]. Additional control packet fields can be used to support multipriority traffic [KSLU95].
- The transceiver tuning overhead in reservation protocols can be significantly reduced if the scheduling of data packets is done on the basis of several packets rather than on a per–packet basis [GYZ94]. In addition, making a reservation for several data packets improves the throughput–delay performance of the network [HMH99].
- Receiver collisions are considered more destructive than channel collisions [BCC98]. In terms of throughput, channel collisions avoidance is superior to retransmissions [Lee91].
- The throughput–delay performance of a network can be improved by deploying switched fiber delay lines (FDLs) at the receivers for resolving destination conflicts [CF91].
- Buffer sharing among nodes reduces the number of required buffers and the mean packet delay [CY94].
- Using a switching mode field in the control packet allows for hybrid switching, i.e., both packet and circuit switching. Fair wavelength reservation can be realized by writing the age of the corresponding data packet into the broadcast control packet [CDR90].

Chapter 4

Architectural Comparisons

We have seen in the previous chapter that metro WDM networks typically have either a ring or star topology. In Section 10.2 we will discuss how both ring and star topologies can be combined in order to realize a hybrid network architecture. But for now we focus on star metro WDM networks. The benefits of a star configuration are numerous [BR99b]. Star configurations are easy to install, configure, manage, and troubleshoot. This has advantages in terms of installation, troubleshooting, and reconfiguration, reducing the cost of installation and ownership for the entire network. In addition, a star network based on a PSC or an AWG is reliable due to its passive nature. As opposed to ring (and bus) topologies it does not suffer from tapping loss which grows linearly with the number of nodes (in dB). However, PSC based star networks suffer from splitting loss (as opposed to the AWG which does not introduce splitting loss). But the splitting loss grows only logarithmically with the number of attached nodes (in dB). Note that star networks exhibit a single point of failure, i.e., when the central hub goes down the entire network connectivity is lost. Therefore, for survivability reasons the central hub has to be protected. This issue is addressed in Chapter 9.

In the previous chapter we have seen that star metro WDM networks are typically based on either a PSC or an AWG. In this chapter, we analytically investigate and compare different star metro WDM architecture candidates in order to gain some insight into their respective advantages and drawbacks. More precisely, in our comparisons we consider $N \geq 2$ nodes which are connected by either an AWG or a PSC. Moreover, we assume fixed–size packets and uniform unicast traffic. We will start with the AWG based star network and examine two different node structures in Section 4.1. A node comprises either tunable or fixed–tuned transceivers resulting in (logical) single–hop or multihop networks, respectively. After discussing their pros and cons with regard to network capacity and costs we concentrate on AWG based single–hop WDM networks. To assess their potential we investigate and compare AWG and PSC based single–hop star WDM networks in terms of throughput, delay, and packet loss in Section 4.2. The comparison results are summarized in Section 4.3. Based on these results we answer the question why we favor AWG based single–hop WDM networks for our metro network proposed in Chapter 5.

4.1 Single–hop vs. multihop AWG based networks

Depending on the node structure two types of *logical* network topologies can be embedded on the AWG based *physical* star topology: Single–hop and multihop networks. In this section, we compare both networks in terms of mean hop distance and capacity, i.e., the maximum

Transceiver Type	*Tuning Range*	*Tuning Time*
Electro–optic	10–15 nm	1–10 ns
Acousto–optic	∼100 nm	∼10 μs
Mechanically tunable	500 nm	1–10 ms

Table 4.1: Transceivers: Tuning ranges and tuning times.

achievable throughput [WMW02].

4.1.1 Architecture

We consider a completely passive physical star network based on an $N \times N$ AWG, where $N \geq 2$. The network connects N nodes each consisting of a transmitting part and a receiving part. The transmitting part of a given node is attached to one of the N AWG input ports while the receiving part of that node is located at the opposite AWG output port. For transmission and reception one FSR of the $N \times N$ AWG spanning wavelengths λ_1, $\lambda_2, \ldots, \lambda_N$ is deployed. On this AWG based physical star network two different logical topologies can be embedded : (1) Single–hop and (2) multihop networks. Both are discussed in more depth in the following.

Using one FSR the $N \times N$ AWG simultaneously accepts at each input port a total of N contiguous wavelengths λ_1 through λ_N and routes each wavelength to a different output port without resulting in channel collisions. At each output port arrive N wavelengths, one from each input port. Since wavelength λ_1 connects each node's own transmitting and receiving parts we omit this wavelength. The remaining $(N-1)$ wavelengths are necessary to send (receive) to (from) all destination (source) nodes in one single hop. Consequently, the resulting single–hop network can be achieved by equipping each node either (i) with an array of $(N-1)$ fixed–tuned transmitters and receivers, each operating on a different wavelength, or (ii) with a transmitter and a receiver which are both tunable over the aforementioned $(N-1)$ wavelengths. The option (i) may be considered a special case of a multihop network and is discussed shortly. Henceforth, we focus on a single–hop network in which each node deploys one or more tunable transceivers each consisting of a tunable transmitter and a tunable receiver. Specifically, each node is equipped with $r_S \geq 1$ tunable transceivers such that each node is able to simultaneously communicate with r_S different nodes, with each on a separate wavelength. Clearly, with a population of N nodes at most $(N-1)$ simultaneous transmissions/receptions per node are possible. Therefore, we let r_S be upper bounded such that $1 \leq r_S \leq (N-1)$. Recall from Section 2.1.4 that each transceiver tunes from one wavelength to another one in a nonzero tuning time. The incurred tuning latency largely depends on the type of transceiver in use. Table 4.1 shows typical tuning ranges and tuning times of electro–optic, acousto–optic, and mechanically tunable transceivers, respectively, which are used in our comparison [Muk97]. The required tuning range is determined by the number of used wavelengths $(N-1)$ and the channel spacing. Note that the tuning times of the various transceiver types differ by multiple orders of magnitude. Hence, for a given channel spacing the number of nodes N has a strong impact on the tuning latency.

In the multihop network tunable transceivers are not used. Instead, each node deploys $r_M \geq 1$ fixed–tuned transceivers, each tuned to a different wavelength. As mentioned in the previous paragraph, $r_M = (N-1)$ fixed–tuned transceivers at each node are sufficient to enable full connectivity among the N nodes in one single hop. However, in the more general case

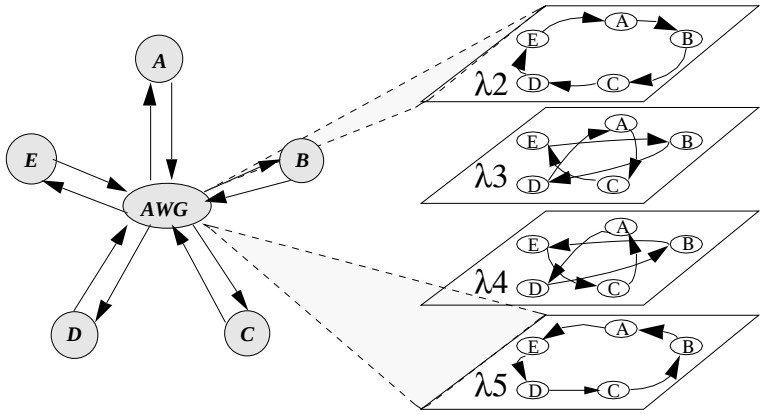

Figure 4.1: Virtual rings in an AWG based multihop network.

with $1 \leq r_M < (N-1)$ a given node can directly send data packets only to a subset of the $(N-1)$ possible destination nodes. As a consequence, in order to reach all destination nodes each node has to forward data packets towards the logical next–hop node until the data packets arrive at the final destination nodes, resulting in a logical multihop network. More specifically, each wavelength forms a virtual ring which interconnects a number of nodes as depicted in Fig. 4.1. In this example, $N = 5$ nodes (A through E) are connected by means of a 5×5 AWG. Recall that wavelength λ_1 is not used since this wavelength simply connects each node's own transmitting and receiving parts. The remaining four wavelengths form virtual rings such that the virtual rings on wavelengths λ_2 and λ_5 are counterdirectional, as are the virtual rings formed by wavelengths λ_3 and λ_4. All four wavelengths can be used for communication between any arbitrary pair of nodes.

In general, however, there may be wavelengths which connect only a subset of the N nodes. This is illustrated for a 4×4 AWG in Fig. 4.2, where only λ_2 and λ_4 interconnect all nodes while λ_3 forms two separate node–disjoint rings (A–C and B–D) leading to logically disjoint subnetworks. It was shown in [Woe98] that these subnetworks exist in all AWG based multihop networks where N is not a prime number. Only if N is a prime number, then the multihop network consists of $(N-1)$ virtual rings each interconnecting all N nodes. Deploying r_M transceivers, each tuned to a different wavelength, every node is able to communicate with r_M nodes simultaneously. The question to which wavelengths those r_M transceivers have to be

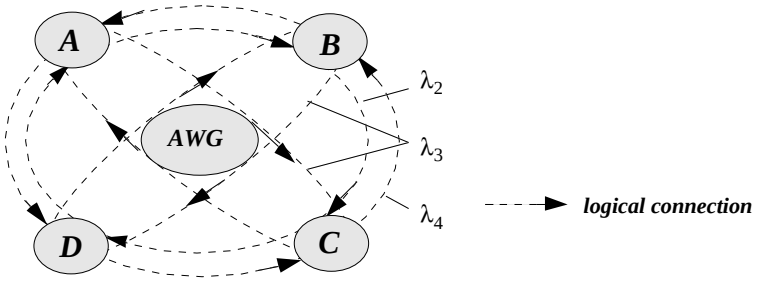

Figure 4.2: Logical topology of a 4×4 AWG based multihop network.

tuned is addressed in the following section.

4.1.2 Mean hop distance

For the calculation of the mean hop distance of both single–hop and multihop networks let one hop denote the distance between two logically adjacent nodes. Moreover, let the mean hop distance be equal to the average value of the minimum numbers of hops a data packet has to traverse on its shortest path from a given source node to all remaining $(N-1)$ destination nodes. Due to the symmetry in both single–hop and multihop networks the mean hop distance is the same for all (source) nodes. In our analysis we make the following assumptions:

- *Uniform unicast* traffic: A given source node sends a data packet to any of the remaining $(N-1)$ destination nodes with equal probability $1/(N-1)$.
- *Wavelength continuity*: Packets have to arrive at and depart from a given forwarding intermediate node on the same wavelength, i.e., nodes cannot perform wavelength conversion.

Clearly, in the single–hop network each source node reaches any arbitrary destination node in one hop. Thus, the mean hop distance of the single–hop network is given by

$$\overline{h}_S = 1. \tag{4.1}$$

The computation of the mean hop distance of the multihop network is more involved. For illustration, let us begin with the simple case $r_M = 1$ where we arbitrarily select one wavelength such that we obtain a unidirectional ring connecting all N nodes. The mean hop distance is then given by

$$\overline{h}_M = \frac{1}{N-1}\sum_{i=1}^{N-1} i = \frac{N(N-1)}{2(N-1)} = \frac{N}{2}, \tag{4.2}$$

where the distance between a given source node and the other $(N-1)$ nodes is equal to $1, 2, \ldots, (N-1)$, respectively. Next, let each node deploy another wavelength, i.e., $r_M = 2$, resulting in another virtual ring. The additional wavelength has to be chosen such that the second virtual ring decreases the mean hop distance as much as possible. To achieve this the second ring has to be counterdirectional to the first ring already in use. In doing so, for odd N we have to traverse at most $(N-1)/2$ hops in each direction. The resulting mean hop distance is then given by

$$\overline{h}_M = \frac{2}{N-1}\sum_{i=1}^{(N-1)/2} i = \frac{2(\frac{N-1}{2}+1)\frac{N-1}{2}}{2(N-1)} = \frac{N+1}{4}, \tag{4.3}$$

where a given source node reaches two different destination nodes for each hop count. For $r_M \geq 3$ the choice of the wavelengths to be added is nontrivial. This problem always arises when designing a logical multihop WDM network. It is sometimes referred to as the node placement problem or the wavelength assignment problem. As it turns out in our case, it is not always the best approach to select counterdirectional virtual rings in order to minimize the mean hop distance. For example, in the case $N = 13$ and $r_M = 4$ the combination of any two pairs of counterdirectional rings such as $\{\lambda_2, \lambda_5, \lambda_{10}, \lambda_{13}\}$ leads to a mean hop distance $\overline{h}_M = 7/3 = 2.3\overline{3}$ while using the noncounterdirectional rings $\{\lambda_2, \lambda_5, \lambda_7, \lambda_{12}\}$ conveys a smaller mean hop distance $\overline{h}_M = 9/4 = 2.25$. Instead of calculating the exact mean hop distance of the

Number of nodes N / *Number of rings r_M*	3	5	7	11	13
1	1.5	2.5	3.5	5.5	6.5
2	1.0	1.5	2.0	3.0	3.5
3		1.25	1.67	2.5	2.92
4		1.0	1.33	2.0	2.25
5			1.17	1.5	1.75
6			1.0	1.4	1.5
7				1.3	1.42
8				1.2	1.33
9				1.1	1.25
10				1.0	1.17
11					1.08
12					1.0

Table 4.2: Optimum mean hop distances of multihop networks with different number of nodes N and different number of fixed–tuned transceivers r_M per node.

AWG based multihop networks with arbitrary N and $1 \leq r_M \leq (N-1)$, we provide its lower bound which is given by

$$\overline{h}_M \geq \sum_{i=1}^{\left\lfloor \frac{N-1}{r_M} \right\rfloor} \frac{r_M}{N-1} \cdot i + \frac{(N-1) \mathrm{mod} r_M}{N-1} \cdot \left\lceil \frac{N-1}{r_M} \right\rceil \tag{4.4}$$

$$= \frac{1}{N-1} \left\{ r_M \cdot \frac{\left\lfloor \frac{N-1}{r_M} \right\rfloor \left(\left\lfloor \frac{N-1}{r_M} \right\rfloor + 1 \right)}{2} + [(N-1)) \bmod r_M] \cdot \left\lceil \frac{N-1}{r_M} \right\rceil \right\}. \tag{4.5}$$

To see this, note that the mean hop distance is minimized if (i) as many different nodes as possible are reached for each hop count starting with one hop, and (ii) the maximum hop distance (diameter) of the network is minimum. Applying this leads us to Eqn.(4.4). Since a given source node sends on r_M wavelengths at most r_M different destination nodes can be reached for each hop count. Each time exactly r_M different destination nodes are reached up to a hop count of $\left\lfloor \frac{N-1}{r_M} \right\rfloor$, which corresponds to the first term of Eqn.(4.4). The second term in Eqn.(4.4) accounts for the remaining (less than r_M) nodes which are $\left\lceil \frac{N-1}{r_M} \right\rceil$ hops away from the given source node. To see whether this lower bound is tight we compare it with optimum mean hop distances obtained through exhaustive search for multihop networks with prime N. The solutions are optimum in the sense that they provide the smallest mean hop distance of all possible ring combinations for a given r_M and N, respectively. Table 4.2 shows the optimum mean hop distances of multihop networks with a prime number of nodes up to $N = 13$ for different r_M. As illustrated in Fig.4.3, the lower bound is tight for $N = 11$ and varying r_M. Apparently, increasing r_M, i.e., adding fixed–tuned transceivers to each node decreases the mean hop distance. The minimum mean hop distance equals one and is achieved for $r_M = (N-1) = 10$. For the other values of N presented in Table 4.2 we observed that the lower bound is tight as well, with the optimum mean hop distance differing from the lower bound by at most 16.8%.

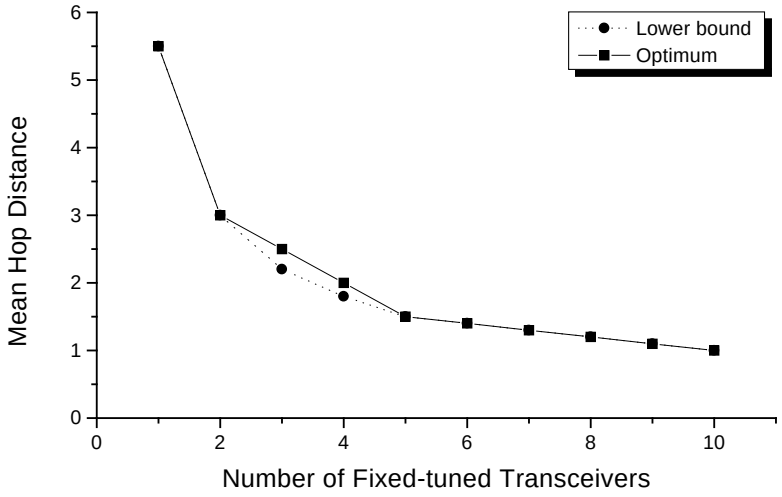

Figure 4.3: Lower bound of mean hop distance of a multihop network with $N = 11$ vs. r_M.

4.1.3 Capacity

According to [AS91] the capacity C of a WDM network – either single–hop or multihop – is defined as

$$C = \frac{r \cdot S \cdot N}{\overline{h}} \tag{4.6}$$

where r denotes the number of transceivers per node each operating on a separate wavelength, S stands for the data rate of each transceiver, N represents the number of nodes in the network, and $\overline{h}$ denotes the mean hop distance of the network. We assume that fixed–size data packets are sent/received.

Recall that in the single–hop network every node is equipped with r_S tunable transceivers, where $1 \leq r_S \leq (N-1)$. For the single–hop network we assume that the transceivers have to be tuned to another wavelength each data packet transmission. Thus, our capacity evaluation is rather conservative in the sense that in general each transceiver can be tuned to the same wavelength while transmitting more than one packet back to back. Let τ denote the nonzero transceiver tuning time which effectively decreases the data rate of each transceiver. The resulting net data rate of each transceiver is then given by

$$S_S = \frac{L}{L+\tau} \cdot S \tag{4.7}$$

$$= \frac{1}{1+\tau_L} \cdot S \quad , \tau_L = \frac{\tau}{L} \tag{4.8}$$

where τ_L denotes the transceiver tuning time normalized by the packet transmission time L. With the mean hop distance $\overline{h}_S = 1$ the capacity of the single–hop network is

$$C_S = \frac{r_S \cdot S \cdot N}{1+\tau_L}. \tag{4.9}$$

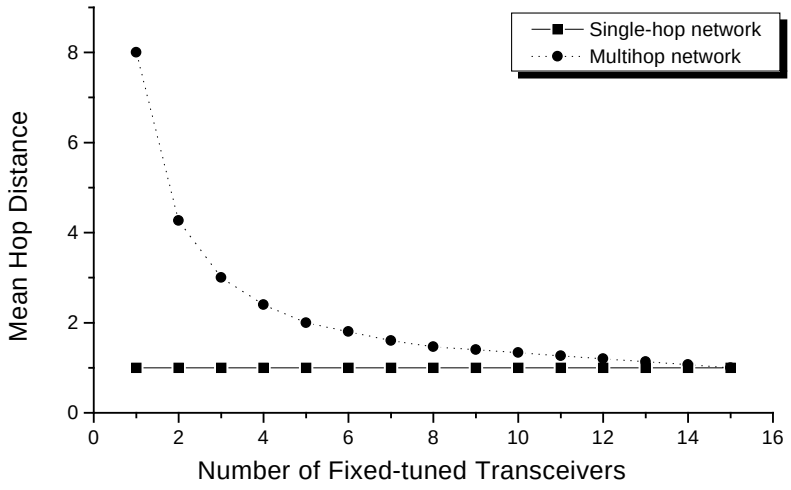

Figure 4.4: Mean hop distance vs. r_M for $N = 16$.

In the multihop network every node has r_M transceivers, where $1 \leq r_M \leq (N-1)$. Since the transceivers are fixed tuned there is no tuning penalty. Consequently, the effective data rate equals S. Using the lower bound of the mean hop distance given in Eqn.(4.5) provides the following upper bound on the capacity of the multihop network

$$C_M \leq \frac{r_M \cdot S \cdot N}{\overline{h}_M} \tag{4.10}$$

$$= \frac{r_M \cdot S \cdot N \cdot (N-1)}{r_M \cdot \frac{\left\lfloor \frac{N-1}{r_M} \right\rfloor \left(\left\lfloor \frac{N-1}{r_M} \right\rfloor + 1 \right)}{2} + [(N-1) \bmod r_M] \cdot \left\lceil \frac{N-1}{r_M} \right\rceil}. \tag{4.11}$$

4.1.4 Results

In the subsequent results we consider packets with a fixed length of 1500 bytes transmitted/received at a data rate of 10 Gbps. This translates into a packet transmission time $L = 1.2$ μs. The channel spacing is assumed to be 100 GHz (0.8 nm at 1.55 μm).

First, we take a look at single–hop networks in which each node has fast tunable electro–optic transceivers with a tuning time of 10 ns. This results in a normalized transceiver tuning time $\tau_L = 8.3\overline{3} \cdot 10^{-3}$. With a channel spacing of 0.8 nm and a limited transceiver tuning range of 10–15 nm (see Table 4.1) the number of available wavelengths is approximately 16. This allows for realizing single–hop and also multihop networks with a population size of up to $N = 16$ nodes. Fig. 4.4 depicts the mean hop distance of both networks as a function of the number of the fixed–tuned transceivers r_M used at each node of the multihop network. Clearly, the mean hop distance of the single–hop network is one, independent of r_M, and serves as a reference for the mean hop distance of the multihop counterpart. For the multihop network, Fig. 4.4 shows the lower bound of the mean hop distance given in Eqn.(4.5). We observe that the lower bound

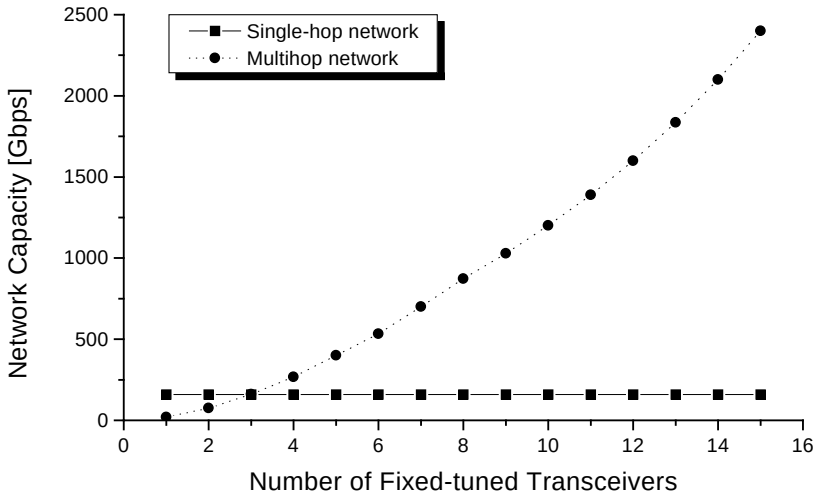

Figure 4.5: Network capacity vs. r_M for $N = 16$ ($r_S = 1$, fixed).

decreases rapidly with increasing r_M. Consequently, a few fixed–tuned transceivers at each node are sufficient to decrease the mean hop distance of the multihop network dramatically and to get close to the mean hop distance of the single–hop network. Adding further transceivers to each node has only a small impact on the resulting mean hop distance of the multihop network. Note that for $r_M = (N - 1) = 15$ both single–hop and multihop networks have the same mean hop distance, namely one.

However, from the network capacity point of view it is very beneficial to equip each node of the multihop network with as many fixed–tuned transceivers as possible. Fig. 4.5 illustrates the upper capacity bound of the multihop network, given in Eqn.(4.11), as a function of r_M. The capacity of the multihop network grows quadratically for increasing r_M. This is due to the fact that a large r_M not only decreases the mean hop distance but also increases the network degree of concurrency by simultaneously using all r_M transceivers at each node. For comparison Fig. 4.5 shows also the capacity of the single–hop network with $r_S = 1$ transceiver at each node given in Eqn.(4.9). Note that the multihop network requires at least four fixed–tuned transceivers per node in order to outperform its single–hop counterpart with one single tunable transceiver per node in terms of capacity. This point is discussed in more detail in Section 4.1.5.

Recall from Section 4.1.1 that for a given channel spacing the number of nodes N determines the required tuning range of the tunable transceivers used in the single–hop network. As we have seen, for a channel spacing of 0.8 nm and a population size of up to $N = 16$ nodes we can use electro–optic transceivers with negligible tuning time. However, for larger populations we have to use different types of transceivers with a wider tuning range at the expense of a significantly larger tuning time (see Table 4.1). Thus, for $N > 16$ acousto–optic transceivers have to be used which exhibit a three orders of magnitude larger tuning time than their electro–optic counterparts, resulting in a normalized tuning time $\tau_L = 8.3\bar{3}$. The impact of the different transceiver tuning times on the single–hop network capacity is shown in Fig. 4.6 for $2 \leq N \leq 32$

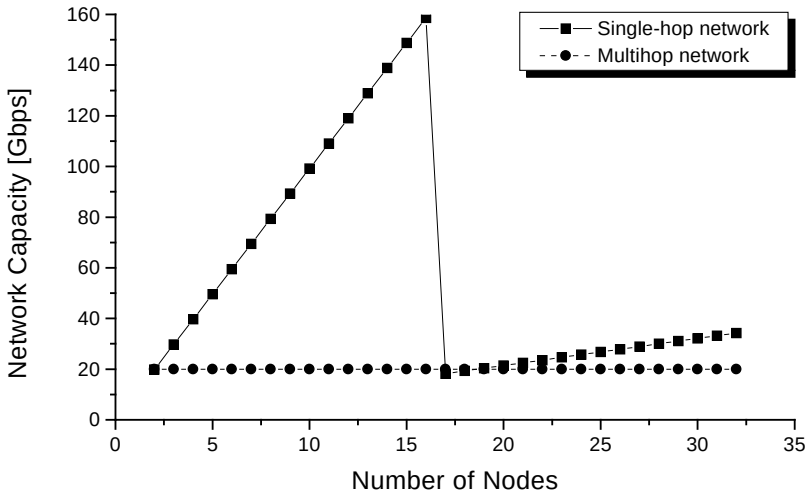

Figure 4.6: Network capacity vs. N ($r_S = 1$, $r_M = 1$ fixed).

and $r_S = 1$. We observe that the network capacity of the single–hop network grows linearly with N for $N \leq 16$. This is because each additional node contributes equally to the aggregate capacity. For $N > 16$ acousto–optic transceivers with a wider tuning range have to be used. The incurred larger transceiver tuning latency dramatically decreases the network capacity. Note that for $N > 16$ the network capacity grows again linearly with N but with a smaller slope.

In addition, Fig. 4.6 depicts the upper capacity bound of the multihop network with $r_M = 1$ fixed–tuned transceiver at each node. Interestingly, this bound remains constant for varying N. This is because with increasing N more nodes add to the network capacity but each node has to forward more packets due to the increased mean hop distance which in turn results in a decreased channel utilization and a lower net data rate per node. This point is reflected in Eqn.(4.11), in which the lower mean hop distance bound of the multihop network is directly proportional to N. As a consequence, the upper capacity bound $C_M = 2 \cdot S = 20$ Gbps is independent of N. Similar observations can also be made for $r_M > 1$.

4.1.5 Discussion

We have investigated logical single–hop and multihop network topologies based on an AWG in terms of mean hop distance and capacity. The single–hop network provides a minimum mean hop distance equal to one but requires tunable transceivers at each node. In contrast, transceivers in the multihop network do not have to be tunable but the mean hop distance is generally larger than one due to the longer hop distances encountered on the above described virtual rings. While equipping each node with additional (fixed–tuned) transceivers decreases the mean hop distance of the multihop network, the mean hop distance of the single–hop network is one independent of the number of (tunable) transceivers per node. However, from the capacity point of view it is beneficial to add transceivers to each node since this increases the number of simultaneous transmissions in both single–hop and multihop networks. The transceiver tuning

penalty significantly impacts the channel utilization. We have seen that the capacity of the single–hop network largely depends on the transceiver type in use which in turn is determined by the number of network nodes for a given channel spacing.

In the remainder of this work we will focus on AWG based single–hop networks. Single–hop networks have several desirable features [Muk92][Ger96]:

- *Transparency*: As opposed to multihop networks, packets in single–hop networks do not have to traverse intermediate nodes along the path to their destinations thus avoiding any optical–electrical–optical (OEO) conversion. Since no electrical processing is involved, single–hop networks are not aware of the structure of the data, and can inherently carry diverse protocols and bit coding schemes. Furthermore, they are able to carry quite a large range of bit rates. Conversely, OEO solutions carry a single form of traffic and require costly conversion devices from other protocols to the supported standard which also complicates the management of the network.

- *Simplified management*: The AWG based single–hop network is completely passive and bits/packets are interpreted only at the border of the network. All network intelligence is moved toward the network periphery simplifying the network management. Single–hop networks are more tolerant to node failures than their multihop counterparts since nodes are not involved in packet forwarding.

- *Future–proofness*: As a consequence of the above, single–hop networks are able to carry most future protocols at many different bit rates without having to replace components of the network. Future services such as rent–a–wavelength are easily feasible. For upgrading the bit rate between a given pair of nodes only the source and destination nodes have to apply the technically advanced devices. Unlike in the single–hop network, all intermediate nodes in the multihop network would have to be upgraded as well.

- *Reduced processing*: In single–hop networks no intermediate nodes are bothered by storing and forwarding packets. Each node has to process only packets which are addressed to itself.

- *Improved throughput–delay performance*: In single–hop networks no bandwidth is wasted due to packet forwarding resulting in a higher channel utilization and throughput. In general, single–hop networks provide a smaller packet delay than multihop networks. In multihop networks the delay may be long since a packet transmission between two nodes may be possible only through multiple hops, each time passing the AWG again. This implies longer routes and thus larger propagation delays, which becomes the dominating delay component in high–speed networks.

The downside of single–hop networks is the limited geographical coverage, the possibly complex transmission coordination between tunable transceivers, and the relatively high system cost because of the expensive tunable transceivers.

Single–hop networks are promissing candidates for local and metropolitan area networks (LANs, MANs), but represent no wide area network (WAN) solutions. Long–haul networks are typically multihop networks using the OEO conversion at intermediate nodes for electrical signal regeneration, packet processing, and buffering. In single–hop networks communication between a given pair of nodes can only take place if the source transmitter and the destination receiver are tuned to the same wavelength. In the previous chapter we have reviewed a wide variety of

MAC protocols which tackle this transmission coordination problem. We will discuss several novel approaches to solve this transmission coordination problem in great detail in Section 5.3.

In the AWG based single–hop network both transmitter and receiver have to be tunable. Given the significantly different tuning latencies of the various transceiver types only electro–optic transceivers with a tuning time of a few nanoseconds are reasonable candidates for realizing efficient packet switched single–hop networks, especially in a high–speed environment. Fast tunable transmitters with a tuning time in the range of a few nanoseconds are already used in metro WDM network testbeds [SWW+00][WSR+00][SWR+01]. With the advent of the Sampled Grating DBR (SG–DBR) laser, tunable transmitters not only with a negligible tuning time but also a significantly enlarged tuning range of several tens of nanometers, high output power, and large Side–Mode Suppression Ratio (SMSR) will be available [Mas00][WRRW00][LRB00]. On the other hand, fast tunable receivers are not yet available at the time of writing. We expect that this technological challenge will attract more attention for the following reasons:

1. Lu and Kleinrock have shown in [LK92b] that in networks with a relatively small number of wavelengths (similar to our AWG based network) it is more advantageous to have both transmitters and receivers tunable rather than having only either one of them tunable. To see this, note that at medium to high loads almost all wavelengths are in use and it is better to have a tunable receiver than one or more fixed–tuned receivers because the wavelengths those fixed–tuned receivers are tuned to may be all in use by other nodes and a given node could not receive any packet even though other wavelengths are free.

2. Deploying tunable transmitters and receivers at each node allows for load balancing since traffic between a given pair of nodes can be sent on any wavelength. In particular for nonuniform traffic, load balancing increases the channel utilization and improves the throughput–delay performance of the network [Sim98].

3. Assigning a separate dedicated home wavelength to the fixed–tuned receiver(s) of one (or more) node(s) leads to low channel utilization since a given wavelength cannot be used by other receivers while it is idle [SS00a].

4. Supporting multicast is less efficient in a network with nodes deploying fixed–tuned receivers. In a system with each node having a separate home channel for reception, a given multicast packet has to be transmitted multiple times. This results in a poor bandwidth efficiency as opposed to a system in which each node can tune its receiver to the corresponding wavelength of the source node.

Wavelength–agile transceivers are expected to be more expensive than their fixed–tuned counterparts. Although it is difficult to predict the prices of future commercially available tunable transceivers, we still are able to make the following observations. For economical reasons, each node in the AWG based single–hop network should not use more than one fast tunable transceiver. Note that one tunable transceiver per node might be enough to produce performance close to the upper bound [LK92b]. This is because for uniform traffic and a small number of wavelengths, the probability that more than one packet is destined to the same destination node and finding a free wavelength is very small. For nonuniform traffic patterns, e.g., servers representing network hot spots, it might be sufficient to equip only the hot–spot nodes with multiple transceivers while all other nodes deploy a single transceiver. In real–world networks each tunable transceiver has to be protected by a secondary transceiver in order to provide survivability in case of failure. Hence, in a real–world single–hop network each node would

consist of two identical tunable transceivers. Whereas the multihop network requires more than one backup transceiver at each node. Every fixed–tuned transceiver has to be protected by another one operating on the same wavelength. As a result, the number of (fixed–tuned) transceivers per node in the multihop network can become quite large [O'D00]. Typically, this array of transceivers is integrated on a single chip for cost reasons. If one transceiver fails, the entire array has to be replaced, wasting transceivers which are still functional. The network costs are not only driven by the number and type of transceivers but also by costs related to operation, maintenance, power consumption, performance monitoring, and management. In particular from the network management and performance monitoring perspective, a smaller number of transceivers per node is preferrable. All these cost factors have to be taken into account when discussing Fig. 4.5. From this figure we observe that an AWG based multihop network requires four or more fixed–tuned transceivers at each node for providing a larger network capacity than a single–hop network which uses one single fast tunable transceiver per node. Whether four fixed–tuned transceivers are less expensive to purchase and operate than one single fast tunable transceiver remains to be seen in the future.

4.2 PSC vs. AWG based single–hop networks

In the previous section we have seen that fast tunable transceivers can be tuned in a negligible amount of time but provide only a limited tuning range. This translates into a relatively small number of available wavelengths. However, all these wavelengths can be used at each AWG port simultaneously without resulting channel collisions. In this section, we analytically show how the throughput–delay performance of an AWG based single–hop network benefits from spatial wavelength reuse. To demonstrate the potential performance gain due to spatial wavelength reuse we compare the AWG based single–hop network with a PSC based single–hop network. To date, most (logical) single–hop WDM networks are embedded on a physical PSC based star configuration which does not allow for spatial wavelength reuse. We show that the AWG clearly outperforms the PSC in terms of throughput, delay, and packet loss [MW00][MW01].

4.2.1 Spatial wavelength reuse

Unlike the PSC the AWG allows for spatial wavelength reuse, i.e., each wavelength can be applied at all AWG input ports simultaneously without resulting in channel collisions at the AWG output ports. Ideally, the AWG routes each wavelength to a different output port without causing any channel crosstalk at the other AWG output ports. However, real AWGs suffer from leakage due to optical path phase errors [TYI95]. As a consequence, each wavelength is routed not only to the intended AWG output port but is also received in part at the remaining AWG output ports. Thus, using the same wavelength at multiple AWG input ports simultaneously leads to interferometric signal–crosstalk beat noise at the AWG output ports [TOT96]. The resulting intrachannel crosstalk has the same nominal wavelength as the proper signal and cannot be removed by a demultiplexer. This homodyne beat noise puts limitations on the network scalability with respect to bit rate, number of wavelength channels per fiber, and number of AWG input/output ports [GLG99].

The signal–crosstalk beat noise has a detrimental impact on the bit error rate (BER). It was shown in [GE95] that independent of the bit rate the worst case power penalty for matched

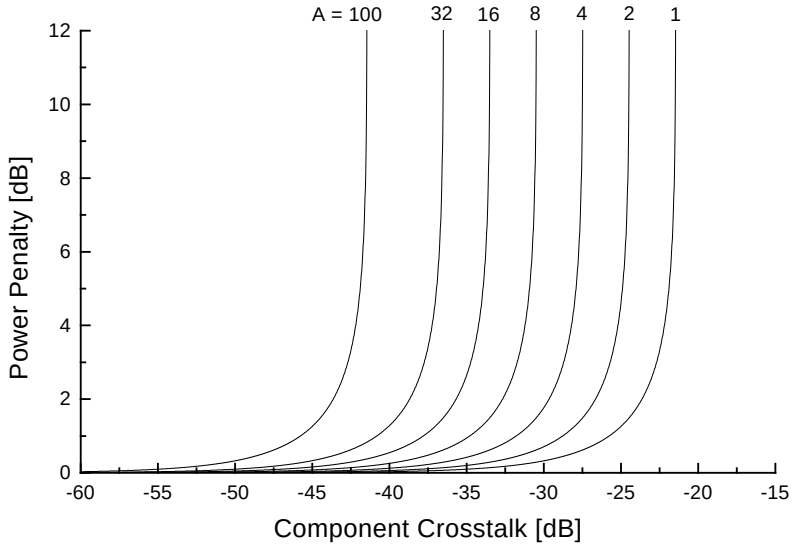

Figure 4.7: Power penalty vs. component crosstalk.

polarization states is given by

$$P = -5 \log_{10} \left[1 - 4q^2 A 10^{\epsilon_{dB}/10} \right] \text{dB} \tag{4.12}$$

where ϵ_{dB} is the component (AWG in our case) crosstalk in dB, A denotes the number of interfering crosstalk terms, and $q = 5.9$ for an error rate of 10^{-9}. Fig. 4.7 depicts the power penalty given in Eqn. (4.12) for $A \in \{1, 2, 4, 8, 16, 32, 100\}$. Clearly, spatial wavelength reuse in an AWG based single–hop network is possible only with a certain penalty. Fig. 4.7 shows that for a realistic AWG with a crosstalk of approximately -35 dB and a power penalty of 1 dB a given wavelength cannot be spatially reused more than eight times. Therefore, in this work we consider only AWGs whose port number is not larger than eight. Note that this limitation holds for AWGs fabricated in planar technology. However, by using free–space AWGs an adjacent channel rejection below -40 dB and a background crosstalk rejection below -60 dB can be achieved [HCA+96]. These free–space AWGs could have a larger number of ports thereby allowing for more extensive spatial wavelength reuse.

Spatial wavelength reuse significantly increases the degree of concurrency, i.e., more transmissions can take place simultaneously. Given N input ports, an AWG which deploys one FSR consisting of N wavelengths is able to support up to N^2 transmissions at the same time as compared to a PSC which allows for only N simultaneous transmissions without resulting channel collisions at the corresponding output ports. Consequently, for N simultaneously transmitting nodes in AWG based single–hop networks the wavelength pool can be kept small by requiring only $\left\lceil \sqrt{N} \right\rceil$ wavelengths which are spatially reused at all AWG input ports, where $\lceil x \rceil$ denotes the smallest integer which is larger than or equal to x. In contrast, in PSC based single–hop networks the number of required wavelengths grows linearly with the number of simultaneously transmitting nodes N, each sending on a separate wavelength. Fig. 4.8 depicts the relation between the number of required wavelengths and the number of simultaneously transmitting

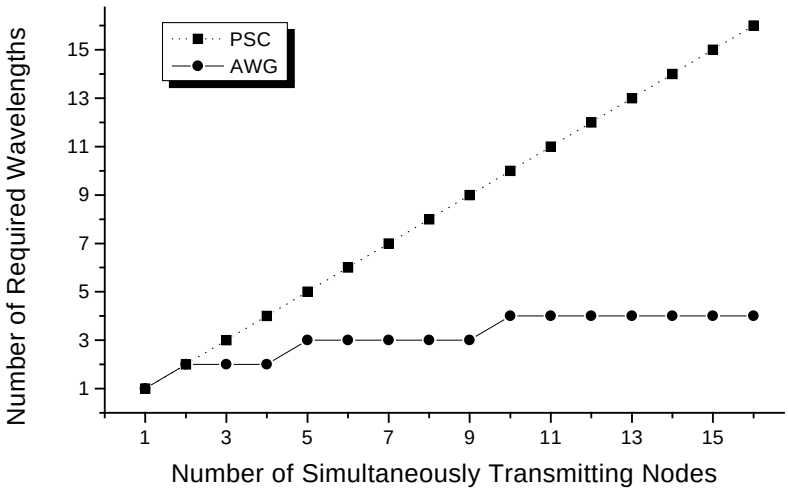

Figure 4.8: Relation between wavelength pool size and population.

nodes for both AWG and PSC based single–hop networks. Apparently, the AWG needs significantly fewer wavelengths than the PSC, especially for larger populations N. This in turn allows AWG based single–hop networks to deploy electro–optic transceivers with a tuning time of a few nanoseconds for populations where PSC based single–hop networks already have to resort to acousto–optic transceivers. We have seen in Section 4.1.1 that acousto–optic transceivers provide a wider tuning range but suffer from a significantly larger tuning latency. For example, for $N = 64$ simultaneously transmitting nodes and a channel spacing of 1.6 nm (200 GHz at 1.55 μm) an 8×8 AWG based single–hop network requires transceivers with a tuning range of 11.2 nm as opposed to 100.8 nm for a PSC based network. Hence, in the 8×8 AWG based network electro–optic transceivers can be used whose tuning time is three orders of magnitude smaller than that of acousto–optic ones which would be necessary in the PSC based network. Due to the smaller tuning penalty, wavelength channels are utilized more efficiently in AWG based single–hop networks than in their PSC based counterparts resulting in an improved performance, as we will see shortly.

4.2.2 Architecture and wavelength assignment

In order to compare PSC and AWG based single–hop networks we consider a given population of N nodes. Both networks are supposed to provide full connectivity and to allow each node to transmit at any given time.

In the PSC based single–hop network each node is attached to a different PSC port, as shown in Fig. 4.9. Every node is equipped with an identical transceiver consisting of one transmitter and one receiver. The transmitter (Tx) of a given node is attached to one of the N PSC input ports while the corresponding receiver (Rx) is located at the opposite output port. Since all nodes are supposed to be active at any time N different wavelengths are required to avoid channel collisions. In the PSC based single–hop network the transmitter and/or the

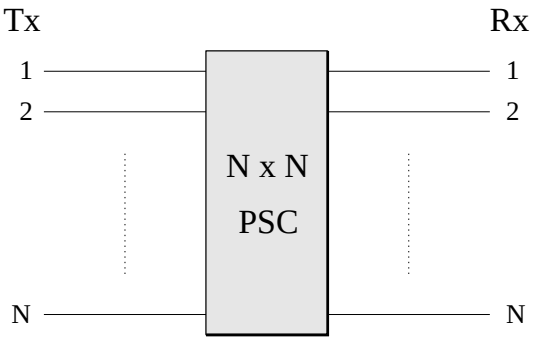

Figure 4.9: PSC based single–hop network architecture.

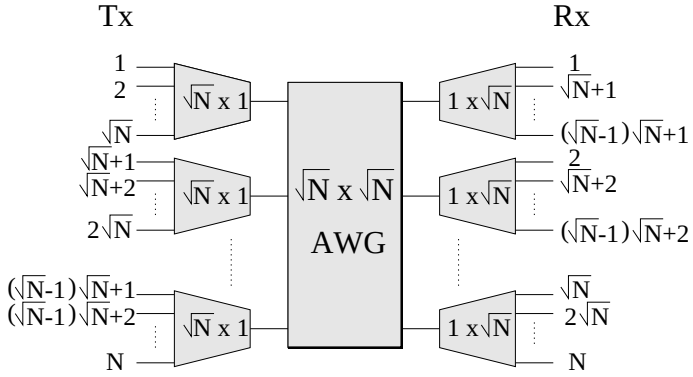

Figure 4.10: AWG based single–hop network architecture.

receiver of each node have to be tunable in order to guarantee full connectivity in one single hop [Muk92]. With a fixed–tuned receiver and a tunable transmitter each node has its own home channel for reception and all other nodes have to tune their tunable transmitters to the destination node's home channel for communication. Similarly, with a fixed–tuned transmitter and a tunable receiver each node transmits on a separate wavelength and the remaining nodes have to tune their tunable receivers to the corresponding source node's wavelength for reception. (Note that both transmitter and receiver can but do not necessarily have to be tunable in order to guarantee full connectivity in one single hop for N simultaneously busy nodes.) In either case each node experiences a certain latency while tuning its transmitter and/or receiver from one to another wavelength.

While in the PSC based single–hop network the degree (port number) of the underlying PSC is equal to the number of nodes N, the degree of the AWG has to be only $\sqrt{N}$ (for the sake of simplicity we assume in the following that $\sqrt{N}$ is an integer; otherwise, we have to take $\left\lceil \sqrt{N} \right\rceil$). This is due to the fact that at each port of the $\sqrt{N} \times \sqrt{N}$ AWG $\sqrt{N}$ wavelengths are spatially reused leading to a total of N channels which are sufficient to support N simultaneously active nodes. Fig. 4.10 depicts the AWG based single–hop network where to each AWG input and output port a $\sqrt{N} \times 1$ combiner and a $1 \times \sqrt{N}$ splitter are attached, respectively. These devices are necessary for attaching $\sqrt{N}$ nodes to each AWG port. Note that all combiners and splitters have to be wavelength insensitive. This is because in an AWG based single–hop network both

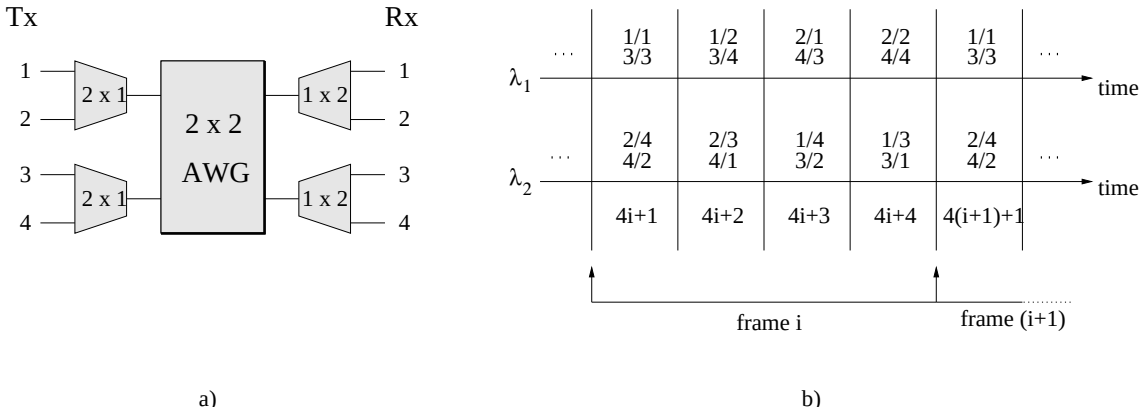

Figure 4.11: a) AWG based single–hop network without cyclic receiver attachment, b) fixed wavelength assignment ($N = 4$).

transmitter and receiver of each node have to be tunable in order to provide full connectivity in one single hop (see Section 4.1.1). Every node is equipped with an identical transceiver consisting of one tunable transmitter and one tunable receiver. The tunable transmitter of a given node is attached to one of the combiner input ports while the corresponding tunable receiver is located at the opposite splitter output port. Each node is able to transmit on $\sqrt{N}$ different wavelengths which have to be fed into the corresponding AWG input port. Similarly, each node receives data from the remaining nodes on $\sqrt{N}$ different wavelengths. Hence, splitters and combiners have to collect and distribute all $\sqrt{N}$ wavelengths identically which implies that both combiner and splitter have to be wavelength insensitive. A positive side effect of splitters is that they enable optical multicasting by distributing the incoming optical signal equally to the attached receivers. The benefit of optical multicasting will be investigated at length in Section 6.3. On the other hand, splitters and also combiners inherently suffer from splitting loss. The resulting splitting loss of the AWG based single–hop network is equal to $2 \cdot 10 \log \sqrt{N} = 10 \log N$ (in dB) which is identical to the splitting loss of the $N \times N$ PSC based counterpart.

Note that in Fig. 4.10 the receivers are attached in a cyclic manner. The reason for this is illustrated in Figs. 4.11 and 4.12 for $N = 4$. Fig. 4.11 a) depicts the architecture without cyclic receiver attachment. Recall that the network is supposed to allow all $N = 4$ nodes to transmit/receive at any time. If nodes attached to the same combiner simultaneously send on the same wavelength a channel collision occurs at the corresponding AWG input port. To avoid not only channel but also receiver collisions we apply a round–robin Time Division Multiplexing (TDM) wavelength assignment scheme as shown in Fig. 4.11 b). Time is divided into frames which are repeated periodically. Each frame consists of $N = 4$ slots whose length is equal to the packet transmission time plus the transceiver tuning time. The length of the packets is assumed to be constant. Transceivers are tuned on a per–packet basis. In each slot communication between $N = 4$ pairs of nodes take place where X/Y denotes node X transmitting to node Y. Clearly, bandwidth is wasted in slots where $X = Y$. Due to the wavelength routing characteristics of the AWG all these slots are assigned only to wavelength λ_1.

Bandwidth can be saved by attaching the receivers to the splitters in a cyclic manner, as shown in Fig. 4.12 a). The resulting wavelength assignment is depicted in Fig. 4.12 b). Note that the first slot of each frame can be omitted since it contains only transmitter–receiver pairs where the transmitter and the receiver belong to the same node. In doing so, the frame length

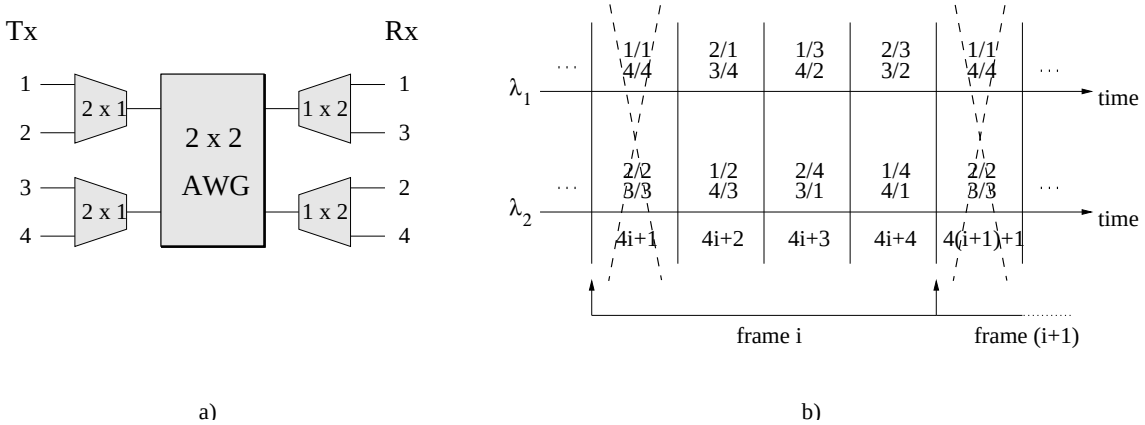

Figure 4.12: a) AWG based single–hop network with cyclic receiver attachment, b) fixed wavelength assignment with reduced frame length ($N = 4$).

is reduced by one slot thus saving bandwidth. This idea is valid for arbitrary N. In general, receiver i is attached to the splitter located at AWG output port j according to the following rule

$$j = \left\lceil (i-1) \bmod \sqrt{N} \right\rceil + 1, \quad i \in \{1, 2, \ldots, N\},\ j \in \{1, 2, \ldots, \sqrt{N}\}. \tag{4.13}$$

The resulting frame contains $(N-1)$ slots providing full connectivity.

The capacity C of both networks is defined as

$$C = \frac{N}{1+\tau} \tag{4.14}$$

where N denotes the number of simultaneously transmitting nodes, time is normalized to the packet transmission time, and τ denotes the normalized transceiver tuning time. The network capacity is identical to the maximum number of simultaneously transmitting nodes at any time. We have already seen in Section 4.1 that the transceiver tuning latency has a large impact on the network capacity. Fig. 4.13 depicts the impact of the tuning penalty of the various transceiver types on the aggregate capacity of the AWG and PSC based single–hop networks as a function of the number of nodes N. For a channel spacing equal to 100 GHz (0.8 nm at 1.55 μm) and a packet length of 10^4 bits the discontinuities represent the necessary transitions from electro–optic to acousto–optic and mechanical transceiver technology for increasing N and thereby larger transceiver tuning ranges. We observe that the AWG based single–hop network clearly outperforms its PSC based counterpart in terms of aggregate capacity. Due to spatial wavelength reuse fast tunable transceivers with a tuning time of a few nanoseconds (see Table 4.1 in Section 4.1.1) can be deployed up to approximately $N = 300$ nodes. In contrast, for $N = 300$ in the PSC based single–hop network, each node has to use mechanically tunable transceivers which provide a sufficiently large tuning range but suffer from a tuning time of several milliseconds. The tuning overhead significantly decreases the aggregate capacity, especially at a higher line rate where the packet transmission time becomes smaller. Moreover, the PSC based network is not able to accomodate more than $N = 626$ simultaneously transmitting nodes due to the limited tuning range of (mechanically) tunable transceivers. Whereas in the AWG based network up to $N = 32^2 = 1024$ nodes can transmit at the same time while still using acousto–optic transceivers with a tuning time in the range of a few microseconds. Note that Fig. 4.13 is intended to

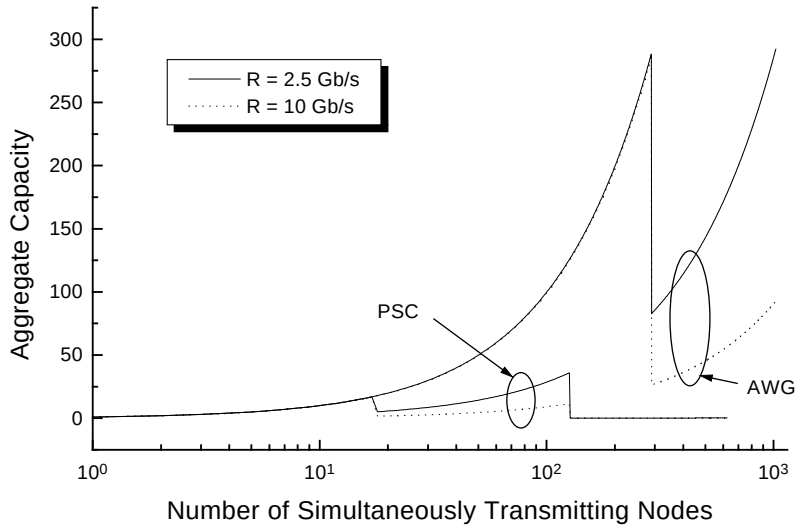

Figure 4.13: Aggregate capacity vs. number of simultaneously transmitting nodes N (R denotes the line rate).

illustrate the fundamental difference between AWG and PSC based single–hop networks and does not account for other network aspects such as power budget and crosstalk.

4.2.3 Analysis

In this section, we take into account that typically nodes do not continuously have (data) packets to send resulting in a smaller channel utilization and an aggregate throughgput less than the network capacity. We investigate the throughput–delay performance of both AWG and PSC based single–hop networks for random data traffic. For the throughput–delay performance analysis the source/destination allocation protocol analytic approach reported in [CG88c] was slightly modified in order to accomodate spatial wavelength reuse and packet loss.

There are N nodes simultaneously transmitting over $\sqrt{N}$ wavelengths in case of the AWG and over N wavelengths in case of the PSC. Each node can transmit/receive one packet at a time on/from any of the wavelengths. Every node has N single–packet buffers, one for reception and $(N-1)$ for transmission. Thus, the interconnections between each pair of nodes can be modeled independently. Each buffer represents an independent (virtual) user (i,j), $1 \leq i,j \leq N$. The packet length is assumed to be constant. Time is normalized to the packet transmission time. Time is divided into cycles (equivalent to the above mentioned frames) which contain $(N-1)$ slots. Each slot is composed of the packet transmission time (unity) and the normalized transceiver tuning time τ. Every pair of virtual users is assigned one slot per cycle. The arrival process is assumed to be Poisson with the average arrival rate of λ packets per time unit per user. An idle user is defined as a user with an empty buffer and a backlogged user is defined as a user with a packet for transmission. Arriving packets are discarded if the user is backlogged, i.e., if the buffer is full. The traffic between any pair of users is assumed to have the same mean arrival rate λ.

The allocation matrix $U(t)$ is an $N \times N$ matrix whose elements $u_{ij}(t)$ represent the channel number (wavelength) on which user (i,j) can transmit in slot t, $t = 1, 2, \ldots, (N-1)$, and $1 \leq i, j \leq N$. $W(t)$ is a binary matrix with the elements $w_{ij}(t) = Ind(u_{ij}(t) > 0)$, where the indicator function is given by

$$Ind(statement) = \begin{cases} 1 & \text{, if } statement \text{ true} \\ 0 & \text{, if } statement \text{ false .} \end{cases} \tag{4.15}$$

The allocation matrix $U(t)$ is subject to the following conditions:

- $\sum_{\forall i} w_{ij}(t) \leq 1$, i.e., no receiver collisions occur
- $\sum_{\forall i} \sum_{\forall j} w_{ij}(t) = N$, i.e., the number of simultaneous transmissions is restricted to N
- In case of the PSC, for every $w_{ij}(t) \neq 0$: $u_{ij}(t) \neq u_{kl}(t)$ if $i \neq k$ and $j \neq l$, i.e., no channel collisions occur
- In case of the AWG, for every $w_{ij}(t) \neq 0$: $u_{ij}(t) \neq u_{kl}(t)$ if $\left\lceil \frac{i}{\sqrt{N}} \right\rceil = \left\lceil \frac{k}{\sqrt{N}} \right\rceil$ and $i \neq k, j \neq l$, i.e., no channel collisions occur
- $\sum_{\forall j} w_{ij}(t) \leq 1$, i.e., a node can transmit on at most one channel.

The system is observed at the regeneration points embedded at the beginning of each slot. The throughput of user (i,j), defined as the number of successfully transmitted packets of user (i,j) per slot is given by

$$S_{ij} = \frac{1}{(N-1)(1+\tau)} \sum_{t=1}^{N-1} w_{ij}(t) \cdot \pi_{ij}(t) \tag{4.16}$$

where $\pi_{ij}(t)$ denotes the steady–state probability that user (i,j) is backlogged (packet in buffer) at the beginning of slot t. Hence, the system throughput, defined as the total number of successfully transmitted packets per slot is

$$S = \sum_{\forall i} \sum_{\forall j} S_{ij}. \tag{4.17}$$

To evaluate the aggregate throughput $\pi_{ij}(t)$ is needed. The probability $\pi_{ij}(t)$ can be expressed as a function of the probability $\pi_{ij}(t-1)$, mean packet arrival rate and the matrix $W(t)$:

$$\pi_{ij}(t) = [1 - \pi_{ij}(t-1)] \left(1 - e^{-\lambda(1+\tau)}\right) + \pi_{ij}(t-1) \left[1 - w_{ij}(t-1)\right] \tag{4.18}$$

$$= \pi_{ij}(t-1) \left[e^{-\lambda(1+\tau)} - w_{ij}(t-1)\right] + \left(1 + e^{-\lambda(1+\tau)}\right), \; 2 \leq t \leq (N-1). \tag{4.19}$$

This leads to the following recursive formula

$$\pi_{ij}(t) = \pi_{ij}(1) \prod_{k=1}^{t-1} \left[e^{-\lambda(1+\tau)} - w_{ij}(k)\right] + \left(1 - e^{-\lambda(1+\tau)}\right) \cdot \cdot \left\{ \sum_{l=2}^{t-1} \prod_{k=l}^{t-1} \left[e^{-\lambda(1+\tau)} - w_{ij}(k)\right] + 1 \right\}, \; 2 \leq t \leq (N-1). \tag{4.20}$$

Assuming the system is in steady state, we equate

$$\pi_{ij}(N) = \pi_{ij}(1). \tag{4.21}$$

Substituting (4.21) in (4.20), we obtain

$$\pi_{ij}(1) = \left(1 - e^{-\lambda(1+\tau)}\right) \cdot \frac{\sum_{l=2}^{N-1} \prod_{k=l}^{N-1} \left[e^{-\lambda(1+\tau)} - w_{ij}(k)\right] + 1}{1 - \prod_{k=1}^{N-1} \left[e^{-\lambda(1+\tau)} - w_{ij}(k)\right]} \tag{4.22}$$

and for $2 \leq t \leq (N-1)$ we finally obtain

$$\begin{aligned}\pi_{ij}(t) = & \left(1 - e^{-\lambda(1+\tau)}\right) \cdot \frac{\sum_{l=2}^{N-1} \prod_{k=l}^{N-1} \left[e^{-\lambda(1+\tau)} - w_{ij}(k)\right] + 1}{1 - \prod_{k=1}^{N-1} \left[e^{-\lambda(1+\tau)} - w_{ij}(k)\right]} \cdot \prod_{k=1}^{t-1} \left[e^{-\lambda(1+\tau)} - w_{ij}(k)\right] + \\ & + \left(1 - e^{-\lambda(1+\tau)}\right) \cdot \left\{ \sum_{l=2}^{t-1} \prod_{k=l}^{t-1} \left[e^{-\lambda(1+\tau)} - w_{ij}(k)\right] + 1 \right\}.\end{aligned} \tag{4.23}$$

Using Little's Law, the mean packet delay of user (i,j), defined as the average time between the arrival of a packet at user (i,j) and the beginning of its transmission is given by

$$D_{ij} = \frac{Q_{ij}}{S_{ij}}, \tag{4.24}$$

where Q_{ij} denotes the mean backlog (over time) at user (i,j). For the average packet delay in the system we obtain

$$D = \sum_{\forall i} \sum_{\forall j} \frac{S_{ij}}{S} D_{ij}. \tag{4.25}$$

For the evaluation of Q_{ij} we introduce the following definitions:

- r_{ij}: the number of transmission permissions per cycle (note that in the considered channel allocation scheme $r_{ij} = 1, 1 \leq i,j \leq N$)
- *idle period*$_{ij}$: the time interval between two consecutive grantings of permissions to user (i,j), or between a permission and the cycle boundary
- n_{ij}: the number of *idle period*$_{ij}$ intervals of user (i,j) in the allocation cycle
- s_{ijl}: the number of idle slots in the l-th *idle period*$_{ij}$ $(1 \leq l \leq n_{ij})$.

First, we calculate $Res(s_{ijl}, \lambda)$ which denotes the expected value of the residual time user (i,j) is backlogged in the l-th *idle period* with the length of s_{ijl} slots.

$$Res(s_{ijl}, \lambda) = \int_0^{s_{ijl}(1+\tau)} \left[s_{ijl}(1+\tau) - t\right] \lambda e^{-\lambda t} dt \tag{4.26}$$

$$= \frac{e^{-\lambda s_{ijl}(1+\tau)} + \lambda s_{ijl}(1+\tau) - 1}{\lambda}. \tag{4.27}$$

Q_{ij} is obtained by the weighted average of the residual time in each idle period

$$Q_{ij} = \sum_{k=1}^{n_{ij}} \frac{s_{ijk}}{(N-1) - r_{ij}} \cdot \frac{Res(s_{ijk}, \lambda)}{s_{ijk} \cdot (1+\tau)} \tag{4.28}$$

$$= \sum_{k=1}^{n_{ij}} \frac{e^{-\lambda(1+\tau) \cdot s_{ijk}} + \lambda(1+\tau) \cdot s_{ijk} - 1}{[(N-1) - r_{ij}] \cdot \lambda(1+\tau)}. \tag{4.29}$$

Thus, using equations (4.17) and (4.25) the network throughput and the average queueing delay can be evaluated.

The blocking probability is equal to the probability that an arriving packet finds the node backlogged. Assuming uniform traffic all virtual users behave identically. Thus, any arbitrary virtual user, say (i, j), can be considered. Using (4.22) and (4.23) the blocking probability P_B can be obtained by

$$P_B = \frac{1}{N-1} \sum_{t=1}^{N-1} \pi_{ij}(t). \tag{4.30}$$

4.2.4 Results

We compare the AWG and PSC based single–hop networks in terms of throughput, delay, and packet loss for different numbers of simultaneously transmitting nodes N. Recall that in the AWG based network $\sqrt{N}$ wavelengths are required as opposed to N wavelengths in the PSC based counterpart. The channel spacing is assumed to be 200 GHz (1.6 nm at 1.55 μm). Fast tunable electro–optic transceivers are assumed to have a tuning range of 10 nm and a tuning time of 10 ns while acousto–optic transceivers are assumed to be tunable over a range of 100 nm with a tuning time of 10 μs. Consequently, fast tunable transceivers can be deployed as long as the number of wavelengths is not larger than 7. Otherwise acousto–optic or mechanically tunable transceivers must be used. Packets are assumed to have a fixed size of 10^4 bits being transmitted/received at a line rate equal to 10 Gb/s. Thus, the normalized tuning time τ equals 10^{-2} for electro–optic transceivers and 10 for acousto–optic transceivers. Recall that wavelengths are assigned in a slotted round–robin TDM scheme. Each slot comprises the packet transmission time and the transceiver tuning time. For illustration, we show plots for small values of $N \in \{4, 9, 16\}$. As we will see, for increasing N the performance difference between AWG and PSC based single–hop networks becomes more dramatic.

Fig. 4.14 depicts the throughput (given in packets/packet transmission time) vs. the mean arrival rate λ (in packet/packet transmission time). The AWG based network clearly outperforms the PSC based network. For $N = 4$ the wavelength pool size is small enough to use fast tunable transceivers in both networks. For $N \in \{9, 16\}$, on the other hand, electro–optic transceivers can be deployed only in the AWG based network. In the PSC based network acousto–optic transceivers whose tuning latency is three orders of magnitude larger have to be used. As a consequence, the channel utilization decreases significantly resulting in a reduced aggregate throughput. In general, the aggregate throughput grows with increasing λ. Note that the PSC based network runs into saturation earlier, i.e., at lower traffic loads. This is because due to the longer slot and thereby frame duration a user is more likely backlogged when the corresponding slot is assigned to it. Owing to the negligible tuning time of electro–optic transceivers the maximum aggregate throughput of the AWG based network is almost equal to N which is identical to the network capacity. For larger N the throughput increases in both networks since more wavelengths lead to a higher degree of concurrency, resulting in an improved aggregate throughput. In addition, the saturation is reached at lower loads due to the longer frame length. For larger N (longer frame) a user is more likely backlogged when the corresponding slot is allocated to it. Note that the throughput difference between the two single–hop networks becomes more evident for increasing N.

In Fig. 4.15 the mean queueing delay (in packet transmission time) vs. the mean arrival rate λ (in packet/packet transmission time) for nonblocked packets is depicted. Again, the AWG based network clearly outperforms its PSC based counterpart, in particular for larger

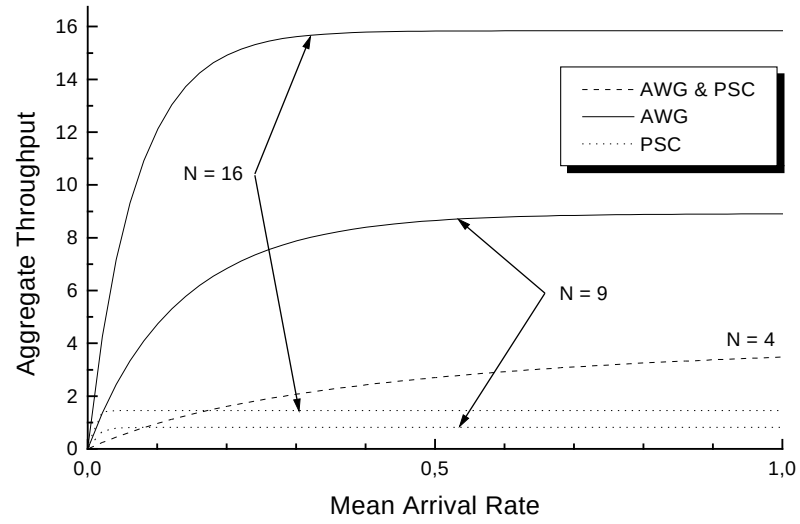

Figure 4.14: Aggregate throughput (packets/packet transmission time) vs. mean arrival rate (packet/packet transmission time).

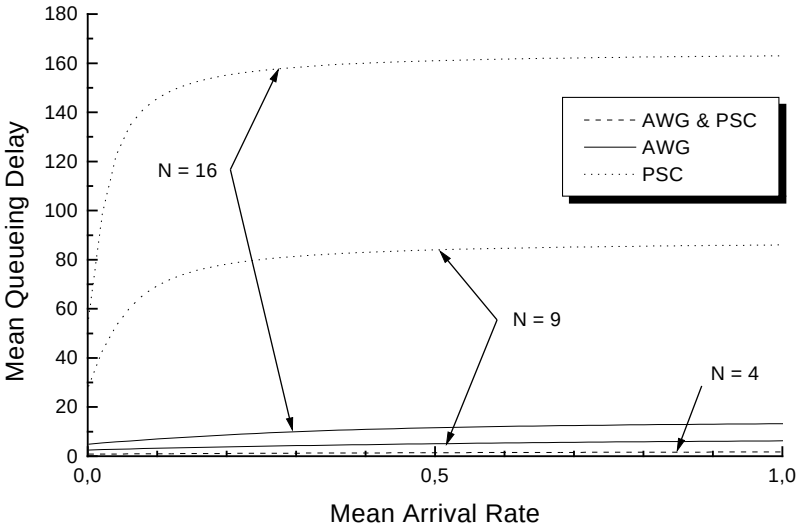

Figure 4.15: Mean queueing delay (packet transmission time) vs. mean arrival rate (packet/packet transmission time).

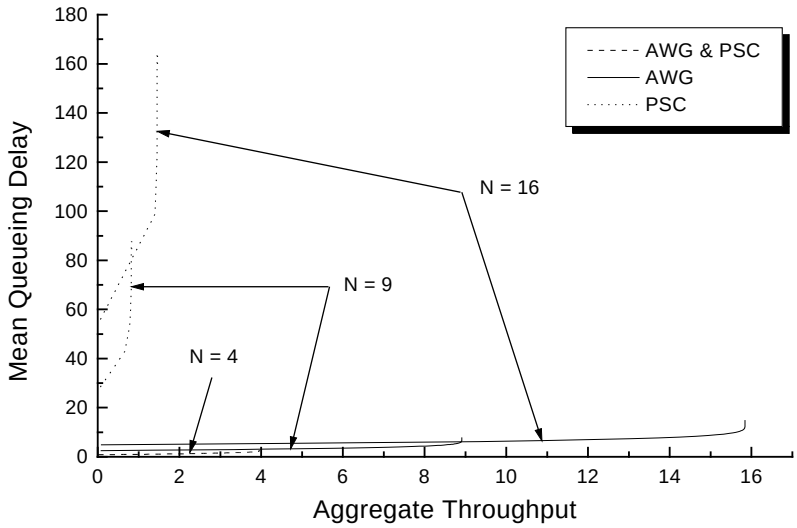

Figure 4.16: Mean queueing delay (packet transmission time) vs. aggregate throughput (packets/packet transmission time).

N. In both networks the mean queueing delay grows and the saturation is reached earlier for increasing N due to the longer frame size. At high loads the mean queueing delay in both networks asymptotically approaches the maximum value which is identical to the frame length $(N-1)(1+\tau)$. Again, for larger N the difference between the two networks becomes more dramatic since for $N \in \{9, 16\}$ fast tunable transceivers can be deployed only in the AWG based network. Note that the mean queueing delay is upper bounded since our analytical model assumes single–packet buffers at each user. If the buffer already contains one packet new arriving packets are discarded and do not add to the mean queueing delay.

Fig. 4.16 shows the mean queueing delay (packet transmission time) vs. aggregate throughput (packets/packet transmission time) for nonblocked packets. The figure clearly shows that for larger N the AWG based network is significantly superior to the PSC based network in terms of throughput and delay.

The blocking probability as a function of the mean arrival rate λ (packet/packet transmission time) is shown in Fig. 4.17. We have seen that for $N \in \{9, 16\}$ backlogged nodes in the PSC based network experience a higher mean queueing delay compared to the AWG based network. As a consequence, in the PSC based network new arriving packets find buffers already full with a higher probability than in the AWG based network which translates into a higher blocking probability. For larger N the blocking probability rises earlier due to the longer frame size.

Note that in real–world systems such high packet loss rates are not acceptable. There are three solutions to this problem:

- The system is run only at light traffic loads. This approach is not attractive since at low loads the aggregate throughput is small as well.

- The single–packet buffers are replaced with larger buffers. In doing so, arrriving packets are

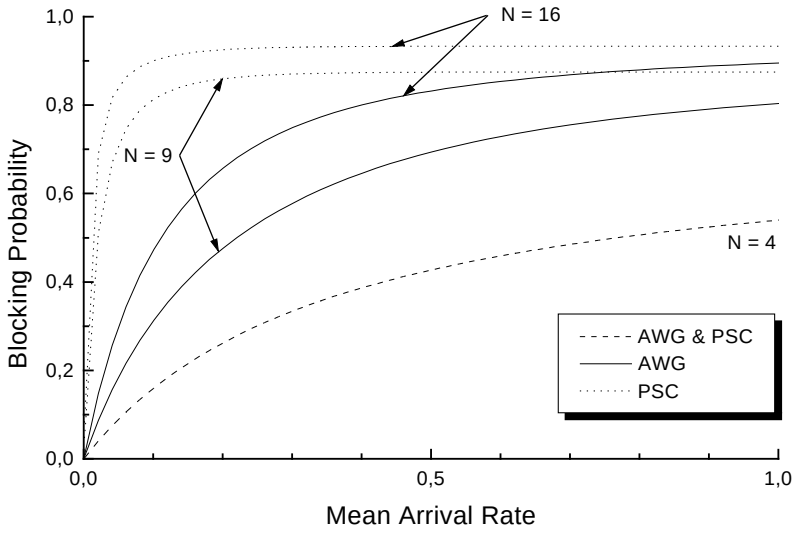

Figure 4.17: Blocking probability vs. mean arrival rate (packet/packet transmission time).

stored resulting in a smaller blocking probability while providing an acceptable throughput. This approach is straightforward and will be investigated Section 6.4.2.

- The third solution is traffic shaping which is discussed in the following.

Recall that the packet arrival process is assumed to be Poisson. Due to the randomness of the arriving process, packets are lost, especially at high arrival rates. Traffic shaping aims at smoothing the arriving traffic such that the interarrival times become deterministic. This could easily be done by using the leaky bucket method which feeds packets into each node's single–packet buffers at a constant rate. This transmission rate of the shaper must not be larger than the service rate per user in order to avoid packet overflow and thereby packet loss in the single–packet buffers. The upper limit of the shaper transmission rate is given by

$$Rate \leq \frac{1}{(N-1)\cdot(1+\tau)}, \tag{4.31}$$

which simply states that the shaper is allowed to put at most one packet in a given single–packet buffer per frame which consists of $(N-1)(1+\tau)$ slots. No packets are lost at each user's single–packet buffer if the shaper puts packets at a rate less than or equal to this limit. Fig. 4.18 shows the maximum shaper transmission rate vs. the number of simultaneously transmitting nodes N. Again, we observe the discontinuities which represent the transition from one to another transceiver technology as a wider tuning range is required to accomodate additional nodes. Larger normalized transceiver tuning times decrease the maximum shaper transmission rate which is further decreased at higher line rates due to the shorter frame length. We observe that in the AWG based network the shaper is allowed to transmit packets at a higher rate than in the PSC based network.

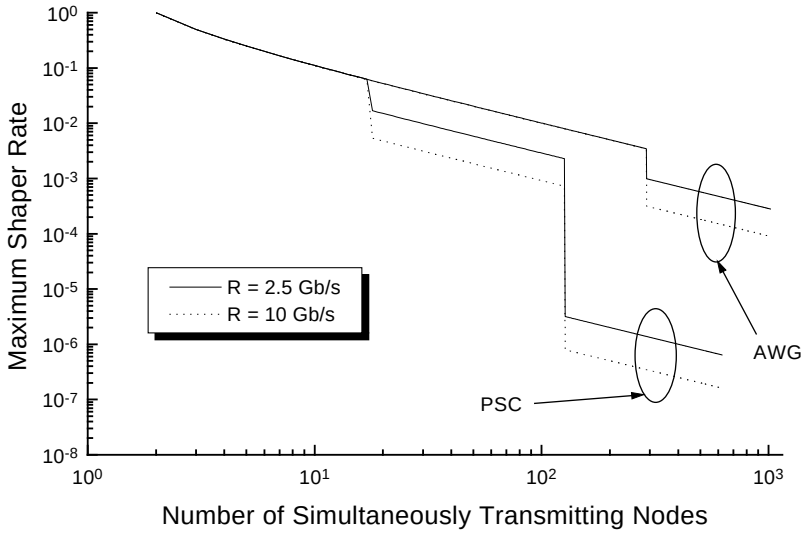

Figure 4.18: Maximum shaper transmission rate (packet/packet transmission time) vs. number of simultaneously transmitting nodes (R denotes the line rate).

Finally, we note that the queueing delay in each single–packet buffer is upper bounded as follows

$$Queueing\ delay \leq (N-1) \cdot (1+\tau) \tag{4.32}$$

which is identical to the number of slots between two successive transmission permissions per user. The upper bound is illustrated in Fig. 4.19. Again, we see that due to spatial wavelength reuse and the implied smaller transceiver tuning penalty the AWG based single–hop network clearly outperforms the PSC based one.

4.2.5 Discussion

Fast tunable transceivers provide a negligible tuning time at the expense of a small tuning range. As a result, channels are utilized efficiently but only a few wavelengths are available at each port of our AWG based single–hop network. Due to spatial wavelength reuse, however, the number of communication channels is significantly increased. This allows a relatively large number of nodes to transmit/receive simultaneously without suffering from large transceiver tuning latencies. The extent of spatial wavelength reuse is mainly determined by the channel crosstalk of the underlying AWG.

For a given number of nodes we have compared an AWG and a PSC based single–hop network under the assumption of a fixed round–robin TDM wavelength assignment scheme. The AWG based single–hop network is superior to its PSC based counterpart in terms of throughput, delay, and packet loss. This is due to the fact that the AWG, as opposed to the PSC, allows for spatial wavelength reuse at each port resulting in a significantly smaller required transceiver tuning range due to the smaller wavelength pool size. Owing to the smaller transceiver tuning times this translates into a higher channel utilization and improved throughput–delay performance of

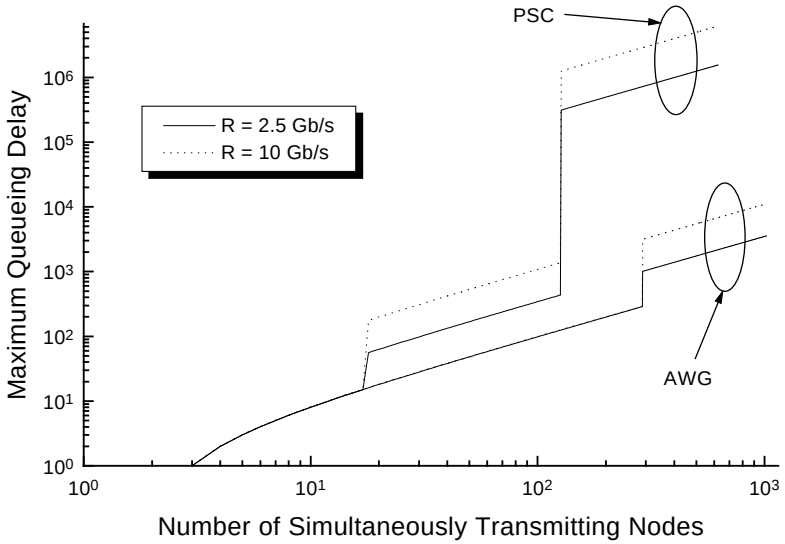

Figure 4.19: Maximum queueing delay (packet transmission time) vs. number of simultaneously transmitting nodes (R denotes the line rate).

the AWG based network compared to the PSC based network.

Note that the above applied fixed (static) wavelength assignment is only well suited for uniform and nonbursty (regular) traffic at medium to high system loads. For bursty data traffic, as typically present in computer communications networks, a large portion of the fixed–assigned slots are not used, resulting in a decreased channel utilization and wasted bandwidth. For bursty traffic, wavelengths have to be allocated more dynamically on demand in order to improve the throughput–delay performance of the network. This problem of efficient resource allocation for both bursty and nonbursty traffic will be tackled in Section 5.3.

4.3 Conclusions

The network under consideration has a physical star topology. The star configuration enables easy installation, configuration, management, and troubleshooting. Furthermore, it does not suffer from tapping loss as present in ring (and bus) topologies resulting in an improved optical power budget. In our first architectural comparison we have considered a physical star network using an AWG as central hub. Two types of logical topologies can be embedded on the physical AWG based star network: Single–hop and multihop networks. As opposed to the multihop network, both transmitters and receivers have to be tunable in an AWG based single–hop network with any–to–any connectivity. We have compared both logical topologies in terms of mean hop distance and network capacity. Due to the significantly different tuning times of the various transceiver types only single–hop networks deploying electro–optic transceivers with a negligible tuning latency of a few nanoseconds appear to be reasonable. The results were discussed taking economical aspects such as initial expenditure and operational costs into account as well. We have observed that if one fast tunable transceiver is less expensive than four fixed–tuned

transceivers the logical single–hop network is the better option.

In the remainder of this work we concentrate on (AWG based) single–hop networks since they provide several advantages such as transparency with respect to protocol, modulation format, and bit rate, simplified management, future–proofness, reduced nodal processing requirements, and an improved throughput–delay performance due to higher channel utilization and smaller propagation delays. We note that the drawbacks of single–hop networks are the relatively high costs of tunable transceivers, the possibly complex transmission coordination between tunable transceivers, and the limited network diameter which makes single–hop networks promissing candidates for LANs and MANs, but not necessarily for long–haul networks. We have also indicated that equipping each node with one single tunable transceiver is not only economically reasonable but also sufficient to achieve a network performance close to the upper bound.

To investigate the performance potential of AWG based single–hop networks we have compared it to a PSC based network which has been the most common type of single–hop WDM network so far. Due to spatial wavelength reuse the AWG based network requires fewer wavelengths which in turn allows for deploying tunable transceivers with a significantly smaller tuning latency. We have seen that real AWGs suffer from intrachannel crosstalk leading to interferometric signal–crosstalk beat noise and a deteriorated BER. As a consequence, for a power penalty of 1 dB the degree of planar AWGs must not be larger than eight while free–space AWGs allow for a much larger number of ports. Under the assumption of a fixed round–robin TDM wavelength allocation scheme we have demonstrated that the AWG based single–hop network clearly outperforms its PSC based counterpart in terms of throughput, delay, and packet loss. The splitting loss is the same in both AWG and PSC based networks due to the splitters/combiners attached to the AWG and the PSC, respectively. Another interesting comparison between AWG and PSC based single–hop WDM networks was conducted in [Ben99]. In that comparison the AWG based network nodes access the wavelengths in a round–robin TDMA fashion (as in our comparison), in the PSC based network the nodes gain access to the wavelengths by using two different reservation MAC protocols. It was shown that the AWG based network outperforms its PSC based counterpart over a wide range of parameter values for both real–time and nonreal–time traffic.

Note that for bursty traffic the above mentioned static wavelength assignment leads to poor resource utilization. To improve the throughput–delay performance of the network wavelengths have to be assigned dynamically, as we will discuss in the subsequent chapter.

Chapter 5

Architecture and Protocol

In this chapter, we specify our metro WDM network architecture and MAC protocol. Based on the results of the previous chapter, we propose a logical single–hop network embedded on a physical star WDM network using an AWG as central hub. In Section 5.1 we first list several requirements of networks in general and metro networks in particular, which we account for in our network design. In Section 5.2 we present the network and node architecture which makes use of the components introduced in Chapter 2. Prior to describing the architecture we explain the underlying principles. Specifically, we discuss what happens when different light source signals are fed into the AWG. We discuss how they can be used for building an efficient and cost–effective network and node architecture. After fixing the architecture we address the dynamic on–demand assignment of wavelengths. Section 5.3 explains the MAC protocol in detail and outlines how the following network requirements are satisfied [Mai01b][Mai01a][MRW02].

5.1 Network requirements

Networks have to meet a number of requirements at the architecture and/or protocol level. In the following we list the key requirements which have to be satisfied when designing a network architecture and protocol. Special attention is thereby paid to metro networks.

- *Reliability:* The network should be able to provide the required end–to–end functionality making sure the network is available to users for a specified period of time.
- *Survivability:* The network should be able to maintain an acceptable level of performance during network (node and/or link) failures by applying various protection and/or restoration techniques, and the mitigation or prevention of service outages from network failures.
- *Scalability:* It should be possible to add or remove network nodes in an easy and nondisruptive way without significantly degrading the network performance.
- *Connectivity:* Network connectivity enables each node to communicate with all other network nodes. Traffic should not have to traverse a large number of intermediate nodes to ensure smaller resource requirements and smaller propagation delays.
- *Future–proofness:* Future–proof networks are able to support future protocols at different bit rates without having to replace network components, thus preserving the investment value for future, i.e., not yet defined, developments.

- *Quality–of–Service:* Quality–of–Service (QoS) is the ability of the network to provide some level of assurance that the service requirements for different types of traffic, e.g., for delay–sensitive, real–time, and interactive applications, are satisfied. QoS is a measure (maximum, average, variance, etc.) acting on a property (throughput, delay, jitter, loss, etc.) of a cell/packet flow.
- *Fairness:* Fairness is the ability of the network to equally and sufficiently allocate network resources to all nodes which need to send data. In networks with fair channel access control each node ready to send data should have an equal opportunity to transmit.
- *Security:* Network security is the protection of the network and its services from unauthorized modification, destruction, or disclosure. It provides assurance that the network performs its critical functions correctly and there are no harmful side–effects.
- *Simplified OAM:* The operation, administration, and maintenance (OAM) aspects of the network should be as simple as possible in order to reduce network costs and overhead.
- *Multicast support:* The network should be able to provide point–to–multipoint connections in order to support multicast applications such as videoconferences and distributed games in an economical and bandwidth–efficient manner.

In addition, networks especially at the metro level have to exhibit the following properties:

- *Flexibility:* Metro networks collect a large variety of different client signals and connect them to the backbone network. Consequently, metro networks must be able to support a wide range of heterogeneous protocols such as ATM, Frame Relay, SONET/SDH, IP, ESCON, HIPPI, and Fibre Channel. This requires that the network is able to transport variable–size packets.
- *Cost–effectiveness:* Due to the smaller number of cost sharing clients/subscribers, metro networks are more cost sensitive than backbone networks. Therefore, the deployed networking components and the network/node architecture have to be economical and simple. Protocols must not perform complex operations.
- *Efficiency:* To meet the cost constraints, resources (wavelengths, transceivers) in metro networks have to be used efficiently.
- *Upgradability:* Content providers increasingly place proxy caches in metro networks in order to decrease the response time. To cope with the resulting increased local traffic, metro networks have to be easily upgradable. Advanced technologies, e.g., tunable tranceivers with a wider tuning range and a smaller tuning time, have to be incorporated without network service disruption and reconfiguration.

5.2 Architecture

Recall from Section 2.1.4 that transmitters can be either broadband or narrowband light sources. While the LED is an example for a broadband light source, lasers represent narrowband sources operating at a given wavelength. As we will see shortly, in our proposed AWG based network each node deploys both types of light sources. In the following, we explain how the two different light source signals are routed by the wavelength–selective AWG when fed into its input ports. Furthermore, we discuss possible signal interferences and whether the wavelength routing nature of the AWG requires that attached transmitters and receivers have to be tunable or not.

5.2.1 Underlying principles

Optical spectrum slicing

Without loss of generality, we consider a 2×2 AWG for explaining the spectral slicing of broadband signals. Fig. 5.1 depicts a scenario where six equidistant wavelengths are launched into the upper AWG input port. The wavelengths originate from six different laser diodes (nodes) which had to be attached to the AWG input port via a 6×1 combiner, as discussed in Section 5.2.2. The broadband signal has a wide spectrum of typically 10–100 nm covering one or more FSRs of the AWG. In our illustrative example, the broadband spectrum is assumed to span all six wavelengths. Fig. 5.1 shows that the AWG periodically routes every second wavelength to the same AWG output port. The AWG slices the broadband spectrum such that in each FSR one slice is routed to either AWG output port. Hence, by deploying a broadband light source control can be broadcast to all AWG output ports and attached receivers due to spectrum slicing. In general, using R FSRs of the underlying AWG there are R slices at each AWG output port, where $R \geq 1$. All these slices carry the same (control) information. Therefore, receivers attached to the AWG output ports are free to choose one of the R slices in order to retrieve the (control) information.

As shown in Fig. 5.1, the wavelengths and the broadband signal overlap spectrally. This enables in–band signalling where one receiver suffices in order to receive both a given wavelength and the corresponding slice of the original broadband signal. No additional receiver is needed leading to reduced network costs. However, both signals must be distinguished at the receiver. In the next section we discuss an approach which allows for receiving both signals simultaneously.

It can be seen in Fig. 5.2 that all wavelengths and broadband signals can be fed into both AWG input ports simultaneously without resulting in channel collisions at the AWG output ports. Thus, nodes attached to different AWG input ports can use the same set of wavelengths simultaneously. The resulting spatial wavelength reuse increases the degree of concurrency and improves the network efficiency. Note, however, that receiver collisions may occur. A given receiver is able to simultaneously obtain both signals (data and control), but both have to originate from the same AWG input port. Listening to a slice restricts the receiver to wavelengths that emanate from the same AWG input port as the slice while completely missing all remaining wavelengths and slices.

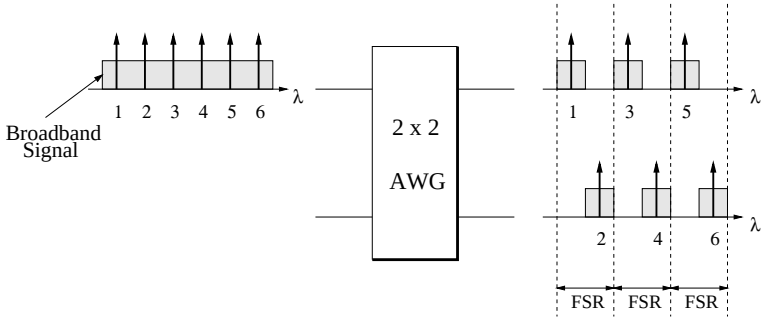

Figure 5.1: Spectral slicing of a broadband signal.

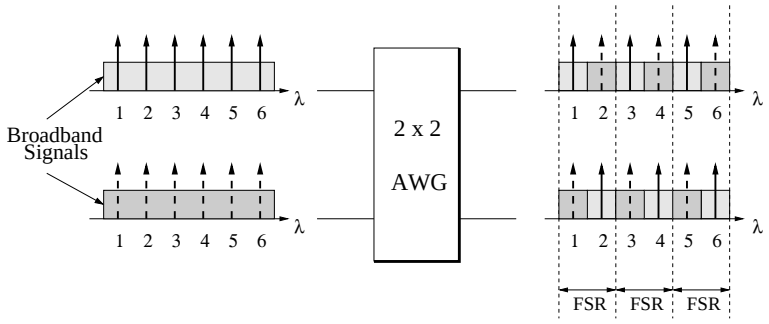

Figure 5.2: Spatial reuse of wavelengths and broadband signals.

Electrical spectrum spreading

The block diagram in Fig. 5.3 illustrates the simultaneous transmission and reception of a given wavelength and the corresponding slice within the same bandwidth interval. For data transmission we deploy a laser diode (LD). Control is broadcast by using a broadband light source. As shown in the figure, the data modulates the LD, whereas the control is spread prior to modulating the broadband light source (we will discuss the benefit of spreading shortly). The control is spread in the electrical domain by means of direct sequence spread spectrum (DSSS) techniques [GGHR98][GGR99]. Both data and control signals are combined and then routed through the AWG based network, which will be further detailed in the next section. At the output of the AWG based network a photodiode (PD) is tuned to the same wavelength as the LD. The PD detects the wavelength and the corresponding slice of the original broadband signal and converts the combined optical signal into the electrical domain. The resulting electrical spectrum is depicted in the lower right corner of Fig. 5.3. The modulation speed and launch power of the broadband signal is such that the control signal has (*i*) a smaller bandwidth and (*ii*) a smaller power level than the data signal. The spread control signal looks like a narrowband noise signal in the time domain whose power level is below that of the data [JM00]. Due to the small power level and the narrow bandwidth of the (spread) control signal the data signal is not distorted significantly and can be received without requiring any further processing (except from possibly some simple highpass filtering). To retrieve the control information a part of the combined data and (spread) control signal is lowpass filtered and subsequently despread. The despreading is done by a decorrelator which multiplies the filtered signal with the corresponding spreading sequence, followed by integration and sampling [GGHR98][GGR99]. In doing so, the power level of the control signal is raised above that of the data signal which allows for detecting the control information.

The spreading of control information has two advantages. First, the spread signal appears as noise and only nodes which have the correct spreading sequence are able to send and receive control information. This prevents malicious users from taking part in the control traffic, resulting in an improved network security [Muk97]. Second, by using more than one spreading sequence, i.e., code division multiple access (CDMA), new nodes can easily join the network. This makes the network scalable by deploying additional spreading sequences as the number of nodes increases. Moreover, by combining CDMA and WDMA the degree of concurrency and thereby the network efficiency can be increased.

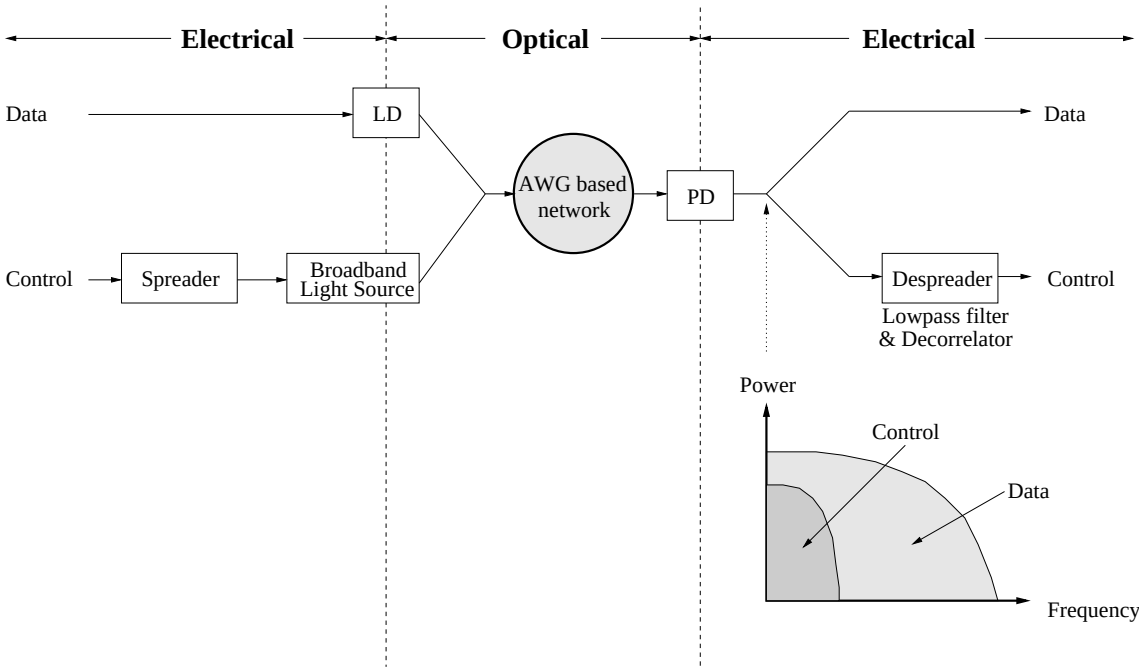

Figure 5.3: Spectral spreading of control information.

5.2.2 Network and node architecture

The proposed network architecture is schematically shown in Fig. 5.4. There are N nodes each attached to the AWG based network via two fibers, where $N \geq 2$. Every node uses one fiber for transmission and the other fiber for reception.

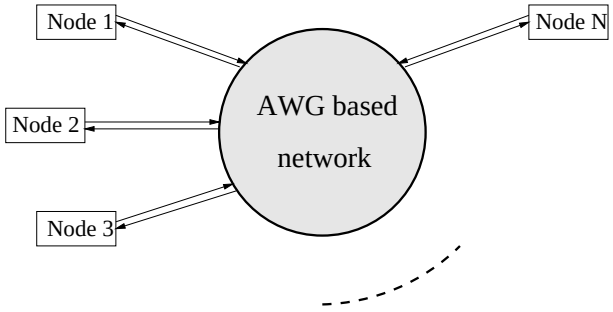

Figure 5.4: Network architecture.

Fig. 5.5 depicts the network and node architecture in more detail. The network is based on a $D \times D$ AWG, where $D \geq 2$. To each AWG input port a wavelength–insensitive $S \times 1$ combiner is attached, where $S \geq 1$. At each AWG output port signals are distributed by a wavelength–insensitive $1 \times S$ splitter. Apart from possibly required optical amplifiers the network does not contain any active devices (e.g., switches and wavelength converters) and is thus completely passive. As such, the network is reliable and the network operation, administration, and maintenance (OAM) is significantly simplified since all active components (nodes) are

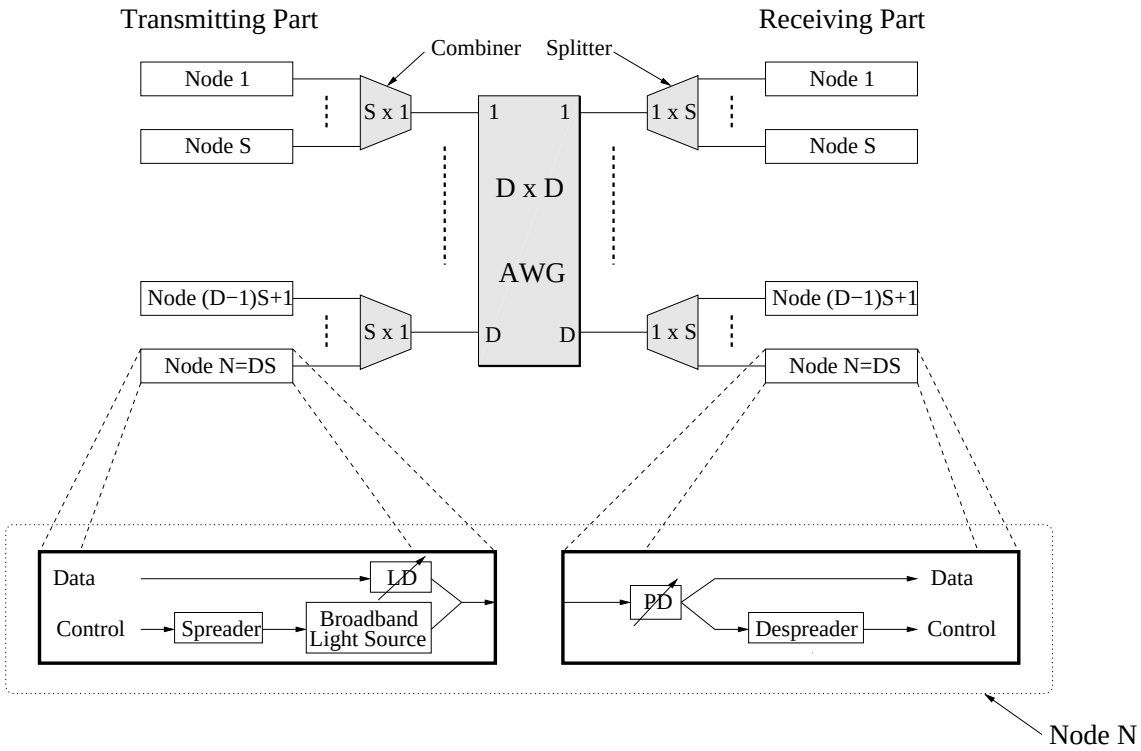

Figure 5.5: Detailed network and node architecture.

located at the network periphery. In addition, the transparent wavelengths provided by the passive network make the network flexible and future proof in that different legacy and future protocols can be supported. Recall from Section 4.2.2 that both splitters and combiners have to be wavelength insensitive in order to collect and distribute all packets from/to the attached S nodes independent of the wavelength. This also enables optical multicasting where a packet can be received by all nodes attached to the same splitter. The combiners and splitters provide additional ports for attaching multiple nodes to each AWG port such that more than one wavelength can be used at each AWG port simultaneously, assuming that each node is equipped with one single transceiver. (Alternatively, several transceivers attached to the same combiner/splitter could belong to one node. This node would then be able to send and receive data on several wavelengths simultaneously. Such a node equipped with multiple transceivers rather than one could operate as a server to cope with the large amount of local hot–spot traffic.) Combiners/splitters at different AWG ports do not necessarily have to have the same degree S. For example, while $(D-1)$ AWG ports accomodate the same number of nodes only one single node might be attached to the remaining AWG port. Moreover, network nodes might be added (or removed) dynamically at different AWG ports, resulting in different combiner/splitter degrees. Note that in the considered single–hop network nodes do not have to forward packets. Consequently, combiners/splitters can be replaced without interrupting the communication between nodes that are attached to the remaining combiners/splitters. Furthermore, node failures do not affect other nodes' communication. This makes the network tolerant against node failures. (However, the central AWG forms a single point of failure. We will address this problem in Chapter 9). Without loss of generality, in this work we focus on combiners/splitters with the

same degree S at all AWG ports. This has the positive side–effect that the splitting loss is the same between any arbitrary pair of nodes.

The network connects N nodes, with $N = D \cdot S$. For a given number of nodes N there are several possible configurations with different values of D and S. For instance, eight nodes can be connected via a 2×2 AWG with two 4×1 combiners and two 1×4 splitters, or via a 4×4 AWG with four 2×1 combiners and four 1×2 splitters. There are also cases, e.g., $N = 7$, where one or more ports are left unused. In these cases, the free ports can be used to attach additional nodes whenever needed. Note that the choice of D and S trades off spatial wavelength reuse and multicast efficiency. From the spectrum reuse point of view it is reasonable to choose a large D for a given N such that all wavelengths can be spatially reused at as many AWG ports as possible, at the expense of a larger number of crosstalk signals. On the other hand, for a given N a small D implies that more nodes are attached to the same splitter, i.e., S becomes large. This has the advantage that a given packet is distributed to more nodes. As a consequence, a given multicast packet has to be transmitted fewer times resulting in an improved bandwidth efficiency, at the expense of a larger splitting loss. If the splitting loss of the combiners and splitters becomes too large for increasing S, optical amplifiers, e.g., Erbium–doped fiber amplifiers (EDFAs), have to be inserted between the combiners/splitters and the corresponding AWG input/output ports [GZS01]. We will examine the different configuration options in detail when evaluating the performance of the network in Chapter 6.

Let us now take a look at the node structure. Each node is composed of a transmitting part and a receiving part. The transmitting part of a given node is attached to one of the combiner input ports. The receiving part of the same node is located at the opposite splitter output port. For cost reasons each node deploys only one single laser diode (LD) for data transmission and one single photodiode (PD) for data reception. Recall from Section 4.1.1 that owing to the wavelength routing characteristics of the AWG both transmitter and receiver have to be tunable over at least one FSR of the underlying $D \times D$ AWG (each FSR comprising D contiguous wavelengths) in order to provide full connectivity in one single hop. Let the tuning range Λ of each transceiver, i.e., both transmitter and receiver, be equal to R adjacent FSRs of the underlying $D \times D$ AWG, where $R \geq 1$. Giving the tuning range as the number of equidistant wavelengths that are used for communication, we can say that Λ comprises $D \cdot R$ contiguous wavelengths, i.e., $\Lambda = D \cdot R$. By using more than one FSR the degree of concurrency and thereby the network efficiency are increased since each FSR provides one additional communication channel between any pair of AWG input and output ports, provided the transceiver tuning range is wide enough. Note that the considered network is easily upgradable: Technologically advanced transceivers with larger tuning ranges (and possibly higher line rates) are able to exploit more FSRs of the underlying AWG without requiring any reconfiguration or upgrade of the network itself. (The transceiver tuning latency should be in the range of a few nanoseconds in order to allow for efficient packet switching. Until these devices become commercially available a viable interim solution might consist of two low–cost, off–the–shelf alternating tunable transceivers with a larger tuning time and tuning range at each node where one transceiver is being tuned to a different wavelength while the other one is busy. Thus, both transceivers mutually mask their tuning latency relaxing the tuning time requirement [TMS94].)

Apart from the tunable transceiver, each node uses a broadband light source for broadcasting control packets. As shown in Fig. 5.5, the control information is spread in the electrical domain before modulating the broadband light source. We have seen in Section 5.2.1 that spreading improves the network security. The modulated optical broadband signal is spectrally sliced by the AWG such that a slice of the original broadband signal is broadcast to all AWG output ports.

At the receiving part the control information is retrieved by despreading a part of the incoming signal in the electrical domain. Both electrical spectrum spreading and optical spectrum slicing were described in more detail in the previous section.

5.3 MAC protocol

A MAC protocol is required in the proposed network for the following three reasons:

- Normally, the network layer is responsible for packet switching. However, since we consider a single–hop network there are no intermediate nodes and alternative routes to choose from. Consequently, in our architecture the network layer is not present and packet switching has to be handled by the MAC sublayer [RS98].
- Remember that each node is equipped with a single tunable transmitter and a single tunable receiver for data transmission and reception, respectively. To allow for efficient packet switching fast tunable transceivers should be deployed. Owing to their limited tuning range there are more nodes than available wavelengths whose shared access has to be controlled by a MAC protocol.
- Due to the routing characteristics of the AWG each transceiver has to be tuned over at least one common FSR in order to provide full connectivity in one single hop (see Section 4.1.1). Hence, every wavelength is accessed by all nodes, again calling for a MAC protocol.

We first formally describe the proposed MAC protocol in Section 5.3.1 and explain its operation by means of an illustrative example in Section 5.3.2. In Section 5.3.3 we outline the rationale behind our protocol and discuss how the protocol meets the requirements listed above.

5.3.1 Protocol

The wavelength assignment at a given AWG port is schematically shown in Fig. 5.6. The y–axis denotes the wavelengths used for transmission and reception. As illustrated, R adjacent FSRs of the underlying $D \times D$ AWG are exploited. Each FSR consists of D contiguous wavelength channels. Transceivers are tunable over the range of $R \cdot D$ contiguous wavelengths. To avoid interferences at the receivers during simultaneous transmissions in different FSRs of the AWG, the FSR of the receivers has to differ from the FSR of the AWG. In our case, the FSR of the receivers is equal to $R \cdot D$ wavelengths. The x–axis denotes the time. Time is divided into cycles which are repeated periodically. Each cycle is further subdivided into D frames.

Nodes are assumed to be synchronized. One method to achieve slot synchronization among all nodes is described in [RS98] and works as follows. Besides the N nodes there is a synchronizer node that is located at the AWG. This node broadcasts a pulse called the sync pulse at the beginning of each cycle by using a broadband light source. Thus, the period of the sync pulses is one cycle. The time of reception of a pulse at a node is taken by the node as the start of a cycle in its receiver. The algorithm that each node uses to achieve cycle synchronization is as follows:

Synchronization Algorithm

1. Each node estimates, or predicts, the time of arrival of the next sync pulse at its receiver. The periodicity of the sync pulses is used in making this prediction quite accurate.

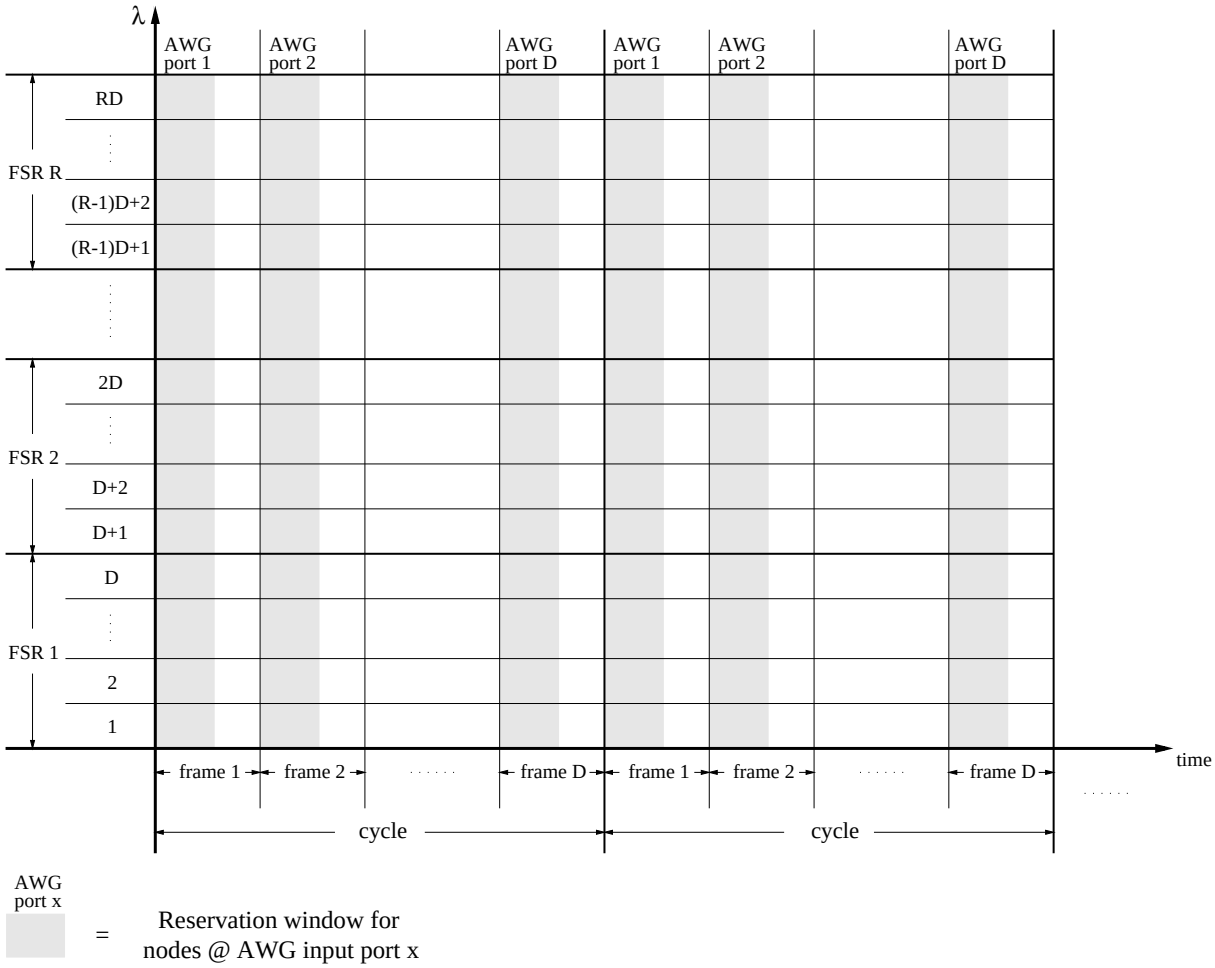

Figure 5.6: Wavelength assignment at a given AWG port.

2. Each node estimates its round–trip propagation delay to the AWG using the fact that the control signals sent by its broadband light source are broadcast by the AWG to it.

3. A node transmits the information for a specific cycle, say, cycle x, one round–trip propagation delay to the AWG prior to the estimated time of arrival of the sync pulse at the start of cycle x.

Once cycle synchronization is achieved, since the number of slots per cycle is known, each node can compute the slot times by suitably dividing the known cycle transmission. The periodic transmission of the start–of–cycle pulses by the synchronizer ensures that any drifts in the individual clocks of the nodes are compensated. Each node also estimates its round–trip propagation delay to the AWG while sending control packets. Thus, each node is able to compensate for changes due to temperature, aging, and similar factors.

The frame format on one wavelength is depicted in Fig. 5.7. A frame contains $F \in \mathbb{N}$ slots with the slot length equal to the transmission time of a control packet (function and format of a control packet will be explained later). If the tuning time of the transceivers is not negligible each slot has an additional transceiver tuning time interval besides the transmission time of a

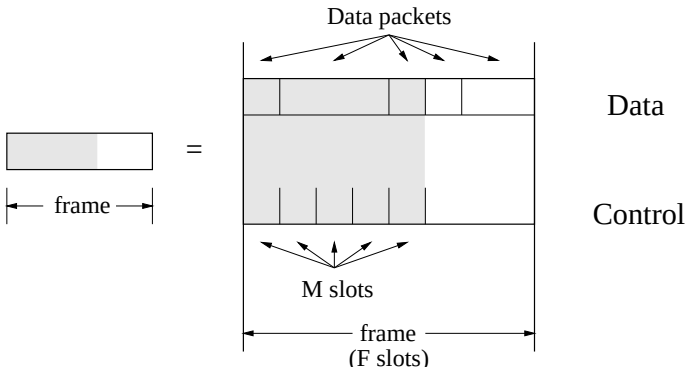

Figure 5.7: Frame format.

control packet. Each frame is partitioned into the first M, $1 \leq M < F$, slots (shaded region) and the remaining $(F - M)$ slots. In the first M slots the pretransmission coordination takes place. Control packets are transmitted during this period and all nodes are obliged to tune their receivers to one of the corresponding slices of the broadband light source (channels) in order to obtain the control information as explained in Section 5.2.1. Owing to the wavelength routing properties of the AWG, in a given frame only nodes that are attached to the same combiner can transmit control packets. Nodes attached to AWG input port i (via a common combiner) send their control packets in frame i of the cycle, where $1 \leq i \leq D$ (see Fig. 5.6). Each frame within a cycle accomodates control packets originating from a different AWG input port. Hence, after D frames (one cycle) all nodes have had the opportunity to send their control packets.

The M slots are not fixed assigned. Instead, control packets are sent on a contention basis using a modified version of slotted ALOHA. Using a random access scheme for control packets without fixed assigned reservation slots makes the entire network scalable. We have chosen a modified version of slotted ALOHA for two reasons. First, it is simple and cheap to implement. Second, and more importantly, ALOHA is independent of the ratio of propagation delay and packet transmission time. In very high–speed optical networks it outperforms other random access schemes such as carrier sense multiple acess (CSMA) [KT75]. Moreover, the line rate can be increased to 40 Gb/s or even higher data rates without degrading the throughput of the control channel. Control packets arrive at the receivers after the one–way end–to–end propagation delay τ that is equal to half the end–to–end round–trip time. Note that the round–trip time is identical to the propagation delay τ from the source to the destination and again back to the source. That is, the round–trip time equals 2τ. In the last $(F - M)$ slots of each frame no control packets are sent, allowing receivers to be tuned to any arbitrary wavelength. This freedom enables transmissions between any pair of nodes. During those slots each node processes the received control packets by executing the same scheduling algorithm. The parameter M trades off two kinds of concurrency. During the first M slots of each frame, control and data packets can be transmitted simultaneously, but only from nodes which are attached to the same AWG input port. In this time interval packets originating from other AWG input ports cannot be received. On the other hand, during the last $(F - M)$ slots of each frame all receivers are unlocked and can be tuned to any arbitrary wavelength. As a consequence, during this time interval data packets from any AWG input port can be received. This allows for spatial wavelength reuse.

The MAC protocol works as follows. First, we consider the transmitting part of a node whose flow chart is depicted in Fig. 5.8. If a node has no data packet in its buffer the broadband light

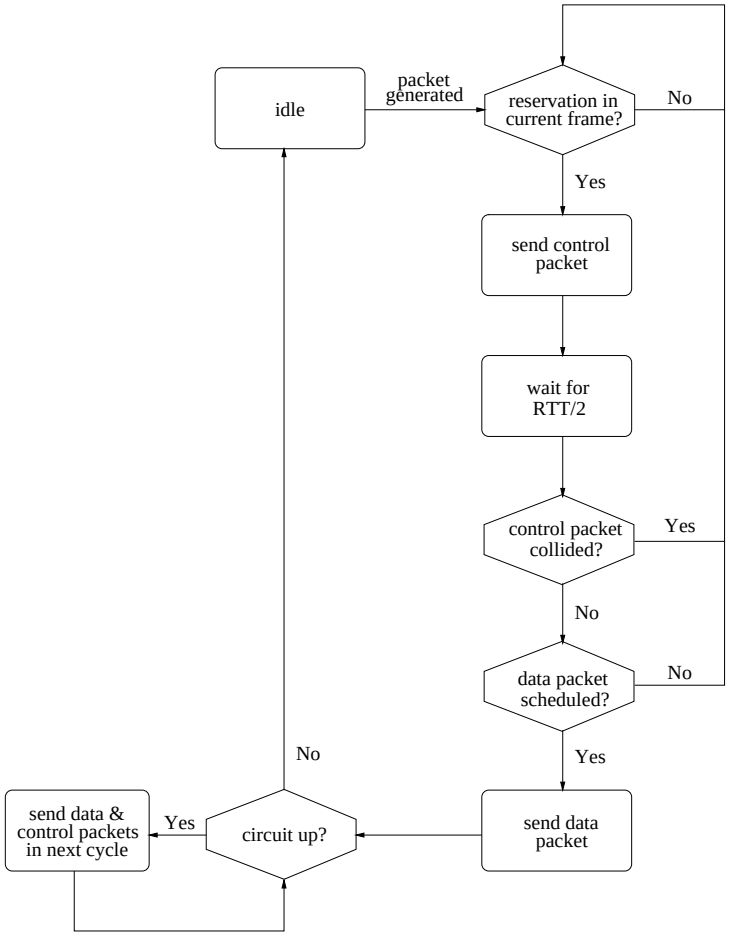

Figure 5.8: Flow chart of a node's transmitting part.

source and LD remain idle. When a data packet destined to node j, $1 \leq j \leq N$, arrives at node $i \neq j$, $1 \leq i \leq N$, node i's broadband light source broadcasts a control packet in one of the M slots of the frame allocated to the AWG input port that node i is attached to. The slot is chosen randomly according to a uniform distribution. A control packet consists of four fields, namely, destination address (unicast or multicast), length and type of the corresponding data packet, and forward error correction (FEC) code. Note that control packets do not have to carry the source address since each source node knows the slot in which it has transmitted the corresponding control packet. As illustrated in Fig. 5.7, the data packet can be of variable size $L, 1 \leq L \leq F$, where L denotes the length in units of slots. The type field contains one bit and is used to enable packet and circuit switching. The FEC is used by the receiver to correct one or more bit errors in the control packet or to find out if the control packet has experienced a collision on the modified slotted ALOHA channel. As we will see shortly, due to the distributed scheduling each control packet has to carry an FEC code which enables the receiver not only to detect errors but also to correct them. Otherwise, source node A could

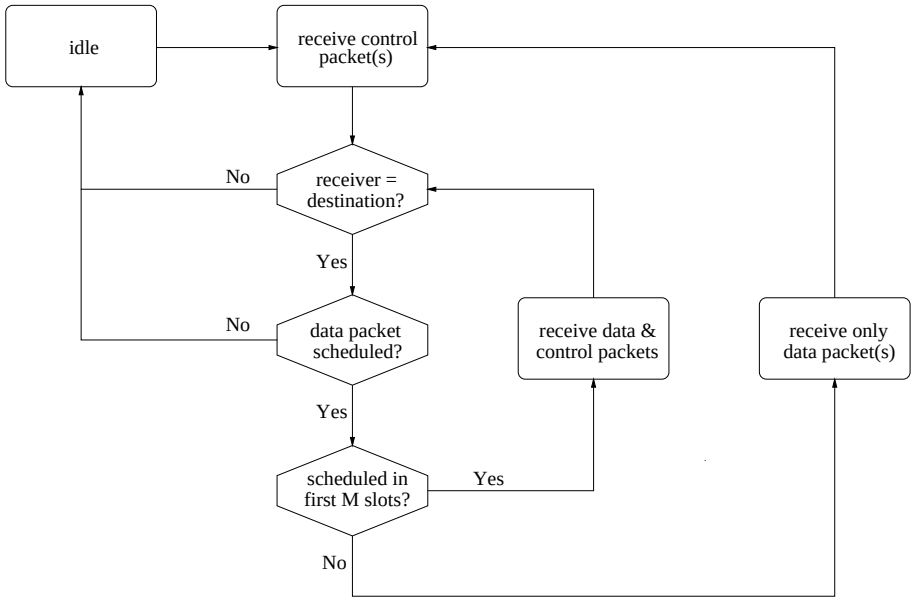

Figure 5.9: Flow chart of a node's receiving part.

receive an intact control packet while destination node B could obtain a damaged copy of the same control packet which B would not consider further for scheduling. As a consequence, node A would send the corresponding data packet and node B would most likely not listen on the appropriate wavelength resulting in a receiver collision and wasted bandwidth. Note that even though the transmission of packets is free of errors, malfunction and/or failure of nodes affect the operation of the distributed scheduling protocol, too. This issue has to be addressed by higher–layer protocols.

Let us now take a look at the receiving part of a node. Fig. 5.9 shows the corresponding flow chart. Every node collects all control packets by tuning its receiver to one of the corresponding channels during the first M slots of each frame. Thus, it learns about all other nodes' activities and whether its own control packet was successful or not by using the FEC field. In frame k, $1 \leq k \leq D$, each receiver collects the control packets which have been sent half a round–trip time ago by nodes that are attached to AWG input port k. If its control packet has collided node i backs off and retransmits the control packet in the next cycle with probability p and with probability $(1-p)$ it will defer the transmission by one cycle, where $0 \leq p \leq 1$. The node retransmits the control packet in this next cycle with probability p, and so forth. Successful control packets are put in a distributed queue at each node.

All nodes process the successfully received control packets by executing the same arbitration (scheduling) algorithm in the last $(F-M)$ slots of each frame. Consequently, all nodes come to the same transmission and reception schedule. Since each node has to process the control packets of all nodes the computational complexity at each node puts constraints on the network scalability. A simple arbitration algorithm is required to relax those constraints [Mod98b]. Therefore, we apply a straightforward greedy algorithm which schedules the data packets on a first–come–first–served and first–fit basis within a finite scheduling window. Note that this algorithm together with the fact that each node randomly selects one of the M reservation slots according to a uniform distribution in a periodic cycle time structure enables fair wavelength

access among all nodes ready to send data packets. After receiving a successful control packet the arbitration algorithm tries to schedule the transmission of the corresponding data packet within the following D frames. Those D frames do not necessarily have to coincide with the cycle boundaries. The data packet is sent in the first possible slot(s) using the lowest available wavelength. If there are not enough slots available within the D frames the data packet is not transmitted and the source node has to retransmit the control packet in the next cycle. Nodes which lose the arbitration are aware of this because all nodes execute the same scheduling algorithm. Note that global knowledge in conjunction with distributed scheduling reduces the delay by avoiding explicit acknowledgements (ACKs) and can achieve a normalized throughput of up to 100%. Since each node has to process the control packets of all nodes to acquire global knowledge the reservation computational overhead can become a serious bottleneck that affects the network scalability. In order to accomodate a large number of nodes and make the entire network scalable it is important to keep the computational complexity at each node small [RS98]. Therefore, in the following we make use of one single spreading sequence without making use of CDMA. This spreading sequence is applied by all nodes. (As a positive side effect, deploying one single spreading sequence also reduces the crosstalk penalty. In Chapter 8 we will get back to CDMA and examine how CDMA is able to improve the throughput–delay performance of our network.)

The length of the scheduling window is equal to D frames for two reasons. First, by limiting the scheduling length to a small number of frames the computational requirements at each node are kept small. Due to the relatively small scheduling window each node has to maintain and update only small schedule tables. Second, every D frames all nodes receive control packets from the same set of nodes. At the same time, due to the wavelength routing properties of the AWG and the requirement that all nodes listen to the control slices, only this set of nodes can transmit data packets. Those data packets were announced by control packets exactly D frames earlier. By making the scheduling window D frames long, data and control packets can be sent simultaneously resulting in an increased degree of concurrency and network efficiency.

Next, we discuss the support for multicasting and circuit switching. Multicasting is realized by the splitters. Each splitter distributes an incoming packet to all attached nodes. By tuning the receivers to the respective wavelength the packet can be obtained by more than one node. The resulting increased receiver throughput has a positive impact on the network performance. Circuit switching is realized by using the type and length fields of the control packet. The length field denotes the required number of slots per cycle. By setting the bit in the type field the source node indicates that this number of slots must be reserved in each cycle. After receiving the control packet the circuit is set up by choosing the first possible free slot(s) at the lowest available wavelength. The slot(s) is (are) reserved in the subsequent cycles until the connection is terminated. The termination of a circuit works as follows. Suppose node i, $1 \leq i \leq N$, has set up a circuit, i.e., node i is granted a certain number of slots per cycle which was specified in the foregoing control packet. Furthermore, suppose j, $1 \leq j \leq (M-1)$, other nodes attached to the same combiner currently hold circuits. Then, in each cycle node i repeats the control packet in slot $(j+1)$ of the corresponding reservation window. To terminate the circuit, node i simply stops repeating the control packet. In doing so, all other nodes notice that the circuit has terminated and the respective slot is freed up for contention. Note that during the holding time of a circuit other circuits can be torn down. As a consequence, the corresponding slot, say k, $1 \leq k \leq (j+1)$, becomes idle. Whenever this happens, all slots with an index larger than k that are used to indicate the existence of circuits are decremented by one. Thus, the first j slots of the corresponding reservation window indicate the existence of circuits while the remaining

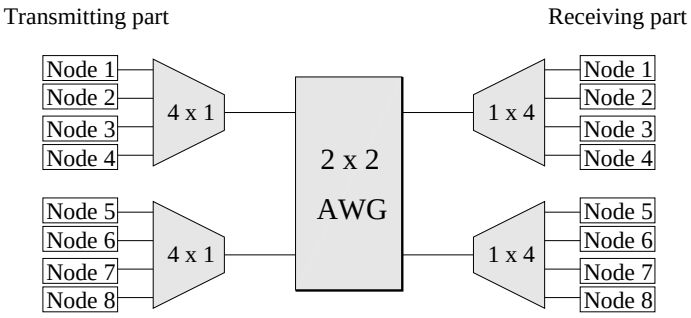

Figure 5.10: Network architecture ($N = 8$, $D = 2$, $S = 4$).

$(M - j)$ slots are free to be used for reservations. A node with a control packet to send chooses one of slots $(j + 1), (j + 2), \ldots, M$ at random. Consequently, while circuits are set up not all M slots are available for reservation resulting in an increased congestion of the modified slotted ALOHA channel. However, each node only has to monitor the control channel to find out the end of a circuit and does not have to maintain and update a lifetime variable for each single connection. Again, this reduces the computational burden on each node which is an important factor especially in high–speed optical networks where channel access control is based on global knowledge.

Finally, we point out that the proposed reservation and circuit set–up is able to provide guaranteed quality of service (QoS) to delay/jitter–sensitive traffic and real–time applications, such as voice, video, and audio. Circuits could also provide QoS to individual flows (following the IntServ paradigm [BCS94]) or flow aggregates (following the DiffServ paradigm [BBC+98]).

5.3.2 An illustrative example

In this section, we illustrate the following features of the proposed MAC protocol:

- Dynamic channel allocation
- Unicast packet switching
- Unicast circuit switching
- Channel collision of control packets and its resolution
- Simultaneous reception of control and data packets
- Variable–size data packets
- Using multiple FSRs

In the example we do not demonstrate spatial wavelength reuse.

As illustrated in Fig. 5.10, we consider $N = 8$ nodes which are connected via a 2×2 AWG ($D = 2$) with attached 4×1 combiners and 1×4 splitters ($S = 4$). Fig. 5.11 depicts the wavelength assignment. Two FSRs are used ($R = 2$), each consisting of two wavelengths. Each frame consists of $F = 5$ slots. The reservation window is $M = 3$ slots long. The upper part shows the transmitters while the lower part indicates the reception of the transmitted packets after a propagation time τ. For the sake of simplicity, we assume that the distance between

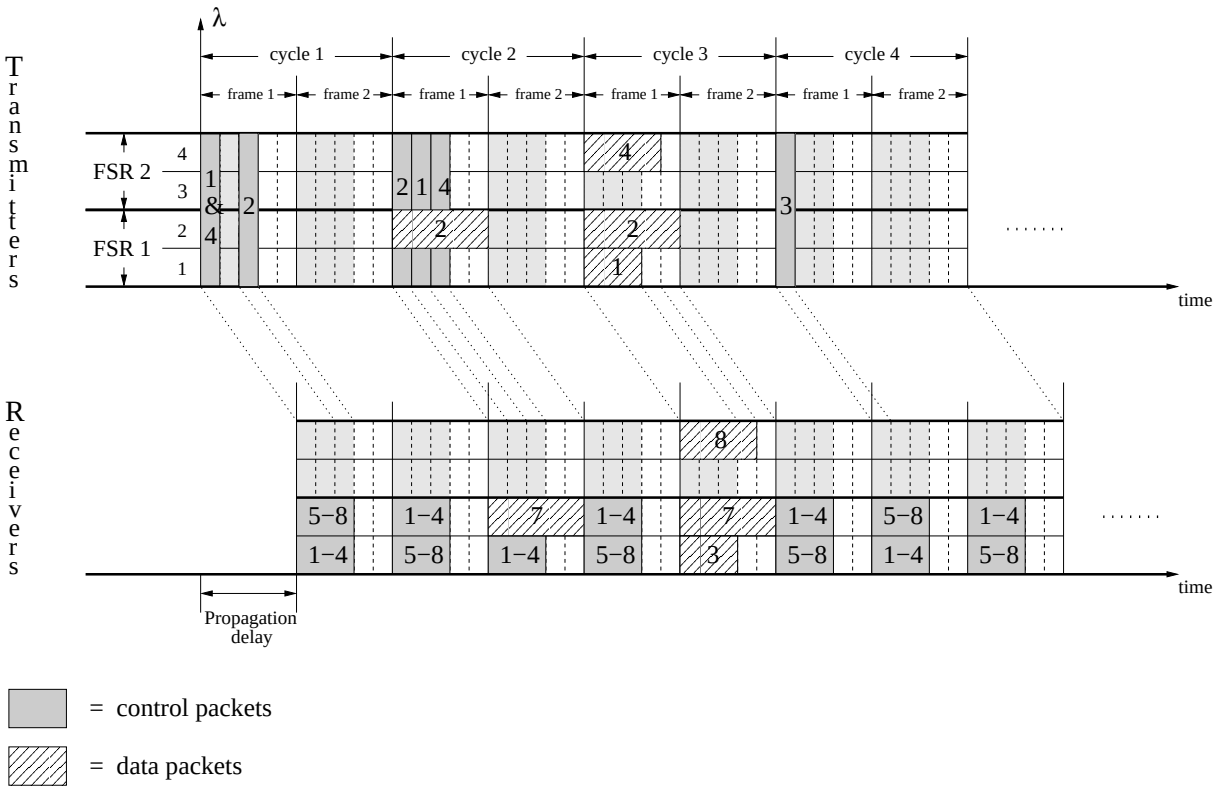

Figure 5.11: Dynamic wavelength allocation ($R = 2$, $F = 5$, $M = 3$, $\tau = 5$ slots = 1 frame).

each node and the AWG is equal, i.e., the propagation delay τ is the same for all nodes. In our example τ is equal to 5 slots, i.e., 1 frame.

Let us start considering the left–most slot of the transmitting part. In frame 1 of cycle 1 only nodes $1-4$ are permitted to send control packets. In the first slot nodes 1 and 4 simultaneously transmit a control packet resulting in a channel collision. Node 2 randomly selects slot 3. After τ slots the control packets arrive at all nodes. Every node has tuned its receiver such that control packets can be obtained, i.e., nodes $1-4$ and nodes $5-8$ are alternately tuned to wavelengths 1 and 2, respectively. Note that nodes could also receive the control packets by tuning their receivers to the corresponding channels of the other FSR, i.e., nodes $1-4$ could tune their receivers to wavelength 3, and nodes $5-8$ could tune their receivers to wavelength 4. The control packet of node 2 is received successfully by all nodes. In our example, that control packet is destined to node 7 and requests a circuit with 5 slots per cycle. The execution of the arbitration algorithm is assumed to take 2 slots. The data packet is transmitted on the first available channel at the earliest possible time. Accordingly, the data packet is sent on wavelength 2 during frame 1 of cycle 2.

Node 2 repeats the control packet in each cycle until the circuit is torn down. Since in our example there are no other circuits currently set up, node 2 sends the control packet in the first slot of frame 1 of cycle 2. As illustrated in Fig. 5.11, slot 1 becomes idle in cycle 3. This tells all nodes that the circuit between nodes 2 and 7 is terminated and that this slot can be used again by all nodes. Node 3 captures that slot in cycle 4 to announce the transmission of a single packet or to set up a circuit in the subsequent cycle.

Now, let us go back to the initially collided control packets of nodes 1 and 4. After τ slots both nodes learn about the channel collision of their control packets. As a consequence, they retransmit their control packets in the next cycle with probability p keeping in mind that the first slot of frame 1 is fixed assigned to node 2 for indicating the corresponding circuit. In our example, nodes 1 and 4 successfully retransmit their control packets in frame 1 of cycle 2. First, we take a look at node 1. Node 1 has one single data packet which is 3 slots long and is destined to node 3. After waiting for 3 slots, node 1 sends the data packet on wavelength 1. Similarly, node 4 has a data packet for node 8 which is 4 slots long. This data packet is sent at the beginning of frame 1 of cycle 3. Note that node 4 chooses wavelength 4 since wavelength 2 is already used by node 2. This is an example for using multiple FSRs for communication between a given AWG input–output port pair. In addition, while receiving the data packets nodes 3, 7, and 8 monitor the control channel. Node 8 thereby receives the corresponding slice of the second FSR. Recall that this is possible due to spreading and spectrally slicing the broadband signal.

5.3.3 Discussion

Since in the proposed AWG based network both transmitter and receiver of each node are tunable it is advisable to control the wavelength access by means of pretransmission coordination. Prior to transmitting data a given node has to make a reservation by using its broadband light source. The reservation process is a hybrid of random (modified slotted ALOHA) and deterministic (TDMA) protocols. By using TDMA receiver collisions of control packets are prevented. All receivers are tuned such that they are able to monitor the entire control traffic. Note that the applied type of TDMA is rather flexible and efficient. Receivers have in each FSR of the underlying AWG one slice to tune to for obtaining the control information. Moreover, while monitoring the control traffic nodes are also able to receive data albeit only from source nodes attached to the corresponding combiner. While round–robin TDMA provides all groups of nodes (each group attached to a different combiner) equal opportunity to make reservation requests, ALOHA is well suited to support random reservation traffic efficiently. Unlike in reservation protocols where each reservation slot is fixed assigned to a different node, every reservation slot in our network can be equally used by all nodes of the same group which wish to send a control packet [SS00b]. In doing so, bandwidth is not wasted by idle nodes and in conjunction with the above mentioned round–robin cyclic timing structure wavelengths are accessed in a fair fashion by nodes ready to send data packets. Using ALOHA for control lets new nodes join the network anytime without any MAC protocol reconfiguration and makes the network scalable. Moreover, since ALOHA is independent of the propagation delay the network is upgradable to higher line rates without decreasing the reservation throughput. However, ALOHA inherently suffers from a relatively small achievable throughput. In our network the throughput of the ALOHA based control channel is increased by two factors: (*i*) To indicate the duration of a given circuit we deploy reservation ALOHA (R–ALOHA), i.e., one ALOHA slot is dedicated to a given source node for signalling to the remaining nodes that the corresponding circuit is still up. This slot must not be used by the remaining nodes and is therefore completely free of collisions resulting in an increased throughput. (*ii*) By using multiple spreading sequences (CDMA) multiple ALOHA control channels are created which work in parallel. The control load can be distributed among these channels (load balancing) resulting in fewer collisions and a larger aggregate throughput, especially at higher (control) traffic loads (as we will see in Chapter 8).

For now, we deploy only one single spreading code in order to keep the system complexity

low. Due to the in–band signalling no additional control channel (wavelength) and receiver are required as opposed to other network designs [LA95]. Consequently, each wavelength is used more efficiently and costs are reduced, which are important issues in metro networks. Conversely, previously investigated architectures with a single tunable receiver at each node suffer from the problem that while receiving control packets no data packets can be received and vice versa resulting in a decreased throughput–delay performance [Mod98a][JM93c]. We have already seen in Section 5.2.1 that spreading of the control information improves the network security. Only nodes which know the spreading sequence are able to participate in the reservation. Since the spread control looks like a low–power noise signal it is difficult for malicious users to detect the spreading sequence. Furthermore, the control information (e.g., MAC destination address) is not present in the corresponding data packet(s). Thus, it is difficult to find out to which source–destination pair a given data packet belongs [DDR98][DR00].

Since the broadcast control traffic is received by all nodes, each node is able to acquire global knowledge. No explicit acknowledgements are required resulting in a decreased response time and an increased bandwidth utilization. Based on this global knowledge all nodes execute the same distributed scheduling algorithm. Since all nodes come to the same conclusion, channel and receiver collisions of data packets can be avoided completely, leading to an improved wavlength utilization. The network is flexible since data packets can have variable size. In addition, through circuit switching guaranteed QoS can be provided. Note that the proposed reservation protocol provides a flexible framework into which specific features can be included. For example, additional information can be sent in each control packet such as the packet age or priority, and different arbitration rules can be devised to conform to specific requirements that include assigning higher priorities to certain types of traffic and certain nodes enabling service differentiation. Different metrics for fairness and optimality can be incorporated. For example, the arbitration algorithm can try to minimize the maximum packet delay [CDR90][LK92a].

5.4 Conclusions

The proposed network architecture consists of combiners and splitters which are attached to a central AWG. Recently, it was shown that athermal polarization–independent AWGs can be realized [KYZ+01]. Such an all–polymer AWG (polymer waveguides on a polymer substrate) is temperature insensitive and does not require wavelength stabilization through electronic control circuits. Due to its passive nature the architecture is reliable. The architecture combines the merits and mitigates the drawbacks of wavelength–routing and wavelength–insensitive components. The AWG not only allows for spatial wavelength reuse but also improves the power budget due to the lack of splitting loss. This allows for realizing single–hop networks with an increased diameter. On the other hand, splitters can be used to achieve bandwidth–efficient multicasting. In the next section, we will take a look at the trade–off between spatial wavelength reuse and multicasting which largely depends on the AWG degree D.

Since all intelligence is located at the network periphery the operation, administration, and maintenance (OAM) of the network is considerably simplified. In addition, our single–hop network not only provides full connectivity but is also immune from node failures since nodes are not involved in packet forwarding, resulting in an improved network survivability. However, the AWG represents a single point of failure, i.e., all nodes are disconnected when the central AWG fails [HBP+98]. We will return to this problem when addressing network protection in Chapter 9.

The network is very flexible and future–proof. The reservation protocol creates transpar-

ent channels between any arbitrary source node and one or more destination nodes. These all–optical connections are able to transport a wide range of present or future heterogeneous protocols at different rates using different modulation formats which can be freely chosen by the corresponding source–destination pair. The network supports data packets with variable size, e.g., IP datagrams. Moreover, the inherent transparency makes the single–hop network easily upgradable. Nodes can deploy technologically advanced transceivers which offer higher data rates and smaller tuning times. Transceivers with an increased tuning range can use additional FSRs of the underlying AWG without needing any upgrade of the (passive) network itself [LRI00][FWLB01]. Due to the fact that the reservation slots are not fixed assigned new nodes can be added (removed) to (from) the network without protocol operation disruption and reconfiguration, making the network scalable.

The reservation is done by using a spectrally sliced broadband signal. Spreading the control information helps provide network security. Due to in–band signalling it is sufficient to equip each node with one single receiver. Only one single spreading sequence is used and the applied scheduling algorithm is kept simple in order to reduce the computational burden at each node. In conjunction with round–robin TDMA and random reservation slot access, the first–come–first–served and first–fit arbitration algorithm yields fair wavelength assignment among all nodes. Furthermore, circuit set up provides guaranteed QoS for supporting real–time applications. Note that in our network traffic grooming can be done in a distributed fashion. To achieve this, each node applies the MAC protocol and independently sends the corresponding data packets to that node which typically functions as the point of presence (PoP) connecting the metro network to the backbone.

In summary, the proposed network allows for very efficient wavelength channel utilization. The degree of concurrency is significantly increased by simultaneously transmitting control and data within the same channel, deploying multiple FSRs of the underlying AWG, and allowing for spatial wavelength reuse at all AWG ports. All wavelengths are efficiently utilized for data transmission and no bandwidth is wasted due to channel and receiver collisions of data packets, receiver collisions of control packets, and explicit acknowledgements. No wavelength is set aside for control traffic which is advantageous, especially in packet switched WDM networks using fast tunable transceivers whose tuning range allows only for a small number of accessible wavelengths. Wavelengths are dynamically allocated only to nodes which have data to send. Note that the network aims at combining the best of electronics and optics. While storing and protocol processing are done in the electrical domain, transmission takes place optically. The photonic fast–tuning of transceivers can be considered somewhere in between. Finally, we note that the number of nodes in metro networks is rather modest. As a consequence, the total amount of broadcast control traffic is limited and is expected to be on–line processed by each node in real–time. Moreover, due to the modest population size of metro networks the splitting loss in our network can be kept sufficiently small in order to cover metropolitan areas (as we will see in Chapter 8).

The proposed network is characterized by a collection of architecture and protocol parameters which are listed in Table 5.1. The impact of these parameters on the network performance is investigated in detail in the following chapter.

Network parameters

Architecture parameters	
N	Number of nodes
D	Degree of AWG
S	Degree of combiners/splitters ($S = N/D$)
Λ	Transceiver tuning range (number of equidistant wavelengths)
R	Number of used AWG FSRs ($R = \Lambda/D$)
τ	End–to–end propagation delay
Protocol parameters	
F	Number of slots in one frame
M	Number of reservation slots per frame
L	Length of data packet in slots ($1 \leq L \leq F$)
p	Retransmission probability of unsuccessful control packets

Table 5.1: Architecture and protocol parameters.

Chapter 6

Performance Evaluation

After describing the network at length in the previous chapter, we investigate the impact of various network architecture and protocol parameters on the network performance by means of analysis and/or simulation in this chapter. For tractability reasons we assume Bernoulli packet arrival processes in our analyses. We note that Bernoulli traffic is not an appropriate traffic model for LANs and WANs, as shown in [LTWW94] and [PF95], respectively. But at the time of writing it is an open question which traffic model applies in MANs. While the simulations in this chapter make also use of Bernoulli traffic, in Chapter 8 we provide additional simulations which are driven by packet header trace files that were recorded at the metro level. The performance evaluation proceeds in two steps. First, we consider different network aspects separately from each other in order to provide insight in how they influence the network performance. In Section 6.1, we focus on switching of *fixed–size* packets and show the positive impact of using multiple FSRs of the underlying AWG on the throughput–delay performance of the network. Next, in Section 6.2 we consider *variable–size* packets and demonstrate the benefit of spatial wavelength reuse for improving the network flexibility and efficiency. Section 6.3 takes also *multicasting* into consideration. We show that our network is able to support multicasting efficiently and we examine the interplay between multicast and unicast packet switched traffic. Second, in Section 6.4 we investigate the network performance while allowing the aforementioned features to occur simultaneously. We thereby try to better understand the interplay of the various architecture and protocol parameters. By means of extensive simulations we investigate additional network aspects, e.g., self–stability and circuit switching, while relaxing several assumptions made in our analyses. Finally, in Section 6.4.4 we discuss the efficiency of our network and compare its performance with a previously reported reservation MAC protocol designed for a PSC based single–hop metro WDM network.

6.1 Using multiple FSRs

In this section, we show the benefit of using multiple FSRs of the underlying AWG. Using multiple FSRs is motivated by the scalability and upgradability of the considered AWG based network. For an increasing number of nodes it is mandatory to provide more wavelengths by exploiting multiple FSRs in order to cope with the higher amount of traffic. Moreover, additional FSRs can be used after upgrading nodes with technologically advanced transceivers which offer a wider tuning range. In the following we develop a Markovian analytical model. We thereby focus on uniform unicast packet switching without spatial wavelength reuse. Since we do not consider circuit switching in this section, all reservation slots can be used for pretransmission

coordination in each frame. The assumptions made in Section 6.1.1 are verified by simulation in Section 6.1.4 [MRW00][MRW03].

6.1.1 Assumptions

In our analysis we make the following assumptions:

- Each node has a *single–packet* buffer, i.e., each node can store at most one data packet at any given time. (This assumption simplifies the analysis of MAC protocols for WDM networks [JM93a].)
- After transmitting a data packet in a given frame the buffer becomes empty at the end of that frame.
- A node with an empty buffer generates a data packet with probability σ at the end of a frame.
- A data packet has a *fixed size* of F slots, i.e., $L = F$.
- *Uniform unicast* traffic: A data packet is destined to any one of the other $(N-1)$ nodes with equal probability $1/(N-1)$.
- The propagation delay τ is the same for all nodes and is an integer multiple of one frame, i.e., all nodes are *equidistant* from the AWG.
- *Nonpersistency*: Random selection of a destination node among the other $(N-1)$ nodes is renewed for each attempt of transmitting a control packet. (The nonpersistency assumption is needed to obtain a Markovian model [LK92a].)
- *Delayed first–time transmission*: A node sends out its control packet in a frame with probability p, not only for retransmissions but also for first–time transmissions. (This assumption simplifies the calculation of the probability of control packet collisions [JM93a].)

6.1.2 Model

Fig. 6.1 depicts a model of the MAC protocol. Each node can be in one of the $(2\tau + 3)$ modes during any frame. Transitions from one mode to another mode occur only at the beginning of a frame. The modes are defined as follows:

- TH: Nodes in the TH (thinking) mode generate a data packet with probability σ at the end of a frame.
- B: Nodes in this mode are backlogged and send a control packet with probability $p \cdot \beta$ (where β accounts for the cycles in the time structure, as is explained later) at the beginning of the next frame.
- $PQ_1, PQ_2, \ldots, PQ_\tau$: These modes represent the propagation delay of successfully transmitted control packets, i.e., control packets which have not collided in the reservation slots. Nodes move from mode PQ_i to mode PQ_{i+1}, $i = 1, 2, \ldots, (\tau - 1)$, at the beginning of the next frame with probability 1. (Note that the number of PQ modes is equal to the propagation delay given in frames since mode transitions occur on a per–frame basis. The same holds for the following PR modes.)

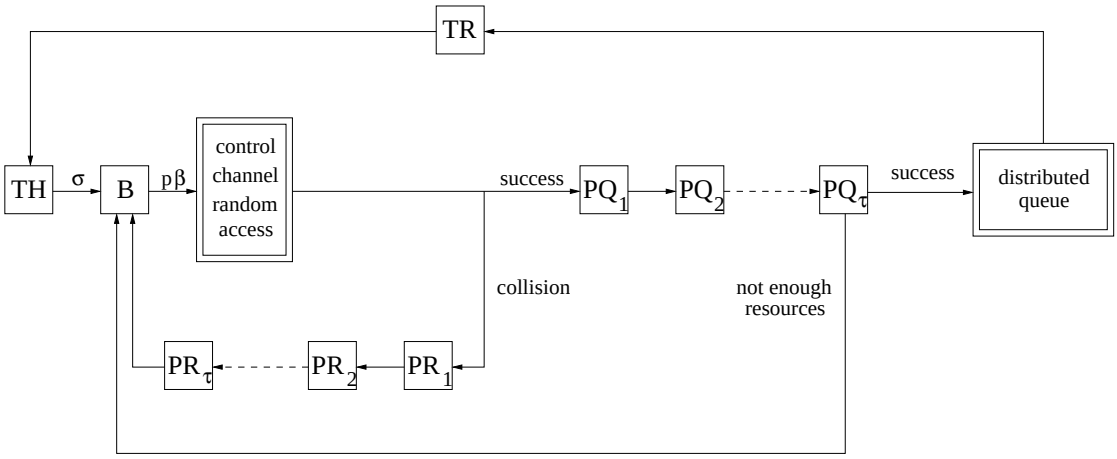

Figure 6.1: Model of MAC protocol.

- $PR_1, PR_2, \ldots, PR_\tau$: These modes are similar to the PQ_i modes, $i = 1, 2, \ldots, \tau$. Nodes whose collided control packets have to be retransmitted enter the mode B after τ frames.
- TR: Nodes in the mode PQ_τ whose data packets are successfully scheduled are put into a distributed queue. A node leaves the distributed queue and moves to mode TR (transmission) in the frame in which it sends control packets. After transmitting the data packet the nodes return to the TH mode. Nodes in mode PQ_τ whose data packets are not scheduled due to the lack of free resources (not enough free slots and/or wavelengths) move to mode B.

The system state in frame n, $n \in \mathbb{Z}$, is completely described by the following state vector:

$$\mathbf{N}(n) = (N_B(n), N_{PR_1}(n), \ldots, N_{PR_\tau}(n), N_{PQ_1}(n), \ldots, N_{PQ_\tau}(n), N_{TR}(n)),$$

where $N_X(n)$ denotes the number of nodes in mode X in frame n. Note that N_{TH} is not included in the state vector since it is linearly dependent on the other modes. With the nonpersistency assumption

$$\{\mathbf{N}(0), \mathbf{N}(1), \ldots, \mathbf{N}(n), \ldots\}$$

is a discrete–time multi–dimensional Markov chain with finite but quite large state space. The exact analysis of that Markov chain would involve the calculation of the state transition probability matrix which is computationally prohibitive. Therefore, we analyze the system at an equilibrium point using the equilibrium point analysis (EPA) approach [FT83].

6.1.3 Analysis

In the EPA method the system is assumed to be always at an equilibrium point, defined as

$$\mathbf{N} = (N_B, N_{PR_1}, \ldots, N_{PR_\tau}, N_{PQ_1}, \ldots, N_{PQ_\tau}, N_{TR}).$$

At an equilibrium point the expected increase in the number of nodes in each mode per unit time (i.e., frame in our analysis) is zero. Applying this condition to all the modes, we get a set of so–called equilibrium point equations.

Equilibrium point equations

By writing the equation for each mode, we get $(2\tau+3)$ equations. Let $\delta_X(\mathbf{N})$ be the conditional expectation of the increase in the number of nodes in mode X in a frame, given that the system is in state $\mathbf{N}$. Since $\delta_{PR_i}(\mathbf{N}) = N_{PR_{i-1}} - N_{PR_i} = 0$, $i = 2, 3, \ldots, \tau$, we can omit the subscript of PR by letting

$$N_{PR} = N_{PR_1} = \cdots = N_{PR_\tau}. \tag{6.1}$$

Similarly, for the modes PQ_i, $i = 1, 2, \ldots, \tau$, we get

$$N_{PQ} = N_{PQ_1} = \cdots = N_{PQ_\tau}. \tag{6.2}$$

For the modes TH and B we have the following equations:

$$\begin{aligned} \delta_{TH}(\mathbf{N}) &= N_{TR} - N_{TH}\sigma \\ &= N_{TR} - [N - N_B - \tau(N_{PR} + N_{PQ}) - N_{TR}]\,\sigma = 0 \qquad (6.3) \\ \delta_B(\mathbf{N}) &= [N_{TH}\sigma + (N_{PQ} - N_{TR}) + N_{PR}] - N_B p\beta = 0. \qquad (6.4) \end{aligned}$$

To obtain the remaining equilibrium point equations for the modes PR_1, PQ_1, and TR we introduce the quantities $Y(\mathbf{N})$ and $Z(\mathbf{N})$. Let $Y(\mathbf{N})$ denote the conditional expectation of the number of nodes that move from mode B to mode PQ_1, given that the system is in state $\mathbf{N}$. $Y(\mathbf{N})$ is the average number of control packets transmitted in a frame without collision. With $Y(\mathbf{N})$ we obtain the following equations for the modes PR_1 and PQ_1:

$$\begin{aligned} \delta_{PR_1}(\mathbf{N}) &= N_B p\beta - Y(\mathbf{N}) - N_{PR} = 0 \qquad (6.5) \\ \delta_{PQ_1}(\mathbf{N}) &= Y(\mathbf{N}) - N_{PQ} = 0. \qquad (6.6) \end{aligned}$$

Let $Z(\mathbf{N})$ denote the conditional expectation of the number of nodes that move from mode PQ_τ to mode TR, given that the system is in state $\mathbf{N}$. $Z(\mathbf{N})$ is the average number of nodes that successfully transmit a data packet in a frame. With $Z(\mathbf{N})$ we obtain the following equation for the mode TR:

$$\delta_{TR}(\mathbf{N}) = Z(\mathbf{N}) - N_{TR} = 0. \tag{6.7}$$

Next, we need to solve for the unknown quantities β, $Y(\mathbf{N})$, and $Z(\mathbf{N})$. Recall that a backlogged node, i.e., a node with a data packet in its buffer, can transmit a control packet only in one frame per cycle that consists of D frames. Let β denote the probability that the next frame is allocated to the backlogged node. Thus, we get

$$\beta = \frac{1}{D}. \tag{6.8}$$

The average number of successfully transmitted control packets per frame is given by[JM93a]

$$\begin{aligned} Y(\mathbf{N}) &= \sum_{i=1}^{N_B} i \left(1 - \frac{1}{M}\right)^{i-1} \binom{N_B}{i} (p\beta)^i (1 - p\beta)^{N_B - i} \qquad (6.9) \\ &= N_B p\beta \left(1 - \frac{p\beta}{M}\right)^{N_B - 1}. \qquad (6.10) \end{aligned}$$

The result can be interpreted such that $p\beta\left(1 - \frac{p\beta}{M}\right)^{N_B-1}$ is the probability that a node's control packet is transmitted collisionfree. The average number of nodes that move from mode B to mode PQ_1 is given by Eqn. (6.10).

Let q be the probability that a given slot of the first M slots of a frame contains exactly one control packet that is to be scheduled. Then,

$$q = \frac{Y(\mathbf{N})}{M}. \tag{6.11}$$

The probability that exactly i control packets are to be scheduled in a frame is

$$P_i = \binom{M}{i} q^i (1-q)^{M-i} \quad , i = 0, 1, 2, \ldots, M. \tag{6.12}$$

Each of the i control packets originates from one of the N nodes with equal probability $1/N$. With i control packets, in each frame the average number of control packets that belong to nodes attached to the same combiner is equal to $i \cdot \beta = i/D$. Recall that due to their fixed size of F slots, data packets can be sent from those nodes only every D frames. Data packets cannot be transmitted in other frames since they are larger than $(F - M)$ slots. Thus, in each frame only nodes attached to the same combiner can send data packets. Control packets emanating from nodes attached to the same combiner aggregate over the interval of D frames until data transmission takes place. As a consequence, in each frame the average number of control packets to be scheduled is given by $i/D \cdot D = i$.

The probability that at least one among those i control packets is destined to a given node under the assumption that a node does not transmit to itself is equal to [LK92a]

$$\begin{aligned} p_o(i) &= 1 - \left[\frac{i}{N} \left(1 - \frac{1}{N-1} \right)^{i-1} + \left(1 - \frac{i}{N} \right) \left(1 - \frac{1}{N-1} \right)^{i} \right] && (6.13) \\ &= 1 - \left(1 - \frac{1}{N-1} \right)^{i-1} \frac{N^2 - 2N + i}{N(N-1)}. && (6.14) \end{aligned}$$

Let $g(i)$ denote the average number of nodes that successfully transmit a data packet in a frame, given that i control packets are to be scheduled. Given this, the number of data packets destined to nodes that are attached to the same splitter is binomially distributed $BIN(S, p_o(i))$. However, no more than R data packets can be simultaneously transmitted to those nodes. This holds for each of the D splitters and we finally obtain

$$g(i) = D \left\{ \sum_{k=0}^{R} k \binom{S}{k} p_o(i)^k \left[1 - p_o(i) \right]^{S-k} + R \sum_{k=R+1}^{S} \binom{S}{k} p_o(i)^k \left[1 - p_o(i) \right]^{S-k} \right\}. \tag{6.15}$$

Note that at most S nodes can transmit data packets in a frame. Hence, $g(i)$ is bounded and the number of actually transmitting nodes is equal to $\min\{g(i),\ S\}$.

The conditional expectation of the number of nodes that successfully transmit a data packet in a frame, given that the system is in state $\mathbf{N}$, is given by

$$\begin{aligned} Z(\mathbf{N}) &= \sum_{i=0}^{M} g(i) \cdot P_i && (6.16) \\ &= \sum_{i=0}^{M} D \left\{ \sum_{k=0}^{R} k \binom{S}{k} p_o(i)^k \left[1 - p_o(i) \right]^{S-k} + R \sum_{k=R+1}^{S} \binom{S}{k} p_o(i)^k \left[1 - p_o(i) \right]^{S-k} \right\} \cdot \\ &\quad \cdot \binom{M}{i} q^i (1-q)^{M-i}. && (6.17) \end{aligned}$$

Using Eqns. (6.5)–(6.8) we can modify Eqns. (6.3) and (6.10). Eqn. (6.3) becomes

$$Z(\mathbf{N}) = \frac{\sigma}{1+\sigma}\left[N - \left(1 + \frac{\tau p}{D}\right) N_B\right] \tag{6.18}$$

and Eqn. (6.10) becomes

$$Y(\mathbf{N}) = N_B \frac{p}{D}\left(1 - \frac{p}{D \cdot M}\right)^{N_B - 1}. \tag{6.19}$$

Eqns. (6.17), (6.18), and (6.19) can be solved simultaneously for the variables N_B, $Y(\mathbf{N})$, and $Z(\mathbf{N})$. N_B, $Y(\mathbf{N})$, and $Z(\mathbf{N})$ can then be used to provide the steady–state solution of the entire system. N_{PR_i}, $i = 1, 2, \ldots, \tau$, is given by equations (6.1) and (6.5). Similarly, N_{PQ_i}, $i = 1, 2, \ldots, \tau$, is given by equations (6.2) and (6.6). According to equation (6.7), N_{TR} is equal to $Z(\mathbf{N})$. And N_{TH} equals N minus the sum of the nodes in all other modes.

Performance measures

The performance measures of interest are throughput and delay at an equilibrium point. The throughput $S(\mathbf{N})$ is defined as the expected number of nodes in the active mode TR:

$$S(\mathbf{N}) = N_{TR}. \tag{6.20}$$

The mean packet delay $D(\mathbf{N})$ is measured from the time the packet is generated at a node until the end of the frame during which it is transmitted. The system shown in Fig. 6.1 is a closed system, and, by Little's law, $N/S(\mathbf{N})$ is the average time that a packet experiences from the moment the packet enters mode TH until the time it returns to mode TH. Also, $1/\sigma$ is the average time that a packet stays in mode TH. Thus, we get the average packet delay as

$$D(\mathbf{N}) = \frac{N}{S(\mathbf{N})} - \frac{1}{\sigma}. \tag{6.21}$$

Note that $D(\mathbf{N})$ is measured in number of frames.

6.1.4 Results

In this section, we investigate the impact of the system parameters on the throughput–delay performance of the network by varying them around the default values. We consider number of nodes N, retransmission probability p, physical degree of the AWG D, propagation delay τ, number of used FSRs R, and number of reservation slots per frame M. Unless stated otherwise, the parameters are set to the following default values: $N = 240$, $p = 0.5$, $D = 2$, $\tau = 10$ frames, $R = 3$, and $M = 8$. The frame length F is assumed to be constant. Since we use the frame length as basic time unit in our performance evaluation we do not have to specify the number of slots per frame. To verify the accuracy of our analysis we simulated a more realistic system. As opposed to the analysis, in the simulation the first–time transmission of a control packet is not delayed and the destination of a collided control packet is not renewed each time it is retransmitted. Each simulation was run for 10^6 cycles including a warm–up phase of 10^5 cycles. We used the method of batch means to obtain confidence intervals for the mean throughput and the mean delay. For all simulation results the 98% confidence interval was less than 1% of the sample mean.

Fig. 6.2 depicts the mean throughput (mean number of transmitting nodes) vs. the mean arrival rate σ (packet/frame) for different populations N. We observe that the maximum mean

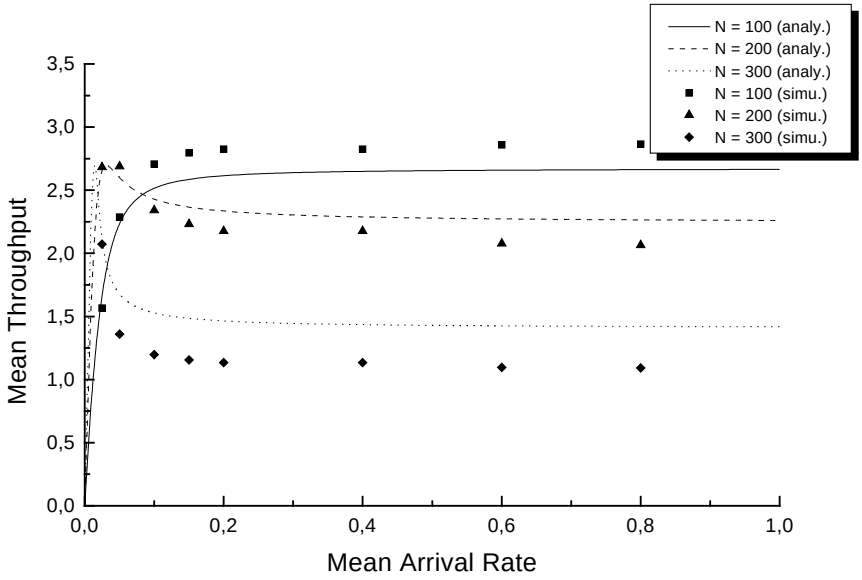

Figure 6.2: Mean throughput (mean number of transmitting nodes) vs. mean arrival rate (packet/frame) for different numbers of nodes $N \in \{100, 200, 300\}$.

throughput is the same for all populations. However, with more nodes this maximum is reached at smaller σ. This is because the mean throughput depends not only on σ but also on the offered traffic load which is equal to $N \cdot \sigma$. Consequently, with a larger N the maximum mean throughput occurs at a smaller σ. For $N \in \{200,\ 300\}$ the average throughput decreases with increasing σ. This is due to the fact that with increasing σ there are more nodes accessing the reservation slots. This causes more collisions and fewer scheduled data packets. The congestion becomes more serious for larger N, resulting in a lower average throughput. Note that with $N = 100$ the mean throughput is not reduced for increasing σ. In this case, N is small enough such that the offered load can be handled well by slotted ALOHA. This indicates that slotted ALOHA does not degrade the mean throughput as long as the population is small enough. Analysis and simulation results match very well for small values of σ, for larger σ there is some discrepancy. Note that generally for $N = 100$ the simulation gives higher throughputs than the analysis; whereas for $N \in \{200,\ 300\}$ we observe the opposite. This is due to the assumption of delayed first–time transmissions of control packets in our analysis. For large populations the number of collisions on the control channel is reduced by delaying also first–time transmissions, resulting in a higher throughput. But for small populations this delay leads to an underutilized control channel and thereby smaller throughput. Hence, this assumption yields accurate results only for moderate loads $N \cdot \sigma$.

The mean packet delay (in frames) vs. the mean arrival rate is shown in Fig. 6.3. For all populations N the average delay is bounded because we consider single–packet buffers at each node without accounting for queueing delays. New packets arriving at a backlogged node are discarded and do not contribute to the mean delay. With more nodes the slotted ALOHA channel gets congested already at low loads. This causes high delays due to retransmissions of control packets. We can see that with $N = 100$ the mean delay does not change much with

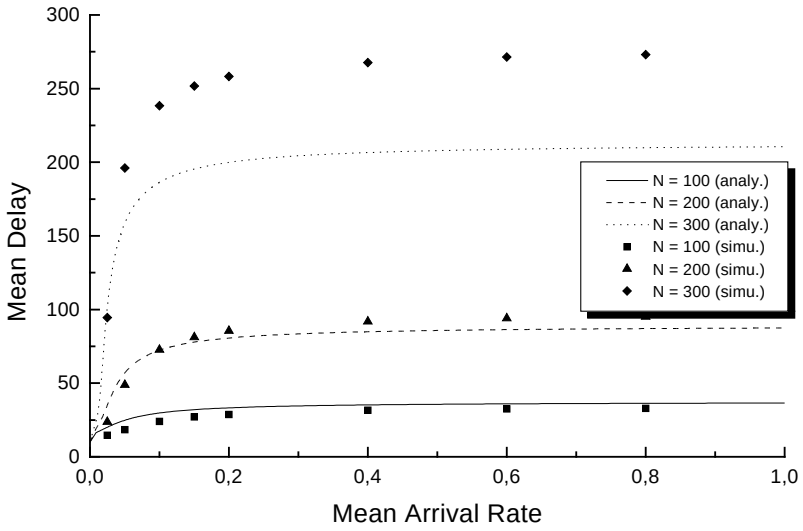

Figure 6.3: Mean delay (frames) vs. mean arrival rate (packet/frame) for different numbers of nodes $N \in \{100, 200, 300\}$.

increasing σ. This is because the offered load does not significantly overload slotted ALOHA leading to only a few collisions and retransmissions which in turn keeps the mean delay small. The figure shows, that for large populations ($N = 300$) the simulation gives larger delays than the analysis. This is again due to the fact that in the simulation first–time transmissions of control packets are not delayed leading to more collisions and more retransmissions of control packets.

Figs. 6.2 and 6.3 can be combined as illustrated in Fig. 6.4. This figure depicts the mean delay vs. mean throughput as σ is varied from 0 to 1. In the following we consider only this type of graph. The subsequent curves are obtained for $N = 240$.

Mean delay vs. mean throughput for different retransmission probabilities p is shown in Fig. 6.5. Apparently, at light traffic larger retransmission probabilities give smaller delays since control packets are more likely to be sent in a given frame. Whereas at higher loads smaller values of p result in a better throughput–delay performance. This is because a smaller p reduces the number of collisions on the busy slotted ALOHA channel and the number of unsuccessfully scheduled data packets. As a consequence, control packets have to be retransmitted fewer times, resulting in lower delay and higher throughput. Note that in the analysis the mean delay is larger for $p = 0.2$ than for $p \in \{0.4,\ 0.6\}$ since the backlogged nodes act too passively. Due to this passive behavior the slotted ALOHA channel is underutilized at light traffic. In the simulation, however, new control packets are not delayed, resulting in a lower delay at light traffic. With increasing σ slotted ALOHA gets congested, deteriorating the throughput–delay performance.

A given number of nodes can be connected by AWGs with different physical degree D. As depicted in Fig. 6.6, a 4×4 AWG yields higher maximum mean throughput and lower mean delay at high traffic loads than a 2×2 AWG. Using an 8×8 AWG, instead, increases the maximum average throughput only slightly but suffers from a larger average delay for lower

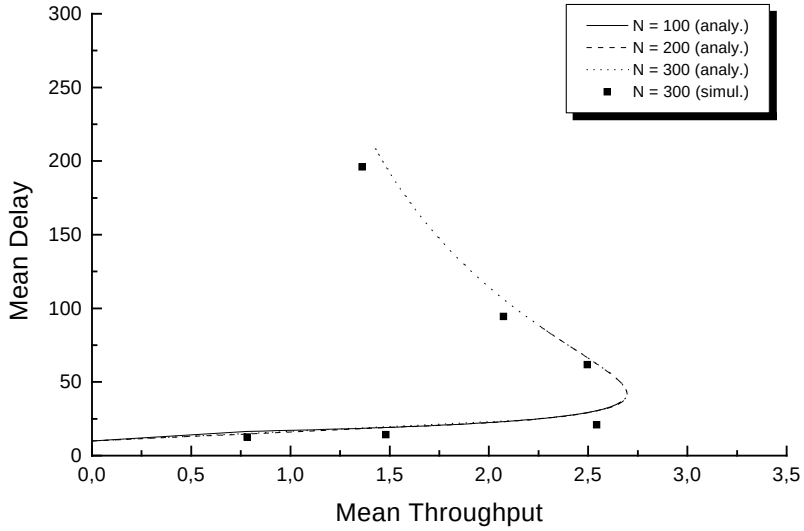

Figure 6.4: Mean delay (frames) vs. mean throughput (mean number of transmitting nodes) for different numbers of nodes $N \in \{100, 200, 300\}$.

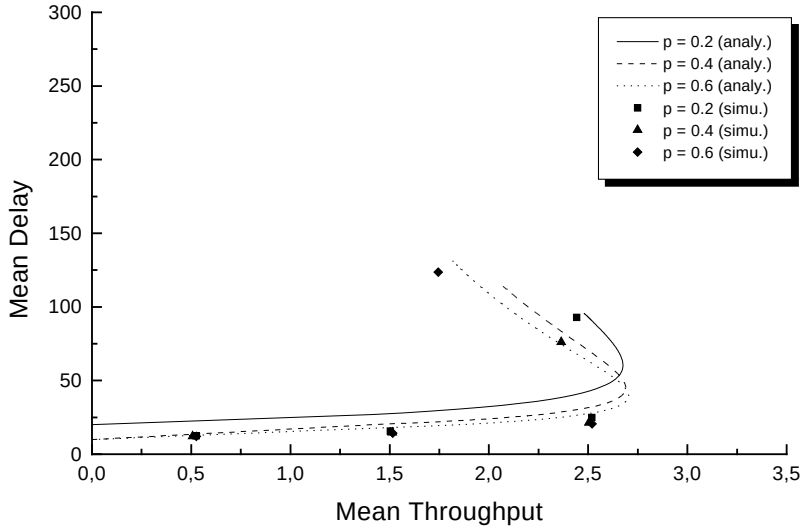

Figure 6.5: Mean delay (frames) vs. mean throughput (mean number of transmitting nodes) for different retransmission probabilities $p \in \{0.2, 0.4, 0.6\}$.

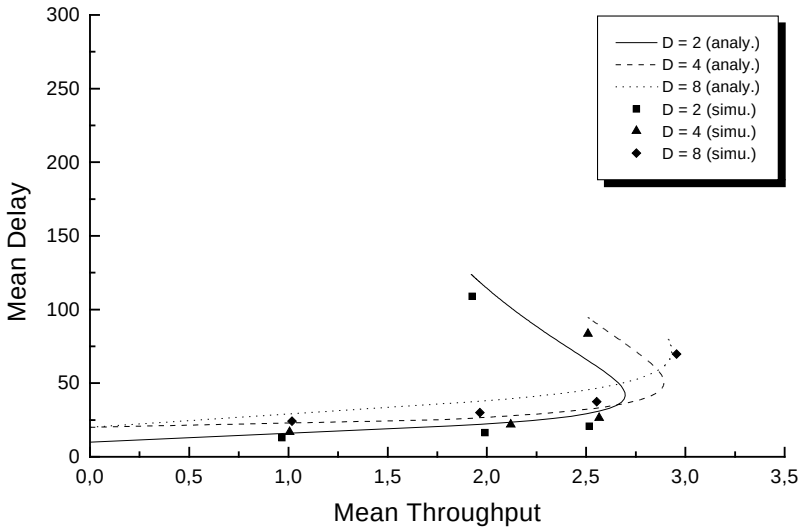

Figure 6.6: Mean delay (frames) vs. mean throughput (mean number of transmitting nodes) for different AWG degrees $D \in \{2, 4, 8\}$.

arrival rates. A large D implies that the degree S of each combiner is small. Accordingly, in a given frame fewer nodes access the slotted ALOHA channel reducing the contention and thereby improving the performance. However, with a D chosen too large the cycle becomes too long and nodes have to wait a longer time period, resulting in an increased delay. In addition, owing to the longer cycle length more nodes are backlogged and try to access the same frame increasing the number of collisions and limiting the throughput improvement. Thus, unless running the system under heavy traffic, it is reasonable to use an AWG with a rather moderate number of ports. Note that this allows to exploit more FSRs of an AWG for a given transceiver tuning range and channel spacing.

Fig. 6.7 depicts the impact of the propagation delay on the network performance. At low traffic loads packets experience less delay for smaller propagation delays. Whereas with increasing traffic, larger propagation delays provide a better throughput–delay performance. This is due to the fact that at low traffic loads almost no collisions of control packets occur and nodes receive the successfully transmitted control packets earlier with smaller propagation delays, resulting in a decreased delay. At higher traffic loads the control channel gets more congested. In this case, a larger propagation delay implies that nodes have to wait a longer time interval for the transmitted control packets. During this time period those nodes do not access the control channel, resulting in less contention and an increased throughput and a decreased delay due to fewer retransmissions. Note that the propagation delay in metro networks with limited diameter is rather small.

The results in Fig. 6.8 clearly demonstrate the benefit of using multiple FSRs of an AWG. Each additional FSR increases the degree of concurrency and thereby alleviates the scheduling bottleneck, resulting in a significantly improved throughput–delay performance of the network. Using three FSRs instead of one improves the maximum mean throughput by approximately

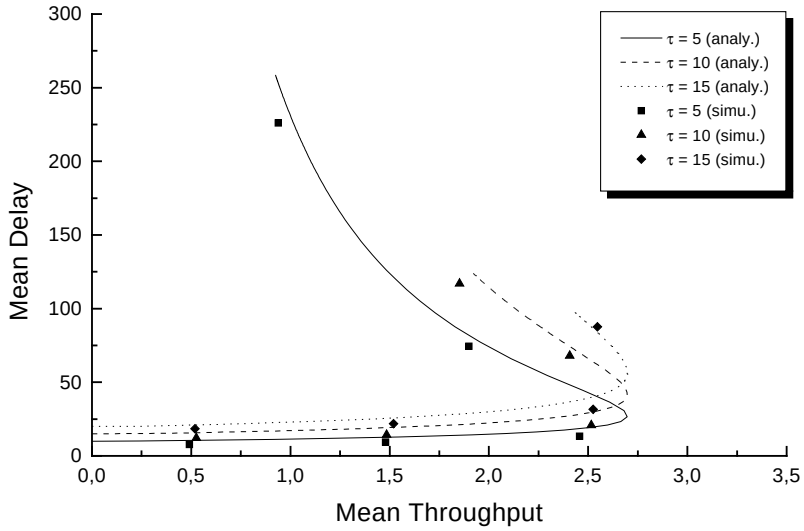

Figure 6.7: Mean delay (frames) vs. mean throughput (mean number of transmitting nodes) for different propagation delays $\tau \in \{5, 10, 15\}$.

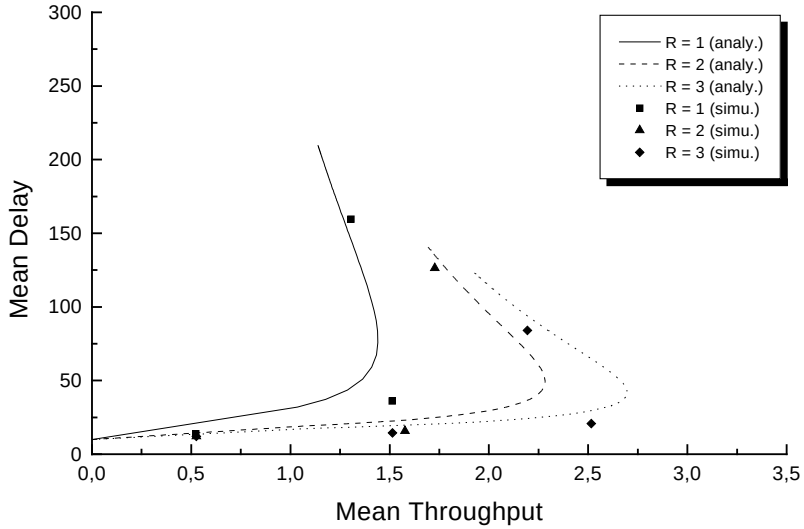

Figure 6.8: Mean delay (frames) vs. mean throughput (mean number of transmitting nodes) for different numbers of used FSRs $R \in \{1, 2, 3\}$.

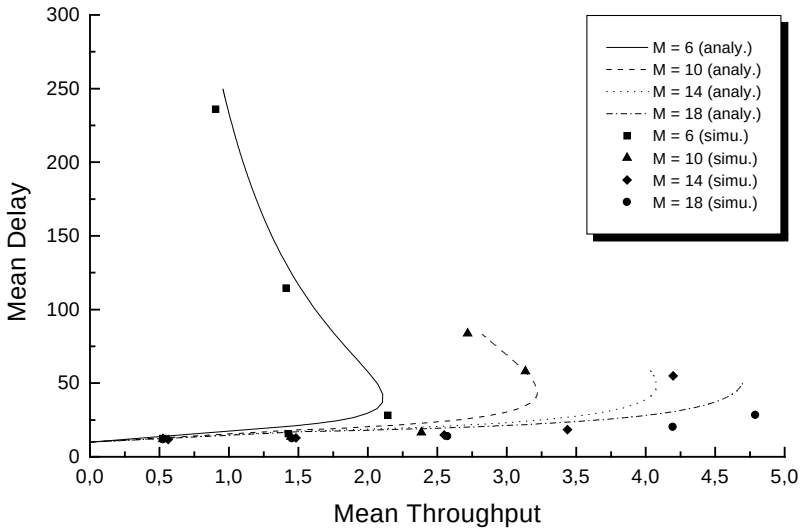

Figure 6.9: Mean delay (frames) vs. mean throughput (mean number of transmitting nodes) for different numbers of reservation slots $M \in \{6, 10, 14, 18\}$.

88%. However, using multiple FSRs requires transceivers with a larger tuning range. With $R = 3$ and $D = 2$ the wavelength pool comprises six wavelengths. Assuming a channel spacing of 200 GHz (1.6 nm at 1.55 μm) each transceiver must be tunable over 8 nm, which is achievable with fast tunable transceivers. Reducing the channel spacing provides additional FSRs at the expense of higher crosstalk. The additional FSRs can be accessed by the nodes in order to further improve the throughput–delay performance of the network.

Fig. 6.9 shows that a large number of reservation slots in each frame has a positive impact on the system performance. By using more reservation slots the collision probability in each slot is reduced and the number of successful control packets is increased. Consequently, fewer control packets have to be retransmitted improving the throughput–delay performance. Note that for $M = 18$ a normalized mean throughput of up to 78% is achieved. However, for a given frame length F increasing M decreases the remaining number of slots $(F - M)$ per frame. Recall from Section 5.3.1 that these $(F - M)$ slots allow for spatial wavelength reuse. Consequently, the length of packets which benefit from spatial wavelength reuse is decreased. Note that in this section we have considered only data packets whose length is equal to one frame. These fixed–size packets are too long to be transmitted in the last $(F - M)$ slots of any frame, and can thus not exploit spatial wavelength reuse. In the next section we allow data packets to have variable size such that spatial wavelength reuse becomes possible.

6.2 Spatial wavelength reuse

In this section, we show the benefit of spatial wavelength reuse. Spatial wavelength reuse is a crucial technique for achieving flexibility, efficiency, and cost–effectiveness in the considered metro WDM network. Flexibility is required since metro networks have to support a wide

range of heterogeneous protocols. This requires, in particular, that metro networks are able to transport variable–size packets. Recall from Section 5.3.1 that data packets which are small enough can also be transmitted in the last $(F - M)$ slots of *each* frame compared to long packets (packet size equals F slots) which can only be sent in one frame per cycle. In these slots *all wavelengths* can be used *at each AWG port simultaneously*, resulting in an increased network efficiency. By increasing the efficiency the utilization of the deployed WDM networking components and the WDM networking resources (in particular wavelengths) is increased, making the network cost effective. In the following we develop a stochastic model. As in Section 6.1.1, we concentrate on uniform unicast packet switching. As opposed to the previous section, however, packets have *variable* size and wavelengths can be *spatially reused*. Furthermore, we drop the assumption of delayed first–time transmission of control packets in order to improve the analytical accuracy. The assumptions made in Section 6.2.1 are verified by extensive simulations in Section 6.2.3 [MSRW02].

6.2.1 Assumptions

In our analysis we make the following assumptions:

- Each node has a *single–control–packet* and a *single–data–packet* buffer, i.e., each node can store at most one control packet and one data packet at any given time.
- After successfully transmitting a control packet and successfully scheduling the corresponding data packet in a given frame the single–control–packet buffer becomes empty at the end of that frame. The node's data buffer may hold the scheduled (but not yet transmitted) data packet while the respective control packet has already been deleted.
- A given data packet is purged from the node's buffer at the end of the frame during which it is transmitted.
- A node with an empty control buffer generates a control packet with probability σ at the beginning of that frame in which the node is allowed to make reservations. If no control packet is generated the node waits for one cycle and then generates a control packet with probability σ, and so on.
- After a data packet is purged from the buffer, the next data packet is placed in the buffer, provided the corresponding control packet is already in the buffer.
- A data packet has *variable size* L: A data packet is long (has size $L = F$ with probability q, and is short (has size $L = K$, where $1 \leq K \leq (F - M)$) with probability $(1 - q)$. That is, $P(L = F) = q$ and $P(L = K) = 1 - q$.
- *Uniform unicast* traffic: A data packet is destined to any one of the N nodes with equal probability $1/N$. (For simplicity, a given node is allowed to transmit data packets to itself.)
- The propagation delay τ is the same for all nodes, i.e., all nodes are *equidistant* from the AWG. τ is assumed to be no larger than one cycle. (This assumption is reasonable for a metropolitan area network.)
- *Length persistency*: The size of a given data packet is not changed if the corresponding control packet is retransmitted.

- *Destination nonpersistency*: Random selection of a destination node among the N nodes is renewed for each attempt of transmitting a control packet.
- A generated control packet is sent with probability one in the frame that is assigned to the node's AWG input port, possibly simultaneously with a previously scheduled data packet. (As opposed to Section 6.1.1 we do not assume delayed first–time transmissions of control packets in order to improve the accuracy of the analysis, as we will see shortly.)
- An unsuccessful control packet is retransmitted with probability p in the next frame assigned to the node's AWG input port.

6.2.2 Analysis

In our analysis we consider a system with a large S. Our analysis is approximate for finite S and exact in the asymptotic limit $S \to \infty$. We define

$$\tilde{\alpha} := \frac{S\sigma}{M} \text{ and } \alpha := \frac{Sp}{M}. \tag{6.22}$$

Our analysis becomes asymptotically exact when $S \to \infty$ and $\tilde{\alpha}$ as well as α (and also M) are fixed (with σ and p chosen so as to satisfy (6.22)).

Now consider the nodes attached to a given (fixed) AWG input port o, $1 \leq o \leq D$. These nodes send their control packets in frame o of a given cycle. We refer to the nodes that at the beginning of frame o hold an old packet, that is, a control packet that has failed in slotted ALOHA or scheduling, as *"old"*. We refer to all the other nodes as *"new"*. Note that the set of "new" nodes comprises both the nodes that have generated a new (never before transmitted) control packet as well as the nodes that have deferred the generation of a new control packet. Let η be a random variable denoting the number of "new" nodes at AWG input port o, and let

$$\nu := \frac{E[\eta]}{S}. \tag{6.23}$$

Let λ_l be a random variable denoting the number of nodes at port o that are to send a control packet corresponding to a long data packet next (irrespective of whether a given node is "old" or "new", and keeping in mind that the set of "new" nodes also comprises those nodes that have deferred the generation of the next control packet; those nodes are accounted for in λ_l if the next generated control packet corresponds to a long data packet). Let

$$\tilde{q} := \frac{E[\lambda_l]}{S} \tag{6.24}$$

denote the expected fraction of long packets to be sent. We expect that $\tilde{q}$ is typically larger than q since long packets are harder to schedule and thus typically require more retransmissions (of control packets).

Slotted ALOHA contention

First, we calculate the number of control packets from nodes attached to AWG input port o, $1 \leq o \leq D$, that are successful in the slotted ALOHA contention in frame o. Let Y_i^n, $i = 1, \ldots, M$, be a random variable denoting the number of control packets that were randomly transmitted

in slot $i,\ i = 1, \ldots, M$, by "new" nodes. Recall that each of the η "new" nodes sends a control packet with probability σ in the frame. Thus,

$$P(Y_i^n = k) = \binom{\eta}{k} \left(\frac{\sigma}{M}\right)^k \left(1 - \frac{\sigma}{M}\right)^{\eta - k}, \quad k = 0, 1, \ldots, \eta. \tag{6.25}$$

Throughout our analysis we assume that S is large and that $\tilde{\alpha}$ and α are fixed. We may therefore reasonably approximate the BIN($\eta,\ \sigma/M$) distribution with a Poisson ($\eta\sigma/M$) distribution, that is,

$$P(Y_i^n = k) \approx e^{-\eta\sigma/M} \frac{(\eta\sigma/M)^k}{k!}, \quad k = 0, 1, \ldots, \tag{6.26}$$

which is exact for $\eta \to \infty$ with $\eta\sigma/M$ fixed. (A refined analysis that does not approximate the binomial distribution by the Poisson distribution is given in Appendix C.) We now recall the definition $\tilde{\alpha} := S\sigma/M$. We also approximate η/S by its expectation ν; this is reasonable since η/S has only small fluctuations in steady state for large S. Thus,

$$P(Y_i^n = k) \approx e^{-\tilde{\alpha}\nu} \frac{(\tilde{\alpha}\nu)^k}{k!}, \quad k = 0, 1, \ldots. \tag{6.27}$$

We note that for $S \to \infty$ the random variables $Y_1^n, Y_2^n, \ldots, Y_M^n$ are mutually independent. This is because a given node places with the miniscule probability σ/M a control packet in a given slot, say slot 1. (Note in particular that the expected value of Y_1^n is small compared to the number of "new" nodes, that is, $\tilde{\alpha}\nu \ll \eta$; this is because in the considered asymptotic limit $S \to \infty$ with α fixed, we have $1 \gg \sigma/M = \tilde{\alpha}\nu/\eta$.) Thus, Y_1^n has almost no impact on $Y_2^n, \ldots, Y_M^n$.

Let $Y_i^o,\ i = 1, \ldots, M$, be a random variable denoting the number of control packets in slot $i,\ i = 1, \ldots, M$, that originate from "old" nodes. Each of the $(S - \eta)$ "old" nodes sends a control packet with probability p in the frame. Thus,

$$P(Y_i^o = k) = \binom{S - \eta}{k} \left(\frac{p}{M}\right)^k \left(1 - \frac{p}{M}\right)^{S - \eta - k},\ k = 0, 1, \ldots, S - \eta. \tag{6.28}$$

Approximating this BIN($S - \eta,\ p/M$) distribution by the Poisson ($(S - \eta)p/M$) distribution we have

$$P(Y_i^o = k) \approx e^{-(S-\eta)p/M} \frac{[(S - \eta)p/M]^k}{k!},\ k = 0, 1, \ldots. \tag{6.29}$$

With $\alpha := Sp/M$ and approximating $(S - \eta)/S$ by its expectation $(1 - \nu)$ we get

$$P(Y_i^o = k) \approx e^{-\alpha(1-\nu)} \frac{[\alpha(1 - \nu)]^k}{k!},\ k = 0, 1, \ldots. \tag{6.30}$$

We note again that for $S \to \infty$ the random variables $Y_1^o, Y_2^o, \ldots, Y_M^o$ are mutually independent. They are also independent of $Y_1^n, Y_2^n, \ldots, Y_M^n$. Hence, we obtain for $Y_i = Y_i^n + Y_i^o$

$$P(Y_i = k) \approx e^{-[\tilde{\alpha}\nu + \alpha(1-\nu)]} \frac{[\tilde{\alpha}\nu + \alpha(1 - \nu)]^k}{k!},\ k = 0, 1, \ldots. \tag{6.31}$$

Henceforth, we let for notational convenience

$$\beta := \tilde{\alpha}\nu + \alpha(1 - \nu), \tag{6.32}$$

i.e.,

$$P(Y_i = k) \approx e^{-\beta}\frac{\beta^k}{k!},\ k = 0, 1, \ldots. \tag{6.33}$$

Let $X_i,\ i = 1, \ldots, M$, be a random variable indicating whether or not slot i contains a successful control packet. Specifically, let

$$X_i = \begin{cases} 1 & \text{if } Y_i = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{6.34}$$

Clearly, from (6.33) $P(X_i = 1) = \beta e^{-\beta}$ and $P(X_i = 0) = 1 - \beta e^{-\beta}$ for $i = 1, \ldots, M$. The total number of successful control packets in the considered frame is $\sum_{i=1}^{M} X_i$, which has a BIN($M,\ \beta e^{-\beta}$) distribution, that is,

$$P\left(\sum_{i=1}^{M} X_i = l\right) = \binom{M}{l}\left(\beta e^{-\beta}\right)^l\left(1 - \beta e^{-\beta}\right)^{M-l},\ l = 0, 1, \ldots, M. \tag{6.35}$$

Recall from Section 6.2.1 that each packet is destined to any one of the D AWG output ports with equal probability $1/D$. Let Z denote the number of successful control packets — in the considered frame — that are destined to a given (fixed) AWG output port $d,\ d = 1, \ldots, D$. Clearly, from (6.35)

$$P(Z = k) = \binom{M}{k}\left(\frac{\beta e^{-\beta}}{D}\right)^k\left(1 - \frac{\beta e^{-\beta}}{D}\right)^{M-k},\ k = 0, 1, \ldots, M. \tag{6.36}$$

(A refined approximation of $P(Z = k)$ which does not approximate the binomial distributions of the Y_i^n's and Y_i^o's by Poisson distributions is given in Appendix C.) Let Z_l denote the number of successful control packets that correspond to long data packets destined to a given AWG output port d. Recall that $\tilde{q}$ is the expected fraction of long packets to be sent. Hence, $Z_l \sim BIN(M,\ \beta e^{-\beta}\frac{1}{D}\tilde{q})$. Similarly, let Z_k denote the number of control packets that are successful in the slotted ALOHA contention and correspond to short data packets destined to a given AWG output port d. Clearly, $Z_k \sim BIN(M,\ \beta e^{-\beta}\frac{1}{D}(1 - \tilde{q}))$.

Packet scheduling

In this section we calculate the expected number of packets that are successfully scheduled. Recall from the previous section that the total number of long packets that (1) originate from a given AWG input port $o,\ 1 \le o \le D$, (2) are successful in the slotted ALOHA contention of frame o (of a given cycle), and (3) are destined to a given AWG output port $d,\ 1 \le d \le D$, is $Z_l \sim BIN(M,\ \beta e^{-\beta}\frac{1}{D}\tilde{q})$. For short packets we have $Z_k \sim BIN(M,\ \beta e^{-\beta}\frac{1}{D}(1 - \tilde{q}))$. Note that these two random variables are not independent. Let $\mathcal{L}$ ($\mathcal{S}$) be a random variable denoting the number of long (short) packets that (1) originate from a given AWG input port $o,\ 1 \le o \le D$, (2) are successful in the slotted ALOHA contention of frame o (of a given cycle), (3) are destined to a given AWG output port $d,\ 1 \le d \le D$, and (4) are successfully scheduled within the scheduling window of D frames (i.e., one cycle).

Consider the scheduling of packets from a given (fixed) AWG input port o to a given (fixed) AWG output port d over the scheduling window (i.e., D frames). Clearly, we can schedule at most R long packets (i.e., $\mathcal{L} \le R$) because the receivers at output port d must tune to the

appropriate spectral slices during the first M slots of every frame. Thus, they can tune to a node at AWG input port o for F consecutive slots, only in the frame, during which the nodes at AWG input port o send their control packets.

Now, suppose that $\mathcal{L}$ ($\leq R$) long packets are scheduled (how $\mathcal{L}$ is determined is discussed shortly). With $\mathcal{L}$ long packets already scheduled, we can schedule at most

$$\mathcal{S} \leq (D-1) \cdot R \cdot \left\lfloor \frac{F-M}{K} \right\rfloor + (R-\mathcal{L}) \left\lfloor \frac{F}{K} \right\rfloor \tag{6.37}$$

short packets. To see this, note that in the frame during which the nodes at AWG input port o send their control packets, there are $(R-\mathcal{L})$ FSRs — channels between AWG input port o and AWG output port d — free for a duration of F consecutive slots. Furthermore, there are $(D-1)$ frames in the scheduling window during which the nodes at AWG output port d must tune (are locked) to the nodes sending control packets from the other AWG input ports for the first M slots of the frame. During each of these frames, the receivers are unlocked for $(F-M)$ slots. The R utilized FSRs provide R parallel channels between AWG input port o and AWG output port d. Note that the $(D-1)R\lfloor (F-M)/K \rfloor$ component in (6.37) is due to the spatial reuse of wavelengths at the considered AWG input port. Without spatial wavelength reuse this component would be zero and we could schedule at most $(R-\mathcal{L})\lfloor F/K \rfloor$ short packets. Continuing our analysis for a network with spatial wavelength reuse, we have

$$\mathcal{S} = \min \left\{ Z_k,\ (D-1) \cdot R \cdot \left\lfloor \frac{F-M}{K} \right\rfloor + (R-\mathcal{L}) \left\lfloor \frac{F}{K} \right\rfloor \right\}. \tag{6.38}$$

In (6.38) we neglect receiver collisions, that is, we do not account for situations where a packet can not be scheduled because its receiver is already scheduled to receive a different packet. This assumption is reasonable as receiver collisions are rather unlikely for large S, which we assume throughout our analysis.

Recall from Section 5.3.1 that we apply a first–come–first–served–first–fit scheduling policy. Data packets are scheduled for the the first possible slot(s) on the lowest available wavelength. To arbitrate the access to the frame which allows transmission for F contiguous slots and the $(D-1)$ frames which allow transmission for $(F-M)$ contiguous slots we adopt the following *arbitration policy*. Our arbitration policy proceeds in one round if there are R or less successful control packets in the slotted ALOHA contention. In case there are more than R successful control packets in the slotted ALOHA contention, our arbitration policy proceeds in two rounds. First, consider the case where R or less control packets are successful in the slotted ALOHA contention and we have one round of arbitration. In this case all the successful packets are scheduled in the frame with F available transmission slots. Next, consider the case where more than R control packets are successful in the slotted ALOHA contention and we have two rounds of arbitration. In this case we scan the M slotted ALOHA slots from index 1 through M. In the first round we schedule the first R successful packets out of the slotted ALOHA contention in the R long (F slots) transmission slots. In this round we schedule only one packet for each of the long transmission slots, irrespective of whether the packet is long or short. At this point (having filled each of the long transmission slots with one data packet) all the remaining successful control packets that correspond to long data packets fail in the scheduling and the transmitting node has to retransmit the control packet. We then proceed with the second round. In the second round we schedule the remaining successful control packets that correspond to short data packets. Provided $F/K \geq 2$, we schedule these short data packets for the long transmission slots that hold only one short data packet from the first round. We also schedule these short data packets

for the short ($(F-M)$ slots) transmission slots. After all the long and short transmission slots have been filled, the remaining short data packets fail in the scheduling and the transmitting node has to retransmit the control packet. (We note that our adopted arbitration policy is just one out of many possible arbitration policies, all of which can be analyzed in a similar fashion.)

With the adopted arbitration policy the expected number of scheduled long packets is

$$\begin{aligned}
E[\mathcal{L}] &= \sum_{k=0}^{M} E[\mathcal{L}|Z=k]\cdot P(Z=k) && (6.39)\\
&= \sum_{k=0}^{M} \min(k,\ R)\cdot \tilde{q}\cdot P(Z=k) && (6.40)\\
&= \tilde{q}\left\{\sum_{k=0}^{\min(R,\ M)} k\cdot P(Z=k) + R\sum_{k=\min(R,M)+1}^{M} P(Z=k)\right\} && (6.41)\\
&= \tilde{q}\left\{\sum_{k=0}^{\min(R,\ M)} k\cdot P(Z=k) + R\left[1-\sum_{k=0}^{\min(R,\ M)} P(Z=k)\right]\right\} && (6.42)\\
&= \tilde{q}\left\{R+\sum_{k=0}^{\min(R,\ M)} (k-R)\cdot P(Z=k)\right\} && (6.43)\\
&= \tilde{q}\left\{R-\sum_{k=0}^{\min(R,\ M)} (R-k)\cdot \binom{M}{k}\left(\frac{\beta e^{-\beta}}{D}\right)^{k}\left(1-\frac{\beta e^{-\beta}}{D}\right)^{M-k}\right\}. && (6.44)
\end{aligned}$$

To see this, note that in case there are $k \le R$ successful control packets in the slotted ALOHA contention, then on average $\tilde{q}k$ of these correspond to long data packets. In case $k \ge R$, then there are on average $\tilde{q}R$ long packets among the first R control packets. (If the arbitration policy does not schedule the long (F slots) transmission slots first, but, say after l frames that allow transmission for $(F-M)$ slots have been scheduled, then an expected number of $\min\{R,\ \tilde{q}(lR\lfloor\frac{F-M}{K}\rfloor + R)\}$ long data packets are scheduled given $(l\lfloor\frac{F-M}{K}\rfloor+1)R$ or more successful control packets in the slotted ALOHA contention.) For notational convenience let

$$\varphi(\beta) := R - \sum_{k=0}^{\min(R,\ M)} (R-k)\cdot \binom{M}{k}\left(\frac{\beta e^{-\beta}}{D}\right)^{k}\left(1-\frac{\beta e^{-\beta}}{D}\right)^{M-k}. \qquad (6.45)$$

Thus,

$$E[\mathcal{L}] = \tilde{q}\cdot\varphi(\beta). \qquad (6.46)$$

We now calculate the expected number of scheduled short packets. Generally,

$$E[\mathcal{S}] = \sum_{k=0}^{M} E[\mathcal{S}|Z=k]\cdot P(Z=k). \qquad (6.47)$$

First, consider the case that there are no more than R successful control packets in the slotted ALOHA contention, i.e., $k \le R$. In this case we have with (6.38)

$$\begin{aligned}
E[\mathcal{S}|Z=k] &= E[Z_k|Z=k] && (6.48)\\
&= (1-\tilde{q})\cdot k, && (6.49)
\end{aligned}$$

since all the successful control packets are scheduled in the long transmission slots.

Next, consider the case $R \leq k \leq M$. Let Θ denote the number of control packets that correspond to short data packets to be scheduled in the second round of arbitration. Note that $Z_k = R - \mathcal{L} + \Theta$, because $(R - \mathcal{L})$ short data packets have been scheduled in the first round of arbitration. With (6.38) we obtain

$$\begin{aligned} E[\mathcal{S}|Z=k] &= E\left[\min\left(R-\mathcal{L}+\Theta, (D-1)R\left\lfloor\frac{F-M}{K}\right\rfloor + \right.\right. \\ &\quad \left.\left.(R-\mathcal{L})\left\lfloor\frac{F}{K}\right\rfloor\right)|Z=k\right] \qquad (6.50) \\ &= E\left[R-\mathcal{L}|Z=k\right] + E\left[\min\left(\Theta, (D-1)R\left\lfloor\frac{F-M}{K}\right\rfloor + \right.\right. \\ &\quad \left.\left.(R-\mathcal{L})\left(\left\lfloor\frac{F}{K}\right\rfloor - 1\right)\right) \mid Z=k\right]. \qquad (6.51) \end{aligned}$$

Note that for a network without spatial wavelength reuse the $(D-1)R\lfloor(F-M)/K\rfloor$ term has to be replaced by zero in (6.50) and (6.51), as well as all the following expressions in this section.

Clearly, $E[R - \mathcal{L}|Z = k] = (1 - \tilde{q})R$, since $k \geq R$. Moreover, note that conditional on $Z = k,\ k \geq R$, the random variables $\mathcal{L}$ and Θ are independent. This is because the first R successful control packets in the slotted ALOHA slots determine $\mathcal{L}$; Θ is determined by the subsequent $(k - R)$ successful control packets. Hence,

$$\begin{aligned} &E\left[\min\left(\Theta,\ (D-1)R\left\lfloor\frac{F-M}{K}\right\rfloor + (R-\mathcal{L})\left(\left\lfloor\frac{F}{K}\right\rfloor - 1\right)\right) \mid Z=k\right] \qquad (6.52) \\ &= \sum_{j=1}^{k-R} P\left(\min\left(\Theta,\ (D-1)R\left\lfloor\frac{F-M}{K}\right\rfloor + (R-\mathcal{L})\left(\left\lfloor\frac{F}{K}\right\rfloor - 1\right)\right) \geq j \mid Z=k\right) \qquad (6.53) \\ &= \sum_{j=1}^{k-R} P(\Theta \geq j|Z=k) \cdot \\ &\quad P\left((D-1)\,R\left\lfloor\frac{F-M}{K}\right\rfloor + (R-\mathcal{L})\left(\left\lfloor\frac{F}{K}\right\rfloor - 1\right) \geq j|Z=k\right). \qquad (6.54) \end{aligned}$$

Now,

$$\begin{aligned} P(\Theta \geq j|Z=k) &= \sum_{m=j}^{k-R} P(\Theta = m|Z=k) \qquad (6.55) \\ &= \sum_{m=j}^{k-R} \binom{k-R}{m} (1-\tilde{q})^m\, \tilde{q}^{k-R-m}. \qquad (6.56) \end{aligned}$$

For notational convenience let

$$\gamma_j := P\left((D-1)\,R\left\lfloor\frac{F-M}{K}\right\rfloor + (R-\mathcal{L})\left(\left\lfloor\frac{F}{K}\right\rfloor - 1\right) \geq j|Z=k\right), \qquad (6.57)$$

and

$$v_j := \min\left(R,\ \frac{(D-1)R\left\lfloor\frac{F-M}{K}\right\rfloor - j}{\left\lfloor\frac{F}{K}\right\rfloor - 1} + R\right). \qquad (6.58)$$

If $\lfloor F/K \rfloor - 1 > 0$ then

$$\begin{aligned}
\gamma_j &= P(\mathcal{L} \leq v_j | Z = k) && (6.59)\\
&= \sum_{\{m:\ m \leq v_j\}} P(\mathcal{L} = m | Z = k) && (6.60)\\
&= \sum_{\{m:\ m \leq v_j\}} \binom{R}{m} \tilde{q}^m \left(1 - \tilde{q}\right)^{R-m}. && (6.61)
\end{aligned}$$

In case $\lfloor F/K \rfloor = 1$ we have

$$\gamma_j = \begin{cases} 1 & \text{if } j \leq (D-1)R \left\lfloor \frac{F-M}{K} \right\rfloor \\ 0 & \text{otherwise.} \end{cases} \tag{6.62}$$

Combining the cases $k \leq R$ and $R \leq k \leq M$, we obtain for $M \geq R$, which is typical for practical networks,

$$\begin{aligned}
E[\mathcal{S}] &= \sum_{k=0}^{R} (1-\tilde{q}) \cdot k \cdot P(Z=k) + (1-\tilde{q}) \cdot R \sum_{k=R+1}^{M} P(Z=k) + \\
&\quad \sum_{k=R+1}^{M} \left(\sum_{j=1}^{k-R} \gamma_j \sum_{m=j}^{k-R} \binom{k-R}{m} \left(1-\tilde{q}\right)^m \tilde{q}^{k-R-m} \right) \cdot P(Z=k) && (6.63)\\
&= (1-\tilde{q}) \sum_{k=0}^{R} k \cdot P(Z=k) + (1-\tilde{q}) \cdot R \left[1 - \sum_{k=0}^{R} P(Z=k) \right] + \\
&\quad \sum_{j=1}^{M-R} \gamma_j \sum_{m=j}^{M-R} \sum_{k=R+1}^{M} \binom{k-R}{m} (1-\tilde{q})^m \tilde{q}^{k-R-m} \cdot P(Z=k) && (6.64)\\
&= (1-\tilde{q}) \left[R - \sum_{k=0}^{R} (R-k) \cdot P(Z=k) \right] + \\
&\quad \sum_{j=1}^{M-R} \gamma_j \sum_{m=j}^{M-R} \sum_{k=m+R}^{M} \binom{k-R}{m} \left(1-\tilde{q}\right)^m \tilde{q}^{k-R-m} \cdot P(Z=k) && (6.65)\\
&= (1-\tilde{q}) \left[R - \sum_{k=0}^{R} (R-k) \cdot \binom{M}{k} \left(\frac{\beta e^{-\beta}}{D} \right)^k \left(1 - \frac{\beta e^{-\beta}}{D} \right)^{M-k} \right] + \\
&\quad \sum_{j=1}^{M-R} \gamma_j \sum_{m=j}^{M-R} \sum_{k=m+R}^{M} \binom{k-R}{m} \left(1-\tilde{q}\right)^m \tilde{q}^{k-R-m} \cdot \\
&\quad \binom{M}{k} \left(\frac{\beta e^{-\beta}}{D} \right)^k \left(1 - \frac{\beta e^{-\beta}}{D} \right)^{M-k} && (6.66)\\
&=: h(\tilde{q},\ \beta). && (6.67)
\end{aligned}$$

Equilibrium conditions

Next, we put the analyses for the individual components of the considered network, namely traffic model, slotted ALOHA contention, and packet scheduling, together. We establish two

equilibrium conditions and solve for the two unknowns $\tilde{q}$ and β. (Alternatively, we may consider the two unknowns $\tilde{q}$ and ν, noting that $\nu = (\beta - \alpha)/(\tilde{\alpha} - \alpha)$ for $\tilde{\alpha} \neq \alpha$; the case $\tilde{\alpha} = \alpha$ is discussed shortly.)

In steady state the system satisfies the equilibrium condition

$$E[\mathcal{L}] = q \cdot (E[\mathcal{L}] + E[\mathcal{S}]). \tag{6.68}$$

To see this, note that in equilibrium the mean number of scheduled long packets from a given (fixed) AWG input port destined to a given (fixed) AWG output port (left–hand side (LHS) of (6.68)) is equal to the mean number of newly generated long packets (right–hand side (RHS) of (6.68)). Inserting (6.46) and (6.67) in (6.68) gives

$$\tilde{q} \cdot \varphi(\beta) = q \cdot [\tilde{q} \cdot \varphi(\beta) + h(\tilde{q},\ \beta)] \tag{6.69}$$

$$\Leftrightarrow (1 - q) \cdot \tilde{q} \cdot \varphi(\beta) = q \cdot h(\tilde{q},\ \beta). \tag{6.70}$$

The second equilibrium condition is

$$\frac{\sigma}{D} \cdot E[\eta] = E[\mathcal{L} + \mathcal{S}]. \tag{6.71}$$

This is because $\sigma \cdot \eta$ new packets are generated in each frame at the nodes attached to a given AWG input port. With probability $1/D$ each of the generated packets is destined to a given (fixed) AWG output port. On the other hand, $E[\mathcal{L}+\mathcal{S}]$ packets are scheduled (and transmitted) on average from a given AWG input port to a given AWG output port in one cycle; in equilibrium as many new packets must be generated. Inserting (6.22) and (6.23) in the LHS of (6.71) and (6.46) and (6.67) in the RHS of (6.71) we obtain

$$\frac{\tilde{\alpha} \cdot M}{D} \cdot \frac{\beta - \alpha}{\tilde{\alpha} - \alpha} = \tilde{q} \cdot \varphi(\beta) + h(\tilde{q},\ \beta). \tag{6.72}$$

Inserting (6.72) in (6.70) we obtain

$$\tilde{q} = \frac{q \cdot \tilde{\alpha} \cdot M}{D \cdot \varphi(\beta)} \cdot \frac{\beta - \alpha}{\tilde{\alpha} - \alpha}. \tag{6.73}$$

Inserting (6.73) in (6.72) we obtain

$$(1 - q) \cdot \frac{\tilde{\alpha} \cdot M}{D} \cdot \frac{\beta - \alpha}{\tilde{\alpha} - \alpha} = h\left(\frac{q \cdot \tilde{\alpha} \cdot M}{D \cdot \varphi(\beta)} \cdot \frac{\beta - \alpha}{\tilde{\alpha} - \alpha},\ \beta\right). \tag{6.74}$$

We solve Eqn. (6.74) numerically to obtain β (noting that by (6.32), $\min(\tilde{\alpha},\ \alpha) \leq \beta \leq \max(\tilde{\alpha},\ \alpha)$). We then insert β in (6.73) to obtain $\tilde{q}$. With β and $\tilde{q}$ we calculate $E[\mathcal{L}]$ (6.46) and $E[\mathcal{S}]$ (6.67).

Performance measures

We define the mean throughput as the mean number of transmitting nodes in a slot. The mean throughput from a given (fixed) AWG input port to a given (fixed) AWG output port is then given by

$$TH_{\text{port}} = \frac{F \cdot E[\mathcal{L}] + K \cdot E[\mathcal{S}]}{F \cdot D}. \tag{6.75}$$

Note that the throughput given in (6.75) may also be interpreted as the average number of transmitted data packets per frame; for convenience we will use this packets/frame interpretation in our numerical work in the next section. The mean aggregate throughput of the network is

$$TH_{\text{net}} = D^2 \cdot TH_{\text{port}}. \tag{6.76}$$

We note that in case $\tilde{\alpha} = \alpha$, i.e., $\sigma = p$, we have from (6.32) $\beta = \alpha$. Inserting this in (6.70) gives an equation for $\tilde{q}$, which we solve numerically.

We now espouse the mean packet delay in the network. We define the mean delay as the average time period in cycles from the generation of the control packet corresponding to a data packet until the transmission of the data packet. Recall that $E[\mathcal{L}]+E[\mathcal{S}]$ is the expected number of data packets that the nodes at a given AWG input port transmit to the nodes at a given AWG output port per cycle. Now, consider a given (fixed) node m, $1 \leq m \leq N$. In the assumed uniform packet traffic scenario, this node m transmits on average $(E[\mathcal{L}]+E[\mathcal{S}])/S$ data packets to the nodes at a given AWG output port per cycle. Thus, node m transmits on average $(E[\mathcal{L}]+E[\mathcal{S}])D/S$ data packets to the N nodes attached to the D AWG output ports per cycle. The average time period in cycles from the generation of a control packet at node m until the generation of the next control packet is therefore $S/[D \cdot (E[\mathcal{L}] + E[\mathcal{S}])]$. Note that the time period from the successful scheduling of a data packet until the generation of the control packet for the next data packet is geometrically distributed with mean $(1-\sigma)/\sigma$ cycles. Hence, the average delay in the network in cycles is

$$\text{Delay} = \frac{S}{D \cdot (E[\mathcal{L}] + E[\mathcal{S}])} - \frac{1-\sigma}{\sigma}. \tag{6.77}$$

where $E[\mathcal{L}]$ and $E[\mathcal{S}]$ are known from the evaluation of the throughput (6.75).

6.2.3 Results

In this section, we show the benefit of spatial wavelength reuse and the impact of the system parameters on the throughput–delay performance of the network. Data packets can have one of two lengths. A data packet is F slots long with probability q and $K = (F - M)$ slots long with probability $(1 - q)$. Specifically, we consider number of used FSRs R, fraction of long data packets q, number of reservation slots per frame M, physical degree of the AWG D, number of nodes N, and retransmission probability p. All results presented in this section assume a channel spacing of 200 GHz, i.e., 1.6 nm at 1.55 μm. To avoid tuning penalties we deploy fast tunable transceivers whose tuning range is typically $10 - 15$ nm. Thus, the number of used wavelengths is assumed to be eight for all subsequent results. By default the parameters take on the following values: $R = 2$, $q = 0.25$, $M = 30$, $D = 4$, $N = 200$, $p = 0.8$, $F = 200$, and $K = 170$. Each cycle is assumed to have a constant length of $D \cdot F = 800$ slots. All numerical results in this section are obtained using the expression (6.36), which approximates the number of successful control packets in the slotted ALOHA contention by a Poisson distribution. (A numerical evaluation of the refined approximation, that does not use the Poisson distribution, but uses directly the binomial distribution, can be found in Appendix C. In summary, we find that using the approximate expression (6.36) gives very accurate results for a wide range of parameter values, as is also demonstrated by the numerical results in this section.) We also provide extensive simulation results of a more realistic network in order to verify the accuracy of the analysis. As opposed to the analysis, in the simulation a given node cannot transmit data packets to itself and both length and destination of a given data packet are not renewed

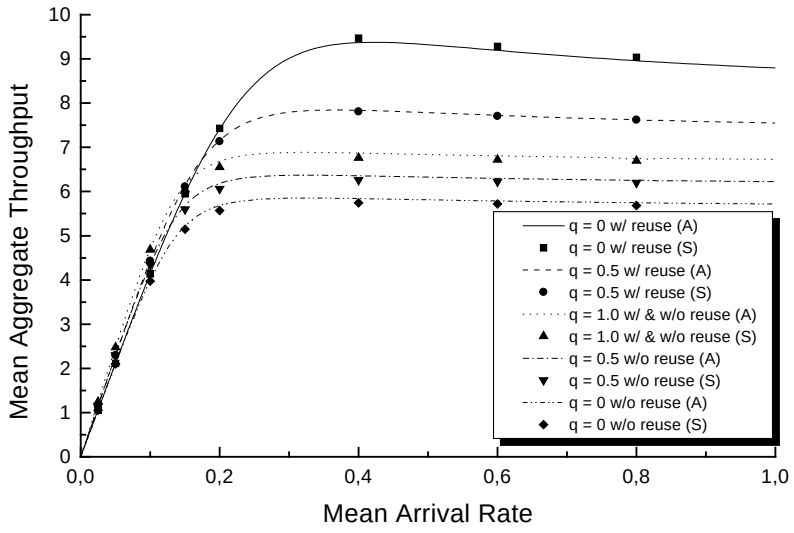

Figure 6.10: Mean aggregate throughput (packets/frame) vs. mean arrival rate with and without wavelength reuse for different fraction of long data packets $q \in \{0, 0.5, 1.0\}$.

when retransmitting the corresponding control packet. In addition, the simulation takes receiver collisions into account, i.e., a given data packet is not scheduled if the receiver of the intended destination node is busy. Each simulation was run for 10^7 slots including a warm–up phase of 10^6 slots. Using the method of batch means we also calculated the 98% confidence intervals for the mean aggregate throughput and the mean delay whose variations from the sample mean were less than 1% for all simulation results.

Figs. 6.10 and 6.11 illustrate that spatial wavelength reuse dramatically improves the throughput–delay performance of the network for variable–size data packets. Fig. 6.10 shows the mean aggregate throughput vs. the mean arrival rate σ with and without spatial wavelength reuse for different fraction of long data packets $q \in \{0,\ 0.5,\ 1.0\}$. Simulation and analysis results match very well. For $q = 1.0$, i.e., all data packets have a length of F slots, the mean aggregate throughput is the same no matter whether wavelengths are spatially reused or not. This is because the data packets are too long for being scheduled in the $(D-1)$ frames in which the corresponding nodes do not send control packets and spatial wavelength reuse would be possible in the last $(F-M)$ slots of the frame. Thus, these frames remain unused for $q = 1.0$. For $q = 0.5$, 50% of the data packets are long (F slots) and the other 50% are short (K slots). Allowing for spatial wavelength reuse the latter ones can now be scheduled in all frames including the aforementioned $(D-1)$ frames. Consequently, with wavelength reuse more data packets are successfully transmitted resulting in a higher throughput. In contrast, without wavelength reuse data packets can be scheduled only in one frame per cycle in which the corresponding nodes also transmit their control packets. Furthermore, since for $q = 0.5$ some successfully transmitted data packets are short (K slots), wavelengths are not fully utilized resulting in a lower throughput compared to $q = 1.0$. For $q = 0$ the benefit of spatial wavelength reuse becomes even more dramatic. In this case there are only short data packets (K slots) which fill up a large number

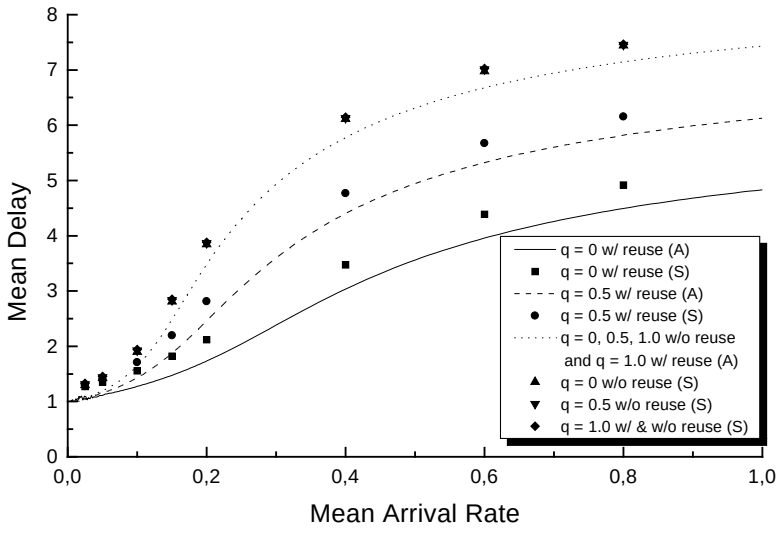

Figure 6.11: Mean delay (cycles) vs. mean arrival rate with and without wavelength reuse for different fraction of long data packets $q \in \{0, 0.5, 1.0\}$.

of frames leading to a further increased mean aggregate throughput. Note that for $q = 0$ spatial wavelength reuse significantly increases the maximum aggregate throughput by more than 60%. All curves in Fig. 6.10 run into saturation since for increasing σ no additional data packets can be scheduled due to busy channels and receivers and an increasing number of colliding control packets.

Fig. 6.11 depicts the mean delay vs. σ with and without wavelength reuse for different fraction of long data packets $q \in \{0,\ 0.5,\ 1.0\}$. We observe that the simulation provides slightly larger delay values than the analysis. (Similar observations will be made in the subsequent results as well.) This is because the simulation takes also the transmission time of data packets into account as opposed to the analysis. In the analysis the mean delay is defined as the time interval between the generation of a given data packet and the end of the cycle in which the given data packet is successfully scheduled but not yet transmitted. All curves have in common that at very light traffic the mean delay is equal to one cycle owing to the propagation delay of the control packet. With increasing σ the mean delay increases due to more unsuccessful control packets. These control packets have to be retransmitted resulting in an increased mean delay. Note that we obtain the largest delay if the aforementioned $(D-1)$ frames per cycle cannot be used for data transmission. This holds not only for the cases where wavelength reuse is not deployed but also for $q = 1.0$ with spatial wavelength reuse. This is due to the fact that for $q = 1.0$ the data packets are too long and do not fit in the last $(F-M)$ slots of the aforementioned $(D-1)$ frames. As a consequence, for these cases fewer data packets can be successfully scheduled and the corresponding control packets have to be retransmitted more often, leading to a higher mean delay. (Note that all cases without spatial wavelength reuse ($q = 0$, 0.5, and 1 w/o reuse as well as $q = 1$ w/ reuse) give the same delay. This is because $\lfloor F/K \rfloor = 1$ in the considered scenario, thus each packet occupies essentially one frame and

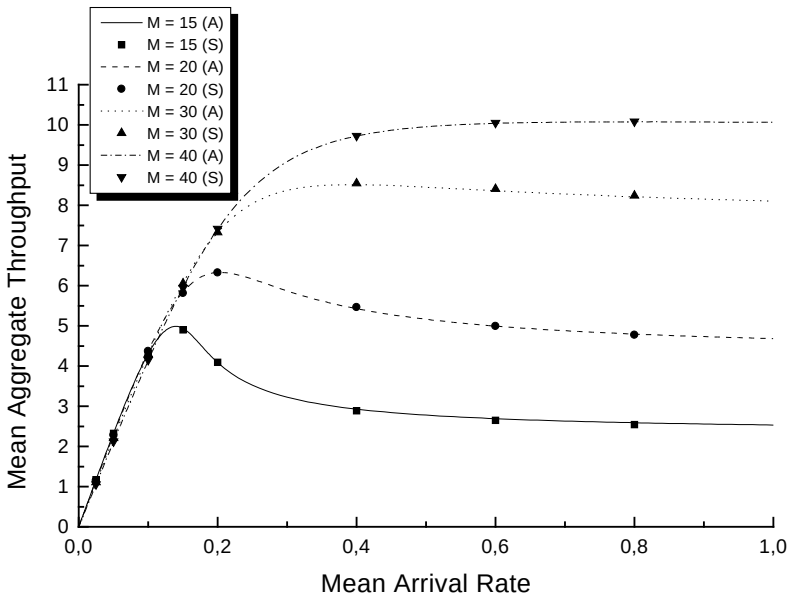

Figure 6.12: Mean aggregate throughput (packets/frame) vs. mean arrival rate for different number of reservation slots $M \in \{15, 20, 30, 40\}$.

the packet delay is independent of the fraction of short packets.) With decreasing q there are more short data packets which can easily be scheduled in the aforementioned $(D-1)$ frames if wavelengths are spatially reused. Due to the resulting wavelength reuse more data packets can be successfully scheduled. Therefore, fewer control packets have to be retransmitted leading to a decreased mean delay. In particular, for $q = 0$, wavelengths are used very efficiently resulting in the lowest delay. Note that for $q = 0$ spatial wavelength reuse reduces the mean delay by approximately 50%.

The impact of the number of reservation slots M per frame on the network throughput–delay performance is shown in Figs. 6.12 and 6.13. The mean aggregate throughput and the mean delay are depicted as a function of σ for $M \in \{15,\ 20,\ 30,\ 40\}$. Recall that by default the frame length F is set to 200 slots. Each frame is composed of M reservation slots and $K = (F - M)$ slots which can be used for transmitting short packets by means of spatial wavelength reuse. Clearly, for a fixed F increasing M decreases the length of short packets K, and vice versa for decreasing M. As shown in Figs. 6.12 and 6.13, increasing the number of reservation slots significantly improves the throughput–delay performance of the considered network. Due to the reduced contention more control packets are successfully transmitted resulting in an increased mean aggregate throughput and a decreased mean delay. Thus, in terms of the throughput–delay performance it is advantageous to use more reservation slots per frame even though this implies a smaller K. This also indicates that the random access reservation scheme can be a severe bottleneck. Note that the network throughput–delay performance could be easily improved by using additional spreading sequences and/or replacing the random access of the reservation slots with a dedicated assignment of the reservation slots. However, such a dedicated assignment does not scale very well.

A given number of nodes can be connected by AWGs with different physical degree D. Figs.

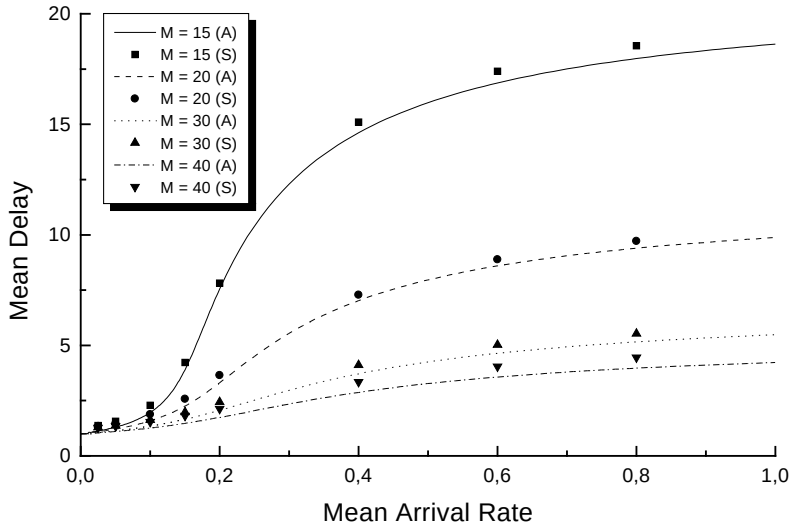

Figure 6.13: Mean delay (cycles) vs. mean arrival rate for different number of reservation slots $M \in \{15, 20, 30, 40\}$.

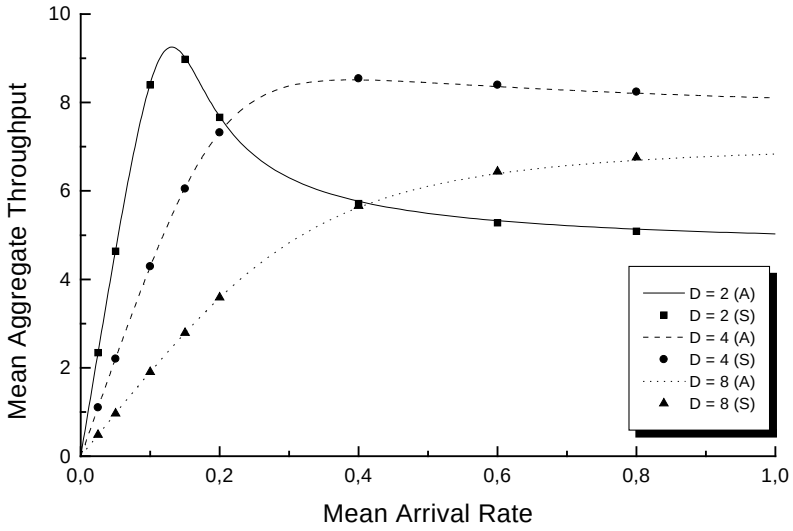

Figure 6.14: Mean aggregate throughput (packets/frame) vs. mean arrival rate for different AWG degree $D \in \{2, 4, 8\}$.

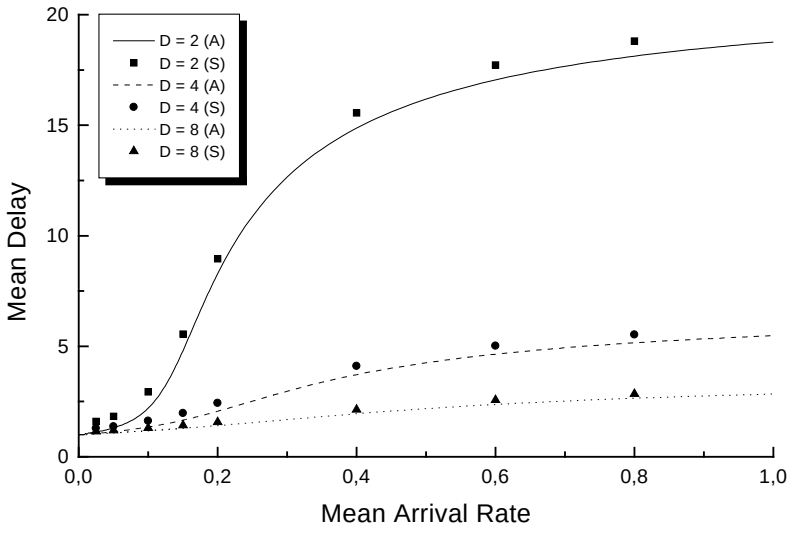

Figure 6.15: Mean delay (cycles) vs. mean arrival rate for different AWG degree $D \in \{2, 4, 8\}$.

6.14 and 6.15 depict for $D \in \{2,\ 4,\ 8\}$ the mean aggregate throughput and the mean delay as a function of σ, respectively. Recall that we have chosen the transceiver tuning range and the channel spacing such that we make use of eight wavelengths. The number of used FSRs R is then determined only by the physical degree D of the underlying AWG and is given by $R = 8/D$. Consequently, for a smaller D more FSRs are exploited, and vice versa for a larger D. Furthermore, for a smaller D each cycle contains fewer but longer frames, and vice versa for a larger D.

As shown in Fig. 6.14, $D = 2$ provides the largest maximum mean aggregate throughput at light traffic. However, with increasing σ the mean aggregate throughput decreases. This is due to the fact that for $D = 2$ short data packets are rather long ($K = 800/D - M = 370$ slots) resulting in a higher channel utilization and thereby a higher throughput at small traffic loads. But a small D also implies that for a given population N more nodes are attached to the same combiner since $S = N/D$. All these S nodes make their reservations in the same frame. For an increasing σ this leads to more collisions of control packets resulting in a lower mean aggregate throughput and a higher mean delay due to more retransmissions of the corresponding control packets (Fig. 6.15).

This problem is alleviated by deploying a 4×4 or 8×8 AWG. For a larger D fewer nodes send control packets in the same frame causing fewer collisions at high traffic loads. However, for $D = 4$ and $D = 8$ only 2 FSRs and 1 FSR can be exploited, respectively. Moreover, a larger D reduces the length of short data packets as well. Fig. 6.14 shows that for $D = 4$ the mean aggregate throughput is rather high for a wide range of σ. Whereas for $D = 8$ the throughput is rather low due to the small number of control packets per frame and the low channel utilization owing to the reduced length of short data packets. Note that for $D = 4$ the mean aggregate throughput gradually decreases for increasing σ. This is because at high traffic loads control packets suffer from collisions and have to be retransmitted, resulting in a slightly higher mean

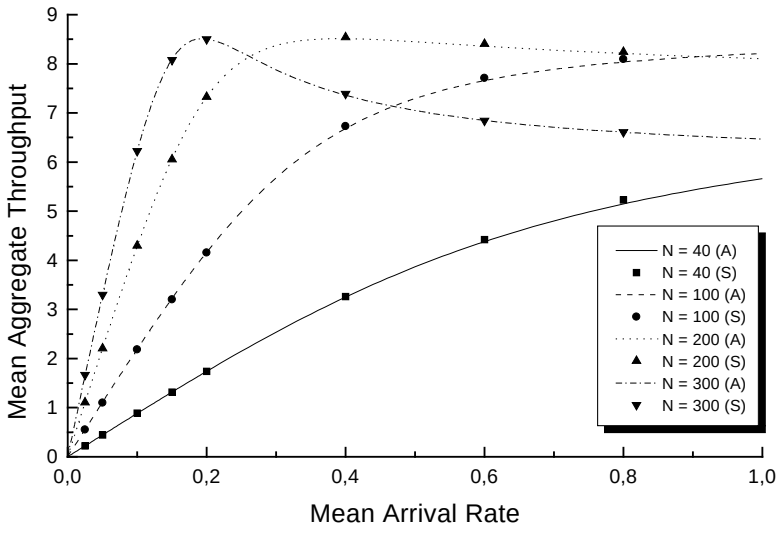

Figure 6.16: Mean aggregate throughput (packets/frame) vs. mean arrival rate for different population $N \in \{40, 100, 200, 300\}$.

delay compared to $D = 8$. Concluding, in terms of throughput–delay performance choosing $D = 4$ seems to provide the best solution for a wide range of traffic loads.

Figs. 6.16 and 6.17 depict the throughput–delay performance of the network for different population $N \in \{40,\ 100,\ 200,\ 300\}$. As shown in Fig. 6.16, increasing N improves the mean aggregate throughput due to more reservation requests and successfully scheduled data packets. However, for $N = 200$ and especially $N = 300$ the throughput decreases for increasing σ. This is because for large populations more control packets suffer from channel collisions resulting in a lower mean aggregate throughput. Accordingly, this leads to higher mean delays as shown in Fig. 6.17. Note that simulation and analysis results match very well even for small populations despite the fact that (i) we have conducted an asymptotic analysis for large S, and (ii) the analysis does not take receiver collisions into account (while the simulation does).

The impact of different AWG degree $D \in \{2,\ 4,\ 8\}$ on the system throughput–delay performance is shown in Figs. 6.18 and 6.19, respectively, for $\sigma = 0.4$. Note that throughput and delay are not given as a function of σ but as a function of the fraction of long data packets $q \in [0,\ 1]$. Recall that we have assumed a constant cycle length of 800 slots and a fixed number of reservation slots $M = 30$. As a consequence, the frame length is given by $F = 800/D$ slots and the length of short data packets is equal to $K = F - M = 800/D - 30$ slots. Moreover, the number of used FSRs of the underlying AWG is given by $R = 8/D$.

Fig. 6.18 depicts the mean aggregate throughput vs. q. For $D = 4$ we observe that the throughput monotonously decreases for increasing q. For $q = 0$ all data packets are short and can be scheduled in any frame resulting in a high mean aggregate throughput. For increasing q more and more data packets are long ($q = 1$ implies that there are only long data packets). However, long data packets can be scheduled only in one frame per cycle. In addition, at most two of them can be scheduled per AWG input–output port pair since $R = 2$. Consequently,

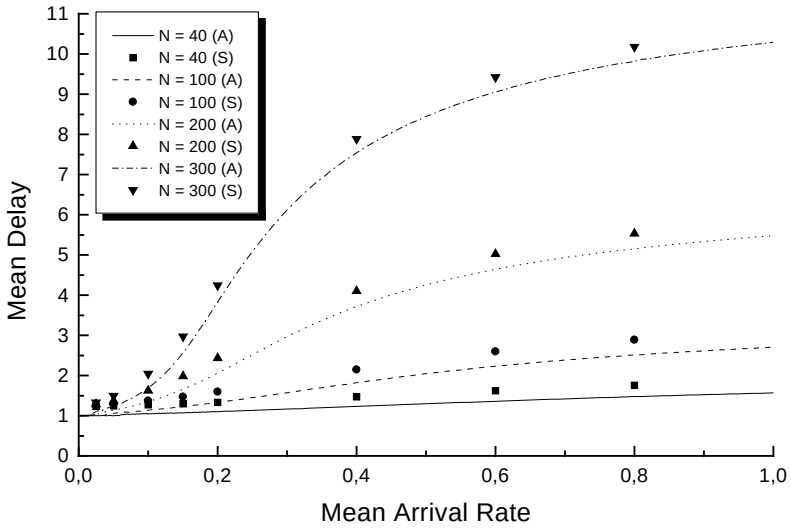

Figure 6.17: Mean delay (cycles) vs. mean arrival rate for different population $N \in \{40, 100, 200, 300\}$.

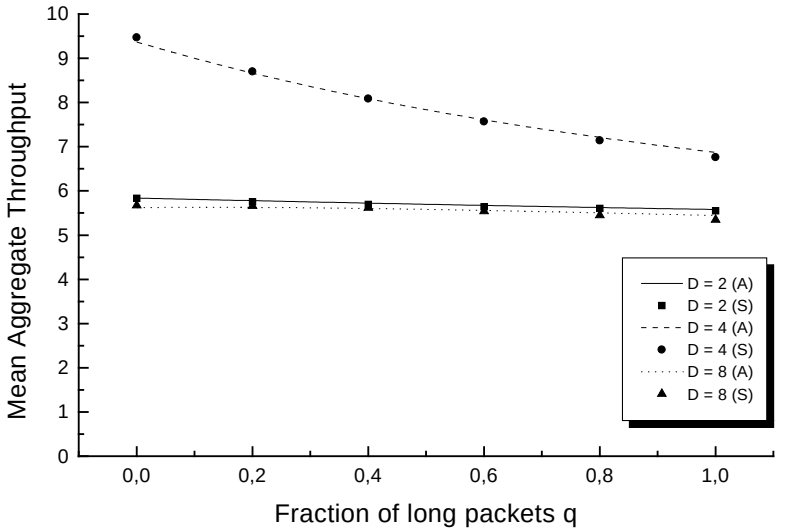

Figure 6.18: Mean aggregate throughput (packets/frame) vs. fraction of data packets q for different AWG degree $D \in \{2, 4, 8\}$.

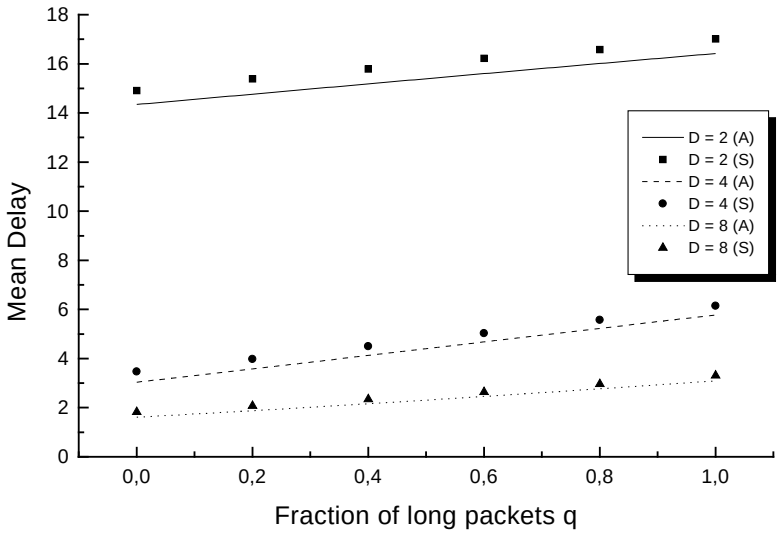

Figure 6.19: Mean delay (cycles) vs. fraction of data packets q for different AWG degree $D \in \{2, 4, 8\}$.

for increasing q fewer data packets can be scheduled resulting in a decreasing mean aggregate throughput (Fig. 6.18) and a higher mean delay as shown in Fig. 6.19. For $D \in \{2,\ 8\}$ the mean aggregate throughput is smaller than for $D = 4$ and, more interestingly, almost independent from q. For $D = 2$, twice as many nodes are attached to each AWG input port compared to $D = 4$. As a consequence, more control packets suffer from collisions and fewer data packets are available for scheduling, resulting in a smaller mean aggregate throughput. Moreover, there are not enough control packets to fully capitalize on spatial wavelength reuse. Thus, for $q = 0$ only slightly more data packets are successfully scheduled than for $q = 1$. Since for $q = 1$ all successfully scheduled data packets are long as opposed to $q = 0$ the mean aggregate throughput is about the same in both cases. Similarly, since for $D = 8$ fewer nodes are attached to each AWG input port there are fewer reservation requests per frame than for $D = 4$ resulting in a smaller throughput. However, these reservation requests experience fewer collisions significantly decreasing the mean delay as illustrated in Fig. 6.19. In contrast, $D = 2$ gives the highest mean delay due to the large number of collided control packets and their retransmissions. Note that for all $D \in \{2,\ 4,\ 8\}$ the mean delay grows with increasing q since for larger q fewer data packets can be scheduled owing to the lack of spatial wavelength reuse. This leads to more retransmissions of control packets and thereby to an increased mean delay. Concluding, while $D \in \{2,\ 8\}$ suffers from a relatively small throughput and $D = 2$ exhibits a large mean delay, choosing $D = 4$ seems to provide the best compromise in terms of throughput–delay performance.

Figs. 6.20 and 6.21 depict the throughput–delay performance of the network as a function of σ for different retransmission probability $p \in \{0.3,\ 0.6,\ 0.9\}$. As shown in Fig. 6.20, for $p = 0.3$ the mean aggregate throughput grows monotonously for increasing σ. We observe that at high traffic loads the slope decreases due to increasingly busy channels and transceivers. For $p = 0.6$ the mean aggregate throughput is larger than for $p = 0.3$. This is because with a larger

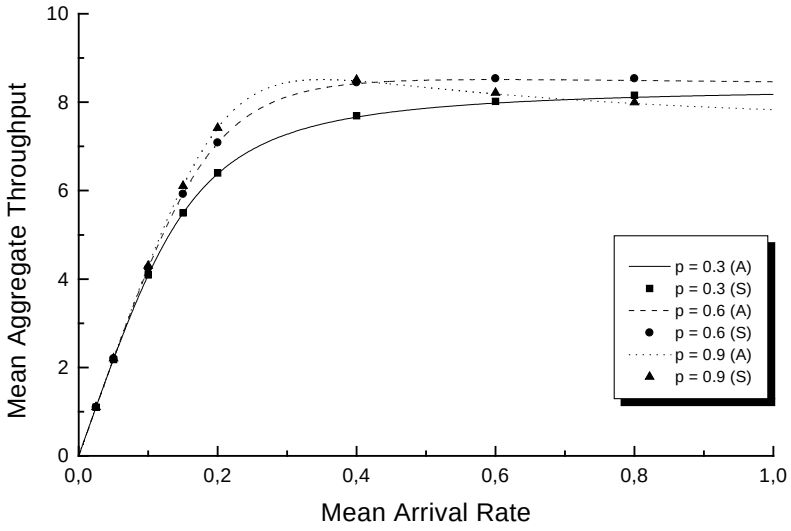

Figure 6.20: Mean aggregate throughput (packets/frame) vs. mean arrival rate for different retransmission probability $p \in \{0.3, 0.6, 0.9\}$.

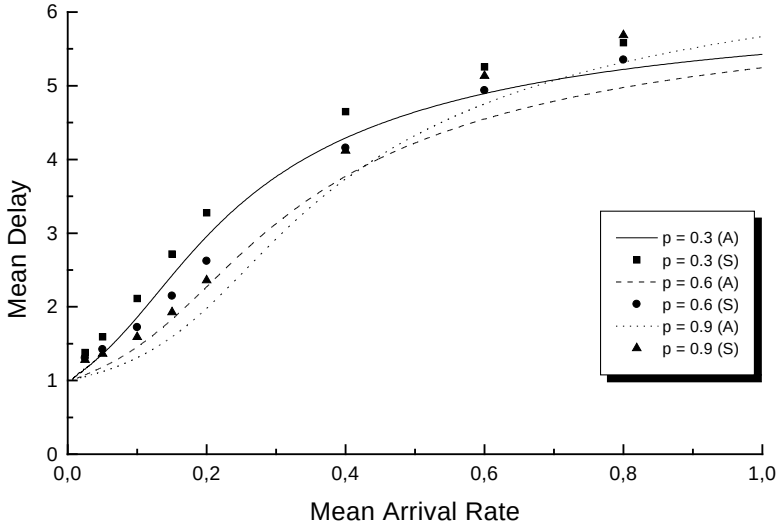

Figure 6.21: Mean delay (cycles) vs. mean arrival rate for different retransmission probability $p \in \{0.3, 0.6, 0.9\}$.

p nodes retransmit collided control packets with a higher probability resulting in more successful control packets and an increased mean aggregate throughput. However, further increasing p has a detrimental impact on the throughput. For $p = 0.9$ nodes retransmit collided control packets after a small time period. As a consequence, at medium to high traffic loads an increasing number of control packets collide leading to a decreased mean aggregate throughput. Fig. 6.21 shows that for all $\sigma \in (0,\ 1]$, $p = 0.3$ yield larger mean delays than $p = 0.6$. With a smaller p nodes defer retransmissions of collided control packets for a larger time interval which in turn increases the mean delay. Note that for $p = 0.9$ nodes experience the smallest mean delay at light to medium traffic loads. At high loads, $p = 0.9$ gives the largest mean delay due to the increasing number of retransmissions of control packets.

6.3 Multicasting

Future optical WDM networks have to support an increasing number of applications which involve multidestination traffic and require a large amount of bandwidth and/or QoS. Examples are teleconferences, video/content distribution, distributed games, mailing list, and news groups. With multicasting, a source node reaches multiple destinations by sending a single multicast data packet, instead of sending multiple unicast packets. Thus, multicasting can significantly increase the efficient resource (transmitter, channel) utilization for multidestination traffic and can improve the cost effectiveness, which is critical for metro networks. To date, multicasting in single–hop WDM networks was investigated only in PSC based networks, see [BM93b][BM95a][BM95b][RA94] [RA97][BSS97][SH97][SM97][TK98][Mod98b][Mod99][TSK99] [ORP00][LW00][LLC00][LW01a][BGL+01][KIS01][LHH02][HK02][HCT02]. In this section, we examine multicasting in our single–hop WDM network that is based on an AWG.

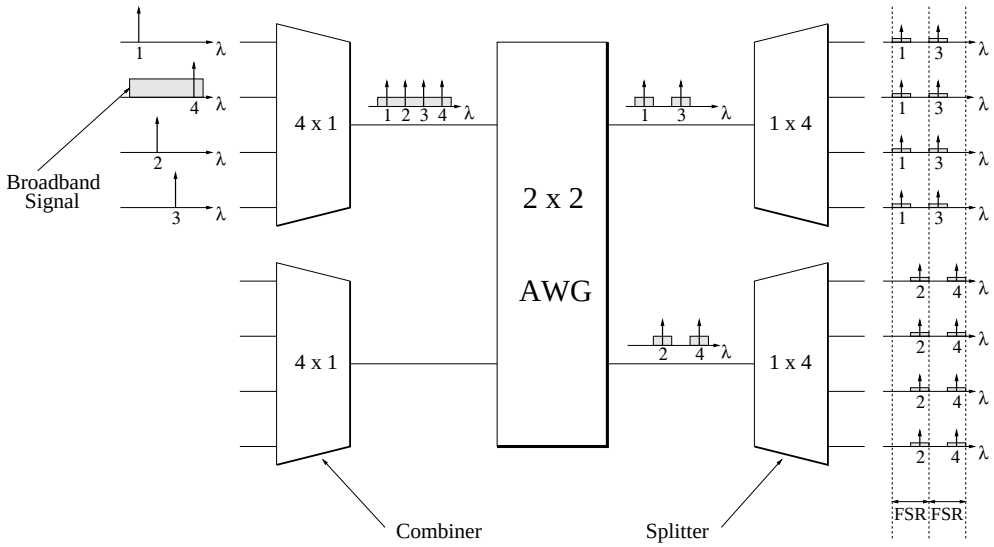

Figure 6.22: Multicasting and combining/splitting loss.

Fig. 6.22 shows how the splitters attached to the output ports of a 2×2 AWG enable optical multicasting. In this example the upper combiner collects four wavelengths and one broadband signal which are routed to both AWG output ports. Each splitter equally distributes all incoming wavelengths (and slices) to the attached receivers. The splitters enable optical

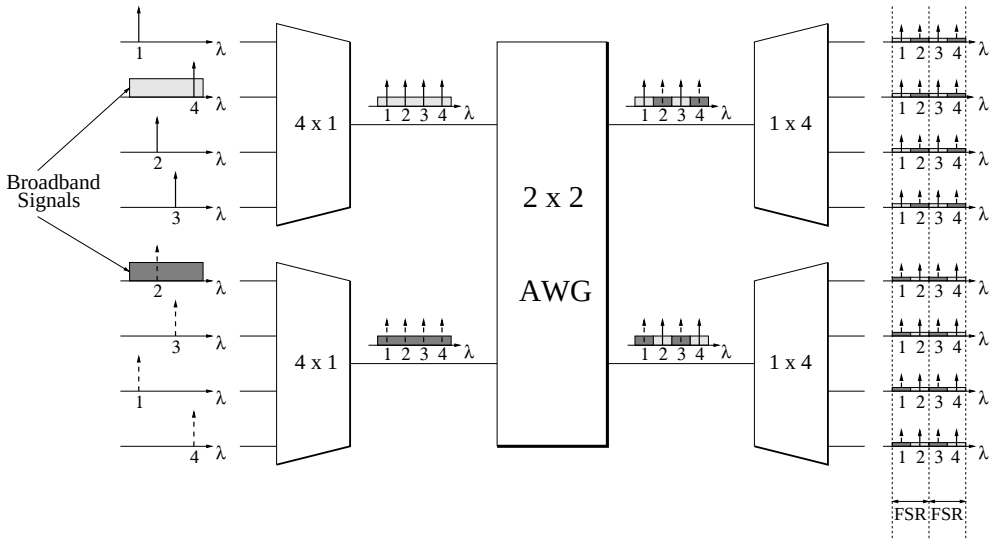

Figure 6.23: Multicasting with spatial wavelength reuse.

multicasting at the expense of splitting loss (similarly, each combiner suffers from combining loss) which might be compensated for by deploying optical amplifiers, as further discussed in Chapter 8. Note in Fig. 6.22 that a source node has to transmit on different wavelengths in order to reach receivers attached to different splitters. Each multicast copy is routed to a different splitter by sending it on a different wavelength. Thus, in general multicast transmissions whose receivers are located at multiple splitters have to be *partitioned*, i.e., the corresponding source node has to send multiple multicast copies (replicas), each on a different wavelength. It was shown in [JM97] that partitioning multicast transmissions achieves an improved throughput–delay performance. To see this, note that with partitioning each multicast copy requires a smaller number of receivers which are more likely to be free. Also, other transmitters can simultaneously send multicast packets to other free receivers. Thus, multiple wavelengths are used resulting in an improved network efficiency and throughput–delay performance. However, recently it was shown that partitioning suffers from a channel (wavelength) bottleneck in PSC based single–hop WDM networks [LW01b]. This is due to the fact that partitioning requires more channels (wavelengths). Since the PSC does not allow for spatial wavelength reuse, the number of available wavelengths is limited. In contrast, the AWG allows for spatial wavelength reuse. As a result, more wavelengths are available for simultaneous multicast transmissions from different AWG input ports, as illustrated in Fig. 6.23 for a 2×2 AWG.

Apart from spatial wavelength reuse, the efficiency of multicast communications can be further increased by exploiting our reservation MAC protocol. More precisely, recall from Section 5.3 that all receivers are obliged to tune to pre–specified wavelengths during the reservation phase of each frame in order to pick up the control traffic and attain global knowledge. Intuitively, multicast packets transmitted during this reservation time interval can be easily received by all nodes leading to a dramatically increased network efficiency and receiver utilization, as we will see shortly.

In this section, we quantify the benefit of (*i*) partitioning multicast transmissions with spatial wavelength reuse and (*ii*) transmitting multicast data packets concurrently with broadcast control packets. In Section 6.3.1 we consider only multicast packets and address partitioning. We demonstrate the positive impact of spatial wavelength reuse on the multicast performance

of the network by means of simulation. In Section 6.3.2 we investigate a mix of unicast and multicast traffic. We focus on multicast packets whose destination receivers are attached to one single splitter, i.e., no partitioning takes place. We present an analytical model which shows the throughput–delay performance gain achieved by sending multicast data and broadcast control simultaneously and examines the interplay between unicast and multicast traffic. The numerical results are verified by simulations.

6.3.1 Multicasting with Partitioning

In this section we study the transmission of multicast packet traffic over the AWG based network by means of simulation. We compare the throughput–delay performance of the AWG based network with the PSC based network. We consider the following commonly studied performance metrics:

- *Mean transmitter throughput* defined as the mean number of transmitting nodes in steady state.
- *Mean multicast throughput* defined as the mean number of multicast completions in steady state. (The mean multicast throughput is equal to the ratio of mean transmitter throughput and mean number of required transmissions in steady state in order to reach all receivers of a given multicast packet. Thus, multicast throughput measures the multicast efficiency of each packet transmission [ORP00].)
- *Mean receiver throughput* defined as the mean number of receiving nodes in steady state.
- *Mean delay* defined as the average time in frames from the generation of a packet until the completion of the multicast transmission.

Assumptions

For this study we consider only *multicast* traffic, i.e., each packet is destined to a multicast group. The size of the multicast group, i.e., the number of destination nodes, and the members of a given multicast group are independently randomly drawn for each packet. The multicast group size is *uniformly* distributed over $[1,\ N-1]$ nodes and the multicast group members are *uniformly* distributed over all network nodes $[1,\ N]$ except the transmitting source node, as is typically considered in multicast studies. The destination nodes of a given multicast packet are *persistent*, i.e., are not renewed when the corresponding control packet fails and is retransmitted. A given packet is *long* (occupies F slots) with probability $q,\ 0 \leq q \leq 1$, and is *short* (occupies $F-M$ slots) with the complementary probability $1-q$. Recall that on the AWG a long multicast packet from a node attached to AWG input port $o,\ 1 \leq o \leq D$, can only be scheduled in frame o of a given cycle, i.e., in the frame in which the node sends control packets. Short multicast packets from a node at port o can be scheduled in frame o as well as the other $(D-1)$ frames of a given cycle according to the adopted first–come–first–served–first–fit scheduling discipline.

For the simulations in this section the network parameters are set to the following default values: Number of nodes $N = 200$, the transceiver tuning range $D \cdot R = 8$ remains constant for varying D and R, retransmission probability $p = 0.5$, number of slots per frame $F = 200$, number of reservation slots per frame $M = 30$, and scheduling window size of 64 frames. (The scheduling window is set to 64 frames to ensure a fair comparison of the considered network configurations with $D = 2$, 4, and 8, of which the $D = 8$ configuration requires the largest scheduling window

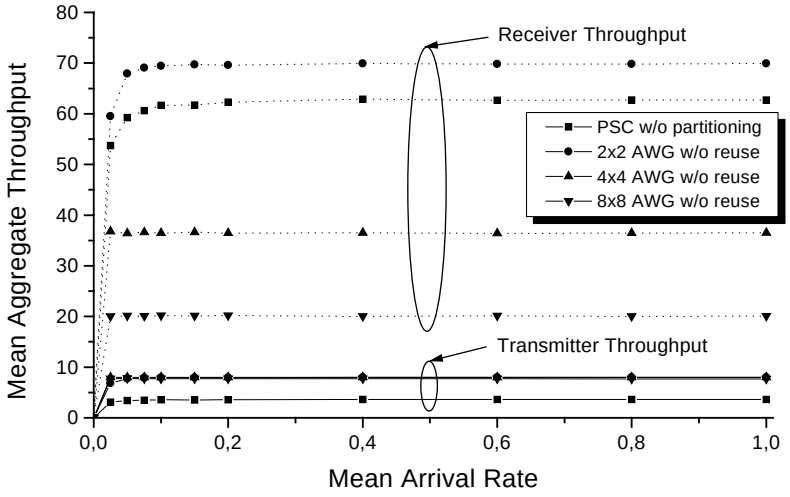

Figure 6.24: Mean transmitter and receiver throughput vs. mean arrival rate (packet/frame) for PSC (without partitioning) and AWG based single–hop networks without spatial wavelength reuse.

of 8 cycles which translates into 64 frames.) The propagation delay is assumed to be no larger than one frame. (The propagation delay is assumed to be no larger than one frame since in the PSC based network there is no cyclic timing structure as opposed to the AWG based network. In the PSC based network each node is assumed to be able to (re)transmit a control packet in every frame.) We assume that all nodes are equidistant from the central AWG (PSC), which can be achieved in practice with standard low–loss fiber delay lines. The mean arrival rate denotes the probability that an idle node generates a multicast packet at the end of a frame. Each simulation was run for 10^6 slots including a warm–up phase of 10^5 slots. The width of the 98% confidence intervals obtained with the method of batch means was always smaller than 5% of the corresponding sample means.

Simulation Results

In Figs. 6.24 and 6.25 we set $q = 1.0$, i.e., we consider only long packets ($L = F = 200$ slots) which cannot benefit from spatial wavelength reuse in the AWG based network. We compare the throughput–delay performance of a $D \times D$ AWG with $D \in \{2, 4, 8\}$ and a PSC based single–hop WDM network. For a fair comparison in both networks each node is equipped with the same pair of one tunable transmitter and one tunable receiver for data transmission. In the PSC based network, control is broadcast by using the inherent broadcast nature of the PSC. Each node is equipped with an additional transceiver fixed–tuned to a separate wavelength. Thus, in the PSC based network there are nine wavelengths, eight for data and one for control transmission. Nodes ready to (re)transmit control packets are allowed to randomly access $M = 30$ reservation slots in each frame of length $F = 200$ slots, using the same retransmission probability $p = 0.5$ as in the AWG based counterpart.

Fig. 6.24 depicts the mean transmitter and mean receiver throughput vs. mean arrival rate.

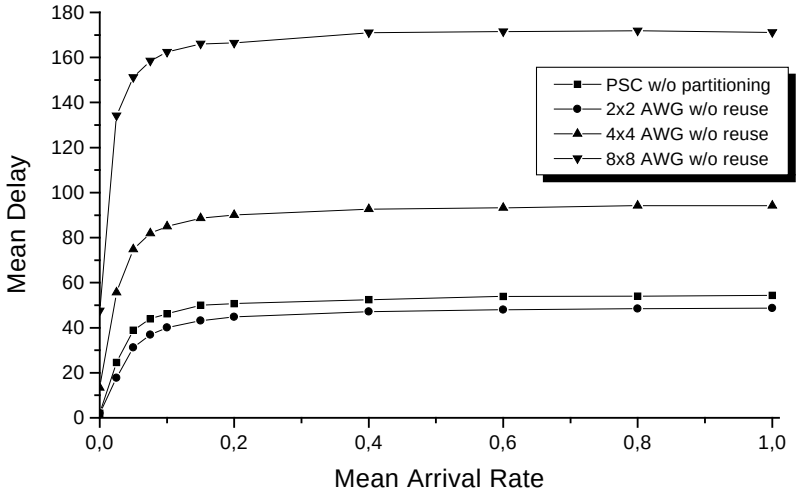

Figure 6.25: Mean delay (frames) vs. mean arrival rate (packet/frame) for PSC (without partitioning) and AWG based single–hop networks without spatial wavelength reuse.

We observe that the transmitter throughput in the AWG based network is about twice as large as in the PSC based network where the PSC network is assumed to operate without partitioning (the case where the PSC supports *logical* partitioning is discussed shortly). This is because due to its wavelength–routing nature the AWG provides (physical) partitioning such that nodes ready to send multicast packets are more likely to find free destination receivers for transmitting the corresponding multicast packets. Note that for all $D \in \{2, 4, 8\}$ the AWG provides the same transmitter throughput of eight. This is due to the fact that with a fixed transceiver tuning range of $D \cdot R = 8$ the number of available wavelength channels is limited such that additional transmissions cannot take place even though the corresponding destination receivers might be free. Hence, this figure confirms that partitioning can cause a channel bottleneck in the network. This channel bottleneck can be alleviated by spatial wavelength reuse, as discussed shortly. However, the physical AWG degree D has an impact on the receiver throughput, as depicted in Fig. 6.24. While a 2×2 AWG yields a larger receiver throughput than the PSC, for $D \in \{4, 8\}$ we observe the opposite. This is due to the channel bottleneck caused by partitioning. To see this, recall that for $D \in \{2, 4, 8\}$ the number of transmitting nodes is equal to the maximum number of available wavelength channels. For increasing D fewer nodes are attached to the same splitter. Consequently, each transmitted multicast copy is received by a smaller number of destination nodes, resulting in a decreased receiver throughput. (We do not show multicast throughput here since in the PSC based network without partitioning multicast and transmitter throughput are the same.)

Fig. 6.25 depicts the mean delay vs. mean arrival rate for the PSC and $D \times D$ AWG based single–hop WDM networks, where $D \in \{2, 4, 8\}$. Only the 2×2 AWG provides a smaller delay than the PSC. This is because in a 2×2 AWG based network with partitioning more nodes can transmit simultaneously than in the PSC based counterpart leading to a smaller delay. On the other hand, for $D \in \{4, 8\}$ the delay is significantly larger, since for increasing D multicast

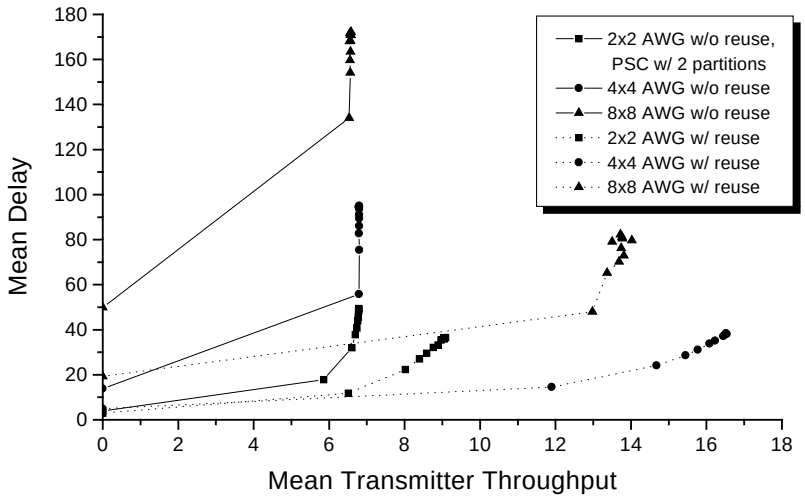

Figure 6.26: Mean delay (frames) vs. mean transmitter throughput for PSC (with 2 partitions) and AWG based single–hop networks with and without spatial wavelength reuse.

packets have to be sent to more splitters. Each of those multicast copies is transmitted in a separate cycle, each consisting of D frames. Therefore, with increasing D not only the average number of required multicast copies but also the cycle length increases, resulting in a larger delay.

In Figs. 6.26 through 6.28 we investigate the impact of spatial wavelength reuse on the transmitter, receiver, and multicast throughput–delay performance of the AWG based single-hop WDM network and compare it with the PSC based network. For this purpose, we set $q = 0$, i.e., all packets have a length of $F - M = 170$ slots which can be transmitted by spatially reusing all wavelengths. To assess the benefits of spatial wavelength reuse in a fair manner we consider partitioning not only in the AWG but also in the PSC based network. More precisely, in the PSC based network receivers are divided into two groups comprising nodes 1 through $N/2$ and $N/2$ through N, respectively. In doing so, the PSC based network operates with the same two partitions as a 2×2 AWG based network, which achieves the best throughput–delay performance without spatial wavelength reuse, as we have seen in Figs. 6.24 and 6.25. Figs. 6.26 through 6.28 are obtained by setting the mean arrival rate to $\{0,\ 0.025,\ 0.05,\ 0.075,\ 0.1,\ 0.15,\ 0.2,\ 0.4,\ 0.6,\ 0.8,\ 1.0\}$.

Fig. 6.26 illustrates the mean delay vs. mean transmitter throughput of both PSC and AWG based networks. Note that compared to Fig. 6.24 the maximum transmitter throughput of the AWG based network without spatial wavelength reuse is smaller than eight, since frames are not fully utilized due to the smaller packet size of $F - M$ slots. We observe that by allowing for spatial wavelength reuse the transmitter throughput–delay performance of all $D \times D$ AWG based networks is significantly improved with $D \in \{2, 4, 8\}$. Note that for $D = 2$ nodes cannot fully capitalize on the increased number of available wavelength channels. This is because with two partitions multicast copies destined to the same splitter are likely to experience receiver conflicts since on average each multicast copy is destined to more receivers for $D = 2$ than $D \in$

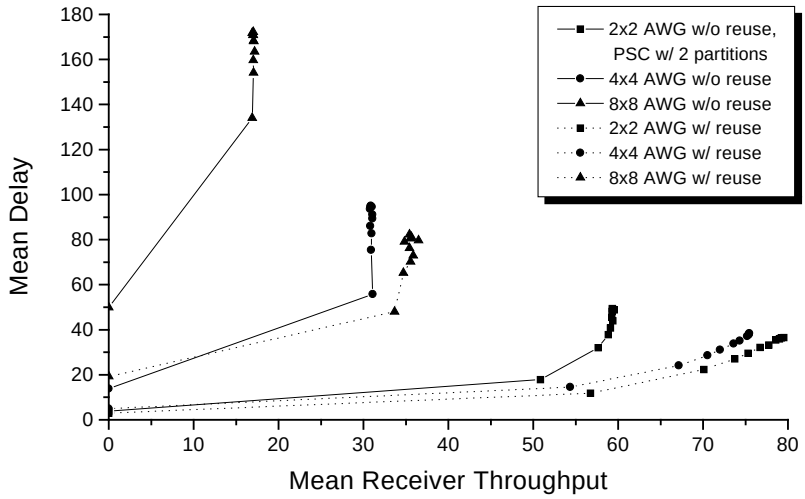

Figure 6.27: Mean delay (frames) vs. mean receiver throughput for PSC (with 2 partitions) and AWG based single–hop networks with and without spatial wavelength reuse.

$\{4, 8\}$. As a result, there are many destination conflicting multicast transmissions resulting in a modest transmitter throughput. The problem of destination conflicts is mitigated by dividing the receivers into more partitions. For $D = 4$ more transmitters are likely to find the corresponding receivers free resulting in a transmitter throughput, which is more than twice as large as the one of a 4×4 AWG based network without spatial wavelength reuse. Further increasing the number of partitions to $D = 8$ reduces the throughput, which is due to the smaller number of wavelength channels $R = 1$ connecting each individual AWG input–output port pair for the fixed transceiver tuning range $R \cdot D = 8$. Overall, we find that with spatial wavelength reuse a 4×4 AWG based network provides the smallest delay and the largest transmitter throughput which is more than twice that of a PSC based network, which operates with two partitions but does not allow for spatial wavelength reuse.

Figs. 6.27 and 6.28 show that spatial wavelength reuse also significantly improves the receiver and multicast throughput–delay performance of AWG based single–hop networks. Again, $D = 8$ is not a good choice to achieve an acceptable network performance. Whereas, $D = 4$ and $D = 2$ exhibit about the same receiver and multicast throughput–delay performance improvement. In terms of multicast throughput, i.e., the mean number of multicast completions, it is advisable to set $D = 2$. That is, with $D = 2$ the transmitter throughput is rather small (see Fig. 6.26) but each transmitted multicast copy is received by more intended destinations attached to the same splitter translating into an increased receiver throughput (see Fig. 6.27) and fewer required transmissions of a given multicast packet. Note that in terms of receiver and multicast throughput a 2×2 AWG based single-hop network outperforms its PSC based counterpart by approximately 30% where the latter one deploys the same partitioning but is unable to provide spatial wavelength reuse.

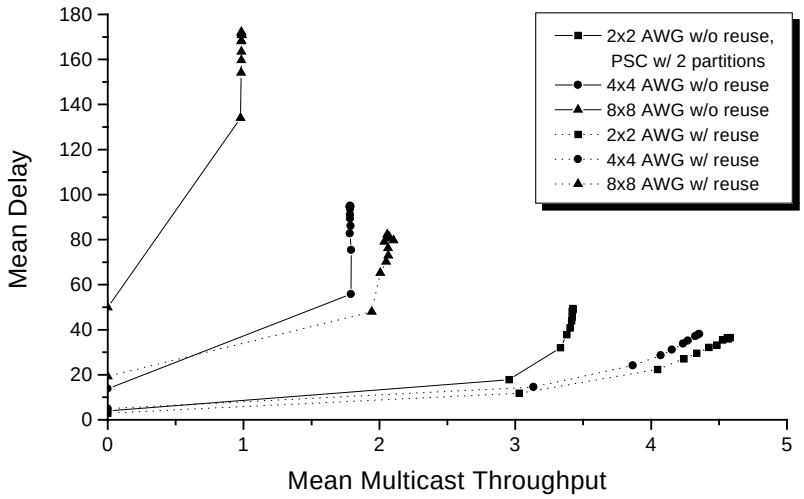

Figure 6.28: Mean delay (frames) vs. mean multicast throughput for PSC (with 2 partitions) and AWG based single–hop networks with and without spatial wavelength reuse.

6.3.2 Multicasting Simultaneously with Control

Up to this point we have considered only multicast packet traffic, i.e., each packet was destined to a random number of $1, \ldots, N-1$ nodes and we have examined the interplay between partitioning and spatial wavelength reuse. In contrast, in this section we analyze the transmission of a typical unicast and multicast traffic mix over the AWG based network. In this traffic mix a certain portion of the traffic is unicast while the remaining traffic is multicast. We focus on the interplay between unicast with spatial wavelength reuse and multicast concurrently with control traffic; we do not consider partitioning in this section. The motivation for this study is as follows. The results of the preceding section demonstrate that spatial wavelength reuse is beneficial for transmitting multicast traffic. Spatial wavelength reuse is not possible during the reservation phase, i.e., the first M slots of every frame when the control packets are transmitted. Thus, the reservation phase prevents the full exploitation of spatial wavelength reuse. Now consider the transmission of a typical mix of unicast and multicast traffic. For unicast traffic the wavelength channels are the primary bottleneck and receiver availability is typically not a problem; hence spatial wavelength reuse (which alleviates the channel bottleneck) brings dramatic benefit for unicast traffic, as shown in Section 6.2. Multicast traffic also benefits from spatial wavelength reuse, but typically receiver availability is its primary bottleneck. This suggests to schedule (i) multicast packets in the frame with the reservation phase, during which all receivers are tuned to a slice carrying the spread control traffic, thus alleviating the receiver availability problem, and (ii) unicast packets in the remaining frames where they can exploit spatial wavelength reuse. In this section we develop an analytical model to study the interplay between unicast and multicast traffic. We examine how spatial wavelength reuse and multicasting concurrently with control improve the overall throughput–delay performance of the AWG based network.

Assumptions

In our multicast analysis we make the same assumptions as in Section 6.2.1 except for packet type in that we consider not only unicast but also multicast data packets. Specifically, we make the following assumptions about the traffic generation at each node:

- *Uniform unicast* traffic: A data packet is short, i.e., occupies K slots long, with $1 \leq K \leq (F-M)$, and unicast with probability $p_{s,1}$. A data packet is long, i.e., occupies F slots and unicast with probability $p_{l,1}$. A unicast data packet – either short or long – is destined to any one of the N nodes (including the sending node, for simplicity) with equal probability $1/N$.
- *Clustered multicast* traffic: A short data packet is multicast with probability $p_{s,a}$. A long data packet is multicast with probability $p_{l,a}$. A multicast data packet – either short or long – is destined to all S nodes attached to any of the D splitters (including the splitter that the sending node is attached to) with equal probability $1/D$.
- *Packet size and type persistency*: The size and type (unicast or multicast) is not changed if the corresponding control packet is retransmitted.
- *Destination nonpersistency*: Random selection of a destination node (unicast) or destination splitter (multicast) is renewed for each attempt of transmitting a control packet.

Analysis

As a shorthand we refer to the four packet types (long, multicast), (long, unicast), (short, multicast), and (short, unicast) with the tuples (l,a), $(l,1)$, (s,a), and $(s,1)$. Thus, the probabilities $p_{l,a},\ p_{l,1},\ p_{s,a}$, and $p_{s,1}$ denote the probabilities that a newly generated packet is of type (l,a), $(l,1)$, (s,a), or $(s,1)$. Note that $p_{l,a}+p_{l,1}+p_{s,a}+p_{s,1}=1$.

Recall from Section 5.3 that nodes attached to a given (fixed) AWG input port o, $1 \leq o \leq D$, send their control packets in frame o of a given cycle. We refer to the nodes that at the beginning of frame o hold an old packet, that is, a control packet that has failed in slotted Aloha or scheduling, as *"old"*. We refer to all the other nodes as *"new"*. Note that the set of "new" nodes comprises both the nodes that have generated a new (never before transmitted) control packet as well as the nodes that have deferred the generation of a new control packet. Let η be a random variable denoting the number of "new" nodes at AWG input port o.

Let $\hat{p}_{l,a},\ \hat{p}_{l,1},\ \hat{p}_{s,a}$, and $\hat{p}_{s,1}$ denote the probabilities that a given node at port o is to send a control packet corresponding to a data packet of type (l,a), $(l,1)$, (s,a), or $(s,1)$ next. Again, note that $\hat{p}_{l,a}+\hat{p}_{l,1}+\hat{p}_{s,a}+\hat{p}_{s,1}=1$. We expect, for instance, that $\hat{p}_{l,a}$ is larger than $p_{l,a}$ since long multicast packets are more difficult to schedule than the other packet types and thus require more re–transmissions (of control packets).

Slotted ALOHA contention First, we calculate the probability κ that a given control slot out of the available M control slots in frame o contains a successful control packet. A given control slot contains a successfully transmitted control packet if either (i) it contains exactly one control packet corresponding to a newly generated data packet (from one of the "new" nodes) and no control packet from the "old" nodes, or (ii) it contains exactly one control packet from an "old" node and no control packet from a "new" node. Hence,

$$\kappa = \eta\frac{\sigma}{M}\left(1-\frac{\sigma}{M}\right)^{\eta-1}\left(1-\frac{p}{M}\right)^{S-\eta}+(S-\eta)\frac{p}{M}\left(1-\frac{p}{M}\right)^{S-\eta-1}\left(1-\frac{\sigma}{M}\right)^{\eta}, \qquad (6.78)$$

where we assume that the number of control packets from "new" nodes is independent of the number of control packets from "old" nodes.

Recall from our traffic model that each packet is destined to any one of the D AWG output ports with equal probability $1/D$. Thus, the number of control packets corresponding to (l,a), $(l,1)$, (s,a), $(s,1)$ data packets that (i) originate from a given AWG input port $o,\ o=1,\ldots,D$, (ii) are successful in the control packet contention of frame o (of a given cycle), and (iii) are destined to a given AWG output port $d,\ d=1,\ldots,D$, are distributed according to the binomial distributions $BIN(M,\ \kappa\hat{p}_{l,a}/D)$, $BIN(M,\ \kappa\hat{p}_{l,1}/D)$, $BIN(M,\ \kappa\hat{p}_{s,a}/D)$, and $BIN(M,\ \kappa\hat{p}_{s,1}/D)$, respectively.

Packet scheduling We now proceed to calculate the numbers of successfully scheduled packets. Recall that the numbers of packets to be considered for the schedule from a given AWG input port to a given AWG output port are distributed according to the binomial distributions given at the end of the preceding section. Let $X_{l,a},\ X_{l,1},\ X_{s,a}$, and $X_{s,1}$ be random variables denoting the number of packets of type (l,a), $(l,1)$, (s,a), and $(s,1)$ that (i) originate from a given AWG input port $o,\ o=1,\ldots,D$, (ii) are successful in the control packet contention of frame o (of a given cycle), (iii) are destined to a given AWG output port $d,\ d=1,\ldots,D$, and (iv) are successfully scheduled within the scheduling window of one cycle. We calculate $E[X_{l,a}],\ E[X_{l,1}],\ E[X_{s,a}]$, and $E[X_{s,1}]$ as functions of $\hat{p}_{l,a},\ \hat{p}_{l,1},\ \hat{p}_{s,a},\ \hat{p}_{s,1}$, and κ (which in turn is a function of η as given in Eqn. (6.78)).

The two critical resources (constraints) for the data packet scheduling are (i) the wavelength channels on the AWG, and (ii) the tunable receiver at each of the nodes:

Channel Constraint First, we examine the wavelength channel constraint. Consider the scheduling of packets from a given (fixed) AWG input port o to a given (fixed) AWG output port d over the scheduling window of D frames. Over this scheduling window the AWG provides R parallel wavelength channels during the long (F slot) transmission slot, i.e., during the frame in which the nodes at port o send their control packets. During each of the remaining $(D-1)$ frames, the AWG provides a short ($F-M$ slot) transmission slot; each again with R parallel wavelength channels.

Now, we consider the scheduling of the four different types of packets in these transmission slots. First, we consider each packet type in isolation. Clearly, we can schedule at most one (l,a)–packet during the scheduling window. To see this note, that a long packet can only be scheduled during the long transmission slot. Also, a multicast will occupy all receivers at the considered destination port d during the transmission slot. Formally, we let $a,\ a=0,\ 1$, denote the number of scheduled (l,a)–packets.

Next, we consider $(l,1)$–packets and let b denote the number of scheduled $(l,1)$–packets. Long packets can again only be scheduled during the long transmission slot. For unicast packets we ignore receiver collision, as we have seen in Section 6.2 that their impact is typically small. Hence, $0 \le b \le R$.

Packets of type (s,a) could be scheduled in the long transmission slot as well as in the short transmission slots. Whereas (s,a)–packets we schedule only in the long transmission slot and let c denote the number of scheduled (s,a)–packets. Note that an (s,a)–packet occupies all receivers at the considered destination splitter for a duration of K slots. Hence, $0 \le c \le \lfloor F/K \rfloor$.

Finally, note that $(s,1)$–packets can be scheduled in both the long and the short transmission slots. We let d denote the number of $(s,1)$–packets that are scheduled in the scheduling window

of D frames. Clearly,

$$0 \leq d \leq \left\{ \left\lfloor \frac{F}{K} \right\rfloor + (D-1) \cdot \left\lfloor \frac{F-M}{K} \right\rfloor \right\} \cdot R. \tag{6.79}$$

We have considered the scheduling of one packet type in isolation so far. To complete our model we need to consider the scheduling of combinations of the different packet types as well as the receiver collisions. Note that receiver collisions due to multicast packets of a given type ((l,a) or (s,a)) from a given AWG input port are accounted for in the above limits for a and b. We examine the receiver collisions due to transmissions by the other ports in the next section and return to the scheduling of combinations of different packet types and receiver collisions due to transmissions from the same port shortly.

Receiver Constraint In our analytical model of the data packet scheduling we account for receiver collisions due to multicast packets. We allow multicast packets to be scheduled from the nodes at a given AWG input port o to the receivers at a given AWG output port d at a given time only if there is not already a multicast or unicast packet from the same input port o or another input port $o' \neq o$ scheduled to output port d at the considered time. Receiver collisions due to the packets from the considered input port o are accounted for in the channel constraints discussed above and the schedulability conditions derived below. In this section the focus is on how the transmissions from the other input ports o' to the considered destination port d interfere with the transmissions from the considered input port o to port d. We note that throughout our analysis we ignore receiver collisions due to unicast packets, i.e., when scheduling a unicast packet we do not verify whether there is already another unicast packet (from the same AWG input port or a different input port) destined to the same destination port at the same time.

Recall from Section 5.3 that the nodes at AWG input port o, $o = 1, \ldots, D$, send their control packets in frame o of a given cycle. Suppose that a sent control packet is successful in the control packet contention. Then, we attempt to schedule the corresponding data packet in the scheduling window that extends from frame o of the next cycle up to and including frame $(o-1)$ of the cycle thereafter, as illustrated in Fig. 6.29. Note that we assume here that the propagation delay is less than one cycle. Also, note that in case $o = 1$, the scheduling window coincides with the cycle boundaries. Now, consider the scheduling window from frame o to frame $(o-1)$ more closely. It consists of the long (F slot) transmission slot and $(D-1)$ short ($F-M$ slot) transmission slots. Packets of types (l,a), $(l,1)$, and (s,a) are only scheduled in the long transmission slot. Packets of type $(s,1)$, on the other hand, are scheduled in the long transmission slot as well as the subsequent $(D-1)$ short transmission slots. (Recall that for scalability reasons the packet scheduling is done on a first–come–first–served and first–fit basis.)

Up to this point we have considered the scheduling of data packets from the nodes at a given AWG input port o, $o = 1, \ldots, D$, to the nodes at a given AWG output port d, $d = 1, \ldots, D$. Now, consider the scheduling of data packets from the nodes at the other AWG input ports o', $o' = 1, \ldots, D$, $o' \neq o$, to the nodes at AWG output port d. The scheduling windows of the other ports o' are staggered with respect to the scheduling window of port o, as illustrated in Figure 6.29 for the $(D-1)$ scheduling windows that precede the considered scheduling window of port o. In each of these preceding scheduling windows, the (l,a), $(l,1)$, and (s,a) packets are again scheduled in the first frame — the long transmission slot — and the $(s,1)$ packets are scheduled in the long transmission slot as well as the subsequent $(D-1)$ short transmission slots. As a consequence, the multicasts from the other ports o' do not interfere with the multicasts from the considered port o. However, $(s,1)$–packets from the other ports o' may have been

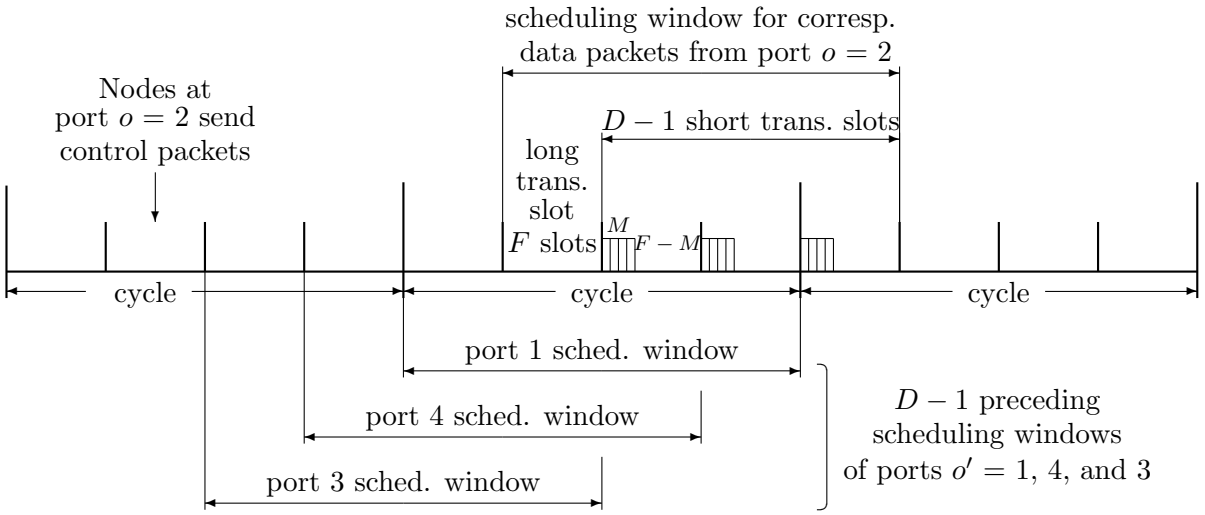

Figure 6.29: Illustration of scheduling of data packets from port $o = 2$ for $D = 4$.

scheduled during the long transmission slot of the scheduling window of port o. More precisely, $(s, 1)$–packets from nodes at the other ports o' may have been scheduled for receivers at the considered destination port d during the last $(F - M)$ slots of the long transmission slot of port o. These already scheduled $(s, 1)$–packets interfere with the scheduling of multicast packets from port o.

We model this interference as follows. We divide the last $(F - M)$ slots in the long transmission slot of the scheduling window of port o into *columns* of width K slots each. Similarly, we divide the $(D - 1)$ short transmission slots into columns of width K slots. Thus, there are $\xi := \lfloor (F-M)/K \rfloor$ columns in the long transmission slot and each of the short transmission slots. We refer to a column as *occupied* if in the $(D - 1)$ preceding scheduling windows of the other ports o' at least one $(s, 1)$–packet has been scheduled in the column. Note that only the columns in the long transmission slot and the first $(D - 2)$ short transmission slots of the scheduling window of port o can be occupied. Port o is the first to schedule data packets in the last short transmission slot of its scheduling window. Formally, we let C be a random variable denoting the number of occupied columns in a given scheduling window of port o. We let

$$\pi_l = P(C = l), \quad 0 \leq l \leq (D - 1)\xi, \tag{6.80}$$

denote the steady state probability that l columns are occupied. We will evaluate the steady state distribution π_l, $l = 0, \ldots, (D - 1)\xi$, from a Markov chain model developed shortly.

For the scheduling of the multicast packets from port o we need to take the number of occupied columns in the long transmission slot (i.e., the first frame) of the scheduling window of port o into consideration. Multicast packets can not be scheduled in any occupied columns. Formally, let Γ denote the number of occupied columns in the long transmission slot of the scheduling window of port o. With the considered first–come–first–served and first–fit scheduling policy the packets from each port are scheduled as early in the respective scheduling windows

as possible. Hence,

$$\Gamma = \min(C,\ \xi). \tag{6.81}$$

We now return to the analysis of the scheduling of combinations of different types of packets. We consider the cases of $\Gamma = 0$ occupied columns and $\Gamma \geq 1$ occupied columns separately.

Scheduling with $\Gamma = 0$ Occupied Columns We denote the scheduling of combinations of packet types by the 5–tuple $(i,\ a,\ b,\ c,\ d)$, which we refer to as *scheduling pattern*. The first element i in the scheduling pattern denotes the index up to which the control slots in the considered frame o have been inspected. Recall that the considered scheduling policy scans the control slots in increasing order of the index, that is, from $i = 1$ to $i = M$. If a control slot is empty or contains two (or more) collided control packets, then no data packet is scheduled. If a control slots contains exactly one control packet, that control packet is considered successful in the control packet contention and we attempt to schedule the corresponding data packet. If the data packet can be scheduled then the corresponding counter a, b, c, or d is incremented by one. If the data packet can not be scheduled (because there are not sufficient free channel and/or receiver resources) then the data packet fails in the scheduling and the counters a, b, c, and d remain unchanged. In summary, the scheduling pattern $(i,\ a,\ b,\ c,\ d)$ indicates that the control slots up to index $i,\ i = 1, \ldots, M$, have been scanned and a packets of type (l, a), b packets of type $(l, 1)$, c packets of type (s, a), and d packets of type $(s, 1)$ have been successfully scheduled.

We now establish *schedulability conditions* to verify whether a given scheduling pattern is feasible. The first schedulability condition is

$$a + b + c + d \leq i. \tag{6.82}$$

Clearly, when we have scanned i control slots we can not have scheduled more than i packets.

The second schedulability condition is

$$a = 1,\ b = 0,\ c = 0,\ 0 \leq d \leq (D-1) \cdot \left\lfloor \frac{F-M}{K} \right\rfloor \cdot R. \tag{6.83}$$

The third schedulability condition is

$$a = 0,\ 1 \leq b \leq R,\ c = 0,\ 0 \leq d \leq (R-b) \cdot \left\lfloor \frac{F}{K} \right\rfloor + (D-1) \cdot \left\lfloor \frac{F-M}{K} \right\rfloor \cdot R. \tag{6.84}$$

The fourth schedulability condition is

$$a = 0,\ b = 0,\ 0 \leq c \leq \left\lfloor \frac{F}{K} \right\rfloor,\ 0 \leq d \leq \left\{ \left\lfloor \frac{F}{K} \right\rfloor - c + (D-1) \cdot \left\lfloor \frac{F-M}{K} \right\rfloor \right\} \cdot R. \tag{6.85}$$

We refer to a scheduling pattern $(i,\ a,\ b,\ c,\ d)$ that satisfies the first schedulability condition (Eqn. (6.82)) *and one out of* the schedulability conditions (6.83), (6.84), (6.85) as *feasible*. Let $P^i_{a,b,c,d}$ denote the probability that the scheduling pattern $(i,\ a,\ b,\ c,\ d)$ arises. For all feasible scheduling patterns we calculate $P^i_{a,b,c,d}$ with the recursion

$$\begin{aligned} P^i_{a,b,c,d} \ = \ & P^{i-1}_{a,b,c,d} \cdot \left\{ \left(1 - \frac{\kappa}{D}\right) + \frac{\kappa}{D}(\alpha + \beta + \gamma + \delta) \right\} + \\ & P^{i-1}_{a-1,b,c,d} \cdot \frac{\kappa \hat{p}_{l,a}}{D} + P^{i-1}_{a,b-1,c,d} \cdot \frac{\kappa \hat{p}_{l,1}}{D} + P^{i-1}_{a,b,c-1,d} \cdot \frac{\kappa \hat{p}_{s,a}}{D} + P^{i-1}_{a,b,c,d-1} \cdot \frac{\kappa \hat{p}_{s,1}}{D}, \end{aligned} \tag{6.86}$$

where

$$\alpha = \begin{cases} 0 & \text{if } (i,\ a+1,\ b,\ c,\ d) \text{ is feasible, i.e., satisfies (6.82) and (6.83)} \\ \hat{p}_{l,a} & \text{otherwise.} \end{cases} \tag{6.87}$$

$$\beta = \begin{cases} 0 & \text{if } (i,\ a,\ b+1,\ c,\ d) \text{ is feasible, i.e., satisfies (6.82) and (6.84)} \\ \hat{p}_{l,1} & \text{otherwise.} \end{cases} \tag{6.88}$$

$$\gamma = \begin{cases} 0 & \text{if } (i,\ a,\ b,\ c+1,\ d) \text{ is feasible, i.e., satisfies (6.82) and (6.85)} \\ \hat{p}_{s,a} & \text{otherwise.} \end{cases} \tag{6.89}$$

$$\delta = \begin{cases} 0 & \text{if } (i,\ a,\ b,\ c,\ d+1) \text{ is feasible, i.e., satisfies (6.82) and} \\ & \text{either (6.83), (6.84), or (6.85)} \\ \hat{p}_{s,1} & \text{otherwise.} \end{cases} \tag{6.90}$$

We initialize this recursion with $P^0_{a,b,c,d} = 1$ if $a = b = c = d = 0$, and $P^0_{a,b,c,d} = 0$ otherwise, and note that all undefined $P^i_{a,b,c,d}$ (e.g., those with negative $a,\ b,\ c$, or d) are set to zero.

Scheduling with $\Gamma \geq 1$ Occupied Columns We assume throughout this section that

$$\left\lfloor \frac{M}{K} \right\rfloor + \left\lfloor \frac{F-M}{K} \right\rfloor = \left\lfloor \frac{F}{K} \right\rfloor. \tag{6.91}$$

If this condition is not satisfied, the analysis of the scheduling with occupied columns becomes more complicated since the specific order of the scheduling of the packets from the considered port o plays a role in the schedulability conditions (see the Appendix D for details). If (6.91) is satisfied, the schedulability conditions with $\Gamma \geq 1$ occupied columns are similar to the conditions discussed in the preceding section, with the differences that (i) $\Gamma \geq 1$ columns in the long transmission slot are not available to (s,a)–packets, and that (ii) (l,a) packets can not be scheduled. Thus, the first schedulability condition is as given by Eqn. (6.82). The second schedulability condition from the preceding section, Eqn. (6.83), is removed from consideration. The third schedulability condition is as given by Eqn. (6.84) since we ignore the receiver collisions due to $(s,1)$ packets from the other ports and $(l,1)$ packets from port o. The fourth schedulability condition is modified to

$$a = 0,\ b = 0,\ 0 \leq c \leq \left\lfloor \frac{F}{K} \right\rfloor - \Gamma,\ 0 \leq d \leq \left\{ \left\lfloor \frac{F}{K} \right\rfloor - c + (D-1) \cdot \left\lfloor \frac{F-M}{K} \right\rfloor \right\} \cdot R. \tag{6.92}$$

This condition accounts for the receiver collisions due to $(s,1)$ packets from the other ports and (s,a) packets from port o. The receiver collisions with $(s,1)$ packets from port o are again ignored.

We modify the definition of the scheduling pattern to the 5–tuple $(\Gamma,\ i,\ b,\ c,\ d)$ which indicates that given Γ occupied columns, the control slots up to index i, $i = 1, \ldots, M$, have been scanned, and b packets of type $(l,1)$, c packets of type (s,a), and d packets of type $(s,1)$ have been successfully scheduled.

We let ${}^{\Gamma}Q^i_{b,c,d}$ denote the probability that the scheduling pattern $(\Gamma,\ i,\ b,\ c,\ d)$ arises. For all feasible scheduling patterns we calculate ${}^{\Gamma}Q^i_{b,c,d}$ with the recursion

$${}^{\Gamma}Q^i_{b,c,d} = {}^{\Gamma}Q^{i-1}_{b,c,d} \cdot \left\{ \left(1 - \frac{\kappa}{D}\right) + \frac{\kappa}{D}(\hat{p}_{l,a} + \beta + \gamma + \delta) \right\} + \tag{6.93}$$

$${}^{\Gamma}Q^{i-1}_{b-1,c,d} \cdot \frac{\kappa \hat{p}_{l,1}}{D} + {}^{\Gamma}Q^{i-1}_{b,c-1,d} \cdot \frac{\kappa \hat{p}_{s,a}}{D} + {}^{\Gamma}Q^{i-1}_{b,c,d-1} \cdot \frac{\kappa \hat{p}_{s,1}}{D}, \tag{6.94}$$

where

$$\beta = \begin{cases} 0 & \text{if } (\Gamma,\ i,\ b+1,\ c,\ d) \text{ is feasible, i.e., satisfies (6.82) and (6.84)} \\ \hat{p}_{l,1} & \text{otherwise.} \end{cases} \tag{6.95}$$

$$\gamma = \begin{cases} 0 & \text{if } (\Gamma,\ i,\ b,\ c+1,\ d) \text{ is feasible, i.e., satisfies (6.82) and (6.92)} \\ \hat{p}_{s,a} & \text{otherwise.} \end{cases} \tag{6.96}$$

$$\delta = \begin{cases} 0 & \text{if } (\Gamma,\ i,\ b,\ c,\ d+1) \text{ is feasible, i.e., satisfies (6.82) and} \\ & \text{either (6.84) or (6.92)} \\ \hat{p}_{s,1} & \text{otherwise.} \end{cases} \tag{6.97}$$

We initialize this recursion with ${}^{\Gamma}Q^{0}_{b,c,d} = 1$ if $b = c = d = 0$, ${}^{\Gamma}Q^{0}_{b,c,d} = 0$ otherwise, and note that all undefined ${}^{\Gamma}Q^{i}_{b,c,d}$ (e.g., those with negative b, c, or d) are set to zero.

Markov Chain Model for Number of Occupied Columns C In this section we derive the steady state probabilities $\pi_l = P(C = l),\ l = 0, \ldots, (D-1)\xi$, that l columns in the scheduling window of the considered port o, $o = 1, \ldots, D$, are already occupied by the other ports o', $o' = 1, \ldots, D$, $o' \neq o$, when port o begins its data packet scheduling. Towards this end we construct an irreducible, positive recurrent Markov chain with the states $C = 0,\ C = 1, \ldots, C = (D-1)\xi$. The Markov chain makes state transitions in every frame. Specifically, we interpret C_n as the number of columns in the scheduling window of port o that are already occupied when the scheduling of the data packets from port o commences. After the data packets from port o have been scheduled, the Markov chain makes a state transition. We interpret C_{n+1} as the number of occupied columns in the scheduling window of port $(o+1)$, that is, upon the state transition the considered scheduling window moves one frame into the future. (If port $o = D$ was originally considered, then C_{n+1} is the number of occupied columns in the scheduling window of port $o = 1$.)

Let Z be a random variable denoting the number of columns in the short transmission slots of the scheduling window of port o that are occupied by $(s,1)$–packets from port o when the scheduling of the data packets from port o is completed. When counting the number of columns occupied by the data packets from port o we ignore whether these columns have already been occupied by some other port or not. The number of columns in the long transmission slot of port o that are occupied by the packets from port o are not included in Z since the scheduling window advances by one frame when port o is done with the scheduling. Thus, the first frame of port o's scheduling window is no longer included in port $(o+1)$'s scheduling window. With Z the state transition probabilities r_{ij} of the Markov chain C_n are given by

$$r_{ij} = P(C_{n+1} = j | C_n = i) = P(\max\{Z,\ i-\xi\} = j | C_n = i) \tag{6.98}$$

$$= \begin{cases} P(Z = j | C_n = i) & \text{if } j > i-\xi \\ P(Z \leq j | C_n = i) & \text{if } j = i-\xi \\ 0 & \text{if } j < i-\xi. \end{cases} \tag{6.99}$$

To see this note that as we make the state transition from the scheduling window of port o to the scheduling window of port $(o+1)$, the considered scheduling window advances one frame into the future and the first ξ columns of the scheduling window of port o are no longer considered. Also, note that the data packets are scheduled in a first–come–first–served and first–fit manner. Hence, the Z first columns in the advanced scheduling window are occupied by packets from port o and the $\max\{i-\xi,\ 0\}$ first columns are occupied by packets scheduled prior to the data packet scheduling from port o. Thus, for $j > i-\xi$ columns to be occupied in the advanced scheduling

window, the data packets from port o must occupy $Z = j$ columns. For $j = i - \xi$ occupied columns in the advanced scheduling window, the packets from port o may occupy $Z = 0, \ldots, j$ columns.

Next, we calculate the probabilities $P(Z = j|C_n = i)$ for $i,\ j = 0, \ldots, (D-1)\xi$. First, we consider these probabilities for $i = 0$. We have for $j \geq 1$

$$\begin{aligned} P(Z = j|C_n = 0) &= \sum_{d=R(j-1)+1}^{Rj} P^M_{1,0,0,d} + \sum_{b=1}^{R} \sum_{d=(R-b)\lfloor F/K \rfloor + R(j-1)+1}^{(R-b)\lfloor F/K \rfloor + Rj} P^M_{0,b,0,d} + \\ &\quad \sum_{c=0}^{\lfloor F/K \rfloor} \sum_{d=(\lfloor F/K \rfloor - c)R + R(j-1)+1}^{(\lfloor F/K \rfloor - c)R + Rj} P^M_{0,0,c,d}. \end{aligned} \tag{6.100}$$

For $j = 0$ we have

$$P(Z = 0|C_n = 0) = 1 - \sum_{j=1}^{(D-1)\xi} P(Z = j|C_n = 0). \tag{6.101}$$

Next we consider the probabilities with $1 \leq i \leq \xi$. We have for $j \geq 1$

$$\begin{aligned} P(Z = j|C_n = i) &= \sum_{b=1}^{R} \sum_{d=(R-b)\lfloor F/K \rfloor + R(j-1)+1}^{(R-b)\lfloor F/K \rfloor + Rj} {}^{i}Q^M_{b,0,d} + \\ &\quad \sum_{c=0}^{\lfloor F/K \rfloor} \sum_{d=(\lfloor F/K \rfloor - c)R + R(j-1)+1}^{(\lfloor F/K \rfloor - c)R + Rj} {}^{i}Q^M_{0,c,d}. \end{aligned} \tag{6.102}$$

Furthermore, we have for $i > \xi$ and $j \geq 1$

$$P(Z = j|C_n = i) = P(Z = j|C_n = \xi), \tag{6.103}$$

and finally for $1 \leq i \leq (D-1)\xi$ and $j = 0$

$$P(Z = 0|C_n = i) = 1 - \sum_{j=1}^{(D-1)\xi} P(Z = j|C_n = i). \tag{6.104}$$

With the calculated state transition probabilities $r_{ij},\ i,\ j = 0, \ldots, (D-1)\xi$, we find the steady state probabilities $\pi_l,\ l = 0, \ldots, (D-1)\xi$, as the solution to

$$\pi_j = \sum_i \pi_i \cdot r_{ij} \tag{6.105}$$

$$\sum_j \pi_j = 1. \tag{6.106}$$

Expected numbers of scheduled packets We obtain the expected number of scheduled packets as

$$E[X_{l,a}] = \pi_0 \cdot \sum_d P^M_{1,0,0,d}, \tag{6.107}$$

$$E[X_{l,1}] = \pi_0 \cdot \sum_{b,d} b \cdot P^M_{0,b,0,d} + \sum_{l=1}^{\xi-1} \pi_l \cdot \sum_{b,d} b \cdot {}^lQ^M_{b,0,d} + \left(\sum_{l \geq \xi} \pi_l \right) \cdot \sum_{b,d} b \cdot {}^\xi Q^M_{b,0,d}, \quad (6.108)$$

$$E[X_{s,a}] = \pi_0 \cdot \sum_{c,d} c \cdot P^M_{0,0,c,d} + \sum_{l=1}^{\xi-1} \pi_l \cdot \sum_{c,d} c \cdot {}^lQ^M_{0,c,d} + \left(\sum_{l \geq \xi} \pi_l \right) \cdot \sum_{c,d} c \cdot {}^\xi Q^M_{0,c,d}, \quad (6.109)$$

$$E[X_{s,1}] = \pi_0 \cdot \sum_{a,b,c,d} d \cdot P^M_{a,b,c,d} + \sum_{l=1}^{\xi-1} \pi_l \cdot \sum_{b,c,d} d \cdot {}^lQ^M_{b,c,d} + \left(\sum_{l \geq \xi} \pi_l \right) \cdot \sum_{b,c,d} d \cdot {}^\xi Q^M_{b,c,d}, \quad (6.110)$$

where the summations are over all feasible scheduling patterns as given by the respective schedulability conditions (6.82), (6.83), (6.84), (6.85), and (6.92).

Equilibrium conditions In this section we establish equilibrium conditions for the network. For ease of notation let

$$E[X] = E[X_{l,a}] + E[X_{l,1}] + E[X_{s,a}] + E[X_{s,1}]. \quad (6.111)$$

Also, note that for large S we may reasonably approximate the expected value $E[\eta]$ by η. This is reasonable since η has only small fluctuations in steady state for large S. With this approximation the first equilibrium condition is

$$\frac{\sigma}{D} \cdot \eta = E[X]. \quad (6.112)$$

This is because $\sigma \cdot \eta$ new packets are generated in each cycle by the nodes attached to a given AWG input port. With probability $1/D$ each of the generated packets is destined to a given (fixed) AWG output port (splitter). On the other hand, $E[X]$ packets are scheduled (and transmitted) on average from a given AWG input port to a given AWG output port in one cycle; in equilibrium as many new packets must be generated.

The four other equilibrium conditions are

$$E[X_{l,a}] = p_{l,a} \cdot E[X], \quad (6.113)$$

$$E[X_{l,1}] = p_{l,1} \cdot E[X], \quad (6.114)$$

$$E[X_{s,a}] = p_{s,a} \cdot E[X], \quad (6.115)$$

$$E[X_{s,1}] = p_{s,1} \cdot E[X]. \quad (6.116)$$

These hold because in equilibrium the mean number of scheduled packets of a given type from a given AWG input port to a given AWG output port in one cycle (LHS in the equations) is equal to the number of newly generated packets of this type in one cycle (RHS in the equations).

The first equilibrium condition (6.112) and any three of the four conditions (6.113), (6.114), (6.115), and (6.116), along with $\hat{p}_{l,a} + \hat{p}_{l,1} + \hat{p}_{s,a} + \hat{p}_{s,1} = 1$ give a system of five linear independent equations which can be solved by standard numerical techniques for the five unknowns η, $\hat{p}_{l,a}$, $\hat{p}_{l,1}$, $\hat{p}_{s,a}$, and $\hat{p}_{s,1}$. These are then used to calculate the expected numbers of scheduled packets from a given AWG input port to a given AWG output port per cycle $E[X_{l,a}]$, $E[X_{l,1}]$, $E[X_{s,a}]$, and $E[X_{s,1}]$ using the recursive approach given in the preceding section.

Performance measures Based on the expected numbers of scheduled packets we evaluate the network performance metrics as follows. The mean aggregate transmitter throughput TH_T is defined as the mean number of transmitting nodes in the network in steady state and is given by

$$TH_T = D^2 \cdot \frac{F \cdot (E[X_{l,a}] + E[X_{l,1}]) + K \cdot (E[X_{s,a}] + E[X_{s,1}])}{F \cdot D}. \tag{6.117}$$

The mean receiver throughput TH_R is defined as the average number of receiving nodes in the network in steady state and is given by

$$TH_R = D^2 \cdot \frac{F \cdot S \cdot E[X_{l,a}] + F \cdot E[X_{l,1}] + K \cdot S \cdot E[X_{s,a}] + K \cdot E[X_{s,1}]}{F \cdot D}. \tag{6.118}$$

The mean delay in the network is defined as the average time in cycles from the generation of the control packet corresponding to a data packet until the transmission of the data packet. Following the arguments in Section 6.2.2 we obtain

$$\text{Delay} = \frac{S}{D \cdot E[X]} - \frac{1-\sigma}{\sigma}. \tag{6.119}$$

Results

In this section, we conduct numerical investigations of the interaction between unicast and multicast traffic. This investigation quantifies the benefits of multicasting concurrently with reservation control traffic in conjunction with unicast with spatial wavelength reuse. The default network parameters are set as follows: Number of nodes $N = 200$, number of available wavelengths at each AWG port $D \cdot R = 8$, cycle length $D \cdot F = 800$ slots, number of reservation slots per frame $M = 40$, retransmission probability $p = 0.8$. We have also conducted extensive simulations of a more realistic network in order to verify the accuracy of the analytical model. As opposed to the analysis, in the simulation a given node cannot transmit unicast packets to itself. Furthermore, in the simulation not only the packet type (length, unicast or multicast) but also the destination of a given unicast or multicast packet are not renewed, i.e., are persistent, when retransmitting the corresponding control packet (recall that the analysis assumes that the type of the packet is persistent while the destination is nonpersistent). In addition, the simulation takes all receiver conflicts into account, i.e., a given unicast or multicast packet is not scheduled if the receiver(s) of the intended destination(s) is (are) busy. Each simulation was run for 10^7 slots including a warm-up phase of 10^6 slots. Using the method of batch means we calculated the 98% confidence intervals for the performance metrics, which were always smaller than 4% of the corresponding sample means.

Fig. 6.30 depicts the mean aggregate transmitter throughput (mean number of transmitting nodes) and receiver throughput (mean number of receiving nodes) in steady state for different fraction $p_{l,a} + p_{s,a} = \{0\%, 10\%, 30\%, 50\%\}$ of multicast packets. In all cases, the fraction of short data packets is $p_{s,a} + p_{s,1} = 0.75$. Accordingly, the fraction of long data packets is $p_{l,a} + p_{l,1} = 0.25$. The AWG degree is set to $D = 4$. Hence, the number of used FSRs is $R = 2$, the frame size equals $F = 200$ slots, and short packets are $K = 160$ slots long. If the fraction of multicast packets is equal to 0% all packets are unicast and transmitter throughput is identical to receiver throughput. As shown in Fig. 6.30, increasing the fraction of multicast packets from 0% up to 50% results in a dramatically larger receiver throughput and a slightly smaller transmitter throughput. This is due to the fact that with an increasing fraction of multicast

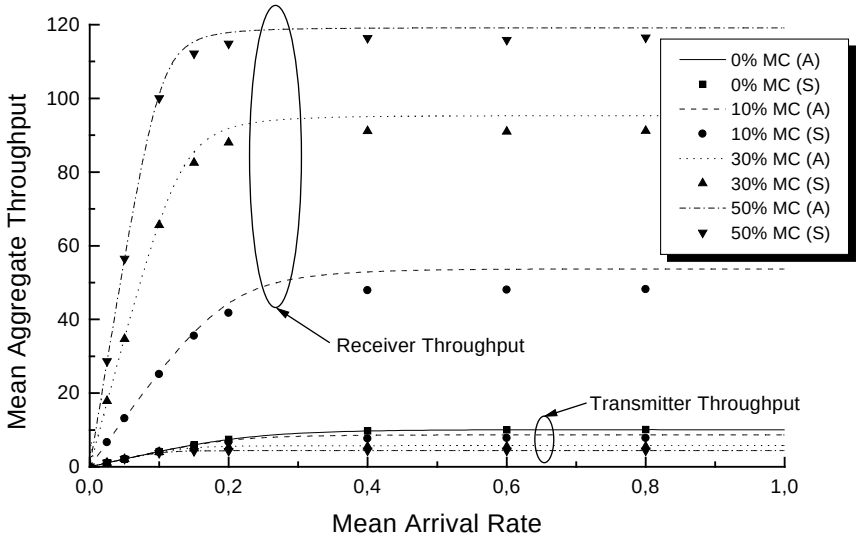

Figure 6.30: Mean aggregate transmitter and receiver throughput vs. mean arrival rate σ (packet/cycle) for different fractions $p_{l,a} + p_{s,a} = \{0\%, 10\%, 30\%, 50\%\}$ of multicast packets (fraction of short packets $p_{s,a} + p_{s,1} = 0.75$, fixed).

packets more receivers are used, resulting in a larger receiver throughput. On the other hand, the transmitter throughput is slightly decreased since nodes are less likely to find free receivers, leading to a smaller number of transmissions and thereby smaller transmitter throughput. Note that analysis and simulation results match very well at low traffic loads. At medium to high loads, on the other hand, the analysis provides a slightly larger receiver throughput than the simulation. This is due to the assumed nonpersistency of the destination(s) in the analysis. As opposed to the simulation, in the analysis unsuccessful control packets renew the destination of the corresponding multicast packets. Consequently, in the analysis previously conflicting multicast packets are less likely to collide again and can be successfully scheduled resulting in an increased receiver throughput. Overall the results clearly illustrate that scheduling multicast packets concurrently with reservation control in each frame significantly improves the receiver utilization. For instance, if 50% of the packets are multicast the mean receiver utilization is almost 60%, as shown in Fig. 6.30.

The mean delay (in cycles) vs. mean arrival rate is shown in Fig. 6.31 for different fractions $p_{l,a} + p_{s,a} = \{0\%, 10\%, 30\%, 50\%\}$ of multicast packets. As expected, with increasing arrival rate the mean delay grows due to more channel and receiver collisions. Moreover, with an increasing fraction of multicast traffic the mean delay becomes larger. Again, this is because with increasing multicast traffic the receiver utilization is higher, resulting in more unsuccessful reservation requests and retransmissions. Note that the analysis yields smaller delay values than the simulation. This is because of two reasons. First, due to the destination nonpersistency in the analysis, control packets are more likely to be successful and have to be retransmitted fewer times resulting in a smaller delay. Second, the definitions of packet delay are slightly different for simulation and analysis. In the simulation the packet delay is defined as the time interval between packet generation and end of packet transmission. In the analysis the packet delay

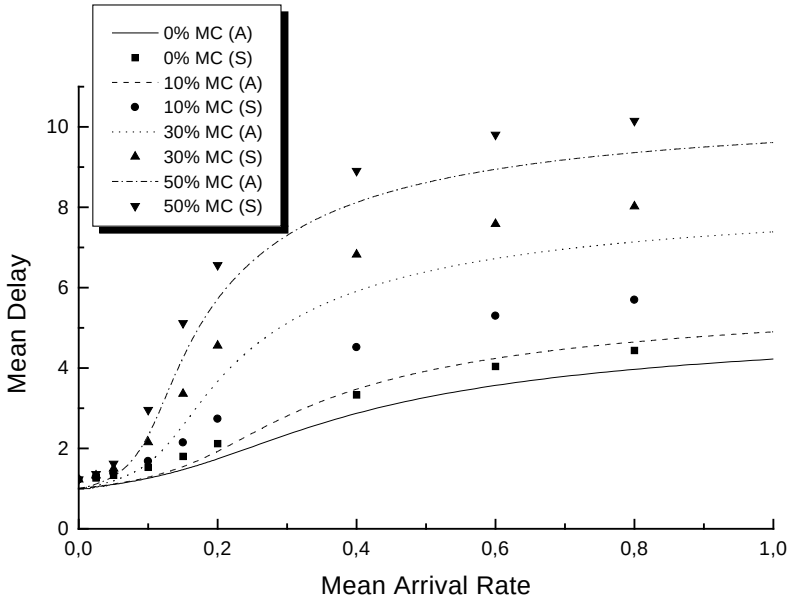

Figure 6.31: Mean delay (cycles) vs. mean arrival rate σ (packet/cycle) for different fractions $p_{l,a} + p_{s,a} = \{0\%, 10\%, 30\%, 50\%\}$ of multicast packets (fraction of short packets $p_{s,a} + p_{s,1} = 0.75$, fixed).

is defined as the time interval between packet generation and the time when the packet is successfully scheduled but not yet transmitted.

In Figs. 6.32 through 6.34 we set the AWG degree to $D = 2$ and the fraction of long data packets to $p_{l,a} + p_{l,1} = 0.25$. Long packets are $L = F = 400$ slots and short packets are $K = 120$ slots long. The number of reservation slots per frame is $M = 40$, $\lfloor F/K \rfloor = 3$, and $R = 4$ FSRs of the underlying AWG are used. 80% of the data packets are unicast, i.e., $p_{l,1} + p_{s,1} = 0.8$. Accordingly, 20% of the data packets are multicast, i.e., $p_{l,a} + p_{s,a} = 0.2$. The multicast packets can be either only short, both short and long, or only long. Specifically, we consider different ratios $p_{s,a}/(p_{s,a} + p_{l,a}) = 0\%,\ 50\%$, and 100% of short multicast packets.

Fig. 6.32 shows how the mean aggregate transmitter throughput is increased by varying the ratio of short and long multicast packets. If 0% of the multicast packets are short, i.e., all multicast packets are long, the transmitter throughput is rather small. By increasing the number of short multicast packets from 0% up to 100% the transmitter throughput is significantly increased; for 100% of short multicast packets the mean aggregate throughput is roughly doubled compared to 0% multicast packets. This is due to the fact that without long multicast packets nodes ready to send short multicast packets are more likely to find free receivers which translates into an increased transmitter throughput. However, increasing the number of short multicast packets leads to a decreased receiver throughput, as depicted in Fig. 6.33. Thus, there is a tradeoff between channel and receiver utilization. Again, analysis and simulation results match very well at low traffic loads. However, at medium to high loads the analysis and simulation results exhibit some discrepancy. While we observe that the discrepancy is not that large for the case of 100% short multicast packets, the mismatch is more pronounced if the amount of long multicast packets is increased. This is again due to the destination nonpersistency assumption made in the analysis which resolves the destination conflicts as opposed to the

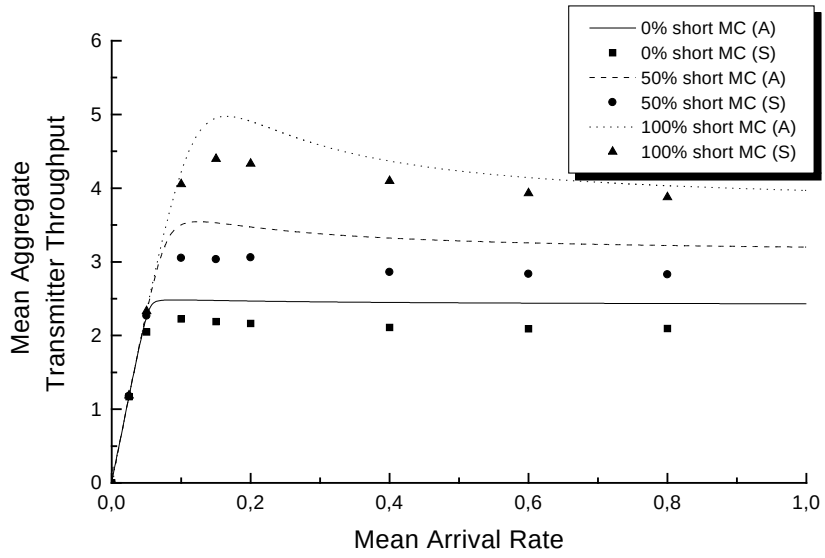

Figure 6.32: Mean aggregate transmitter throughput vs. mean arrival rate σ (packet/cycle) for different ratios $\{0\%, 50\%, 100\%\}$ of short multicast packets (20% multicast traffic, fixed).

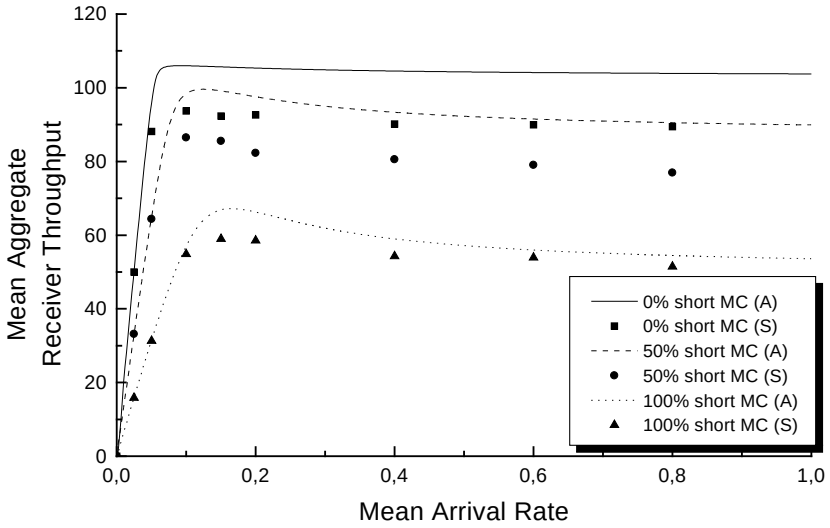

Figure 6.33: Mean aggregate receiver throughput vs. mean arrival rate σ (packet/cycle) for different ratios $\{0\%, 50\%, 100\%\}$ of short multicast packets (20% multicast traffic, fixed).

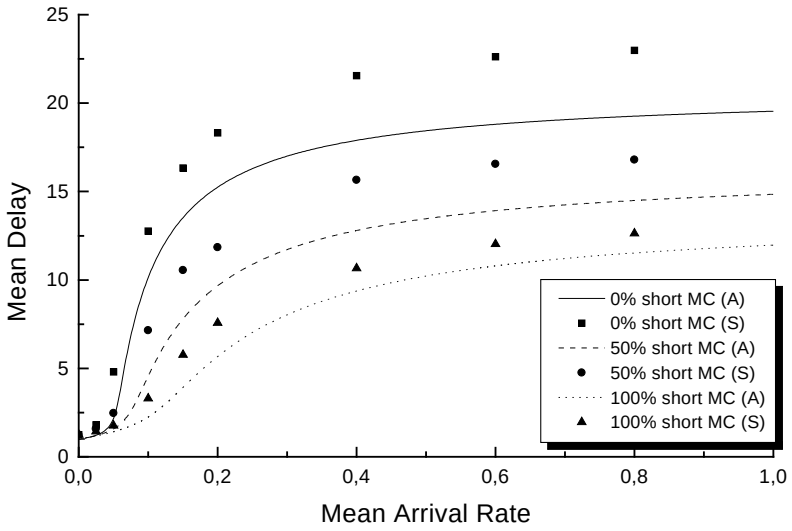

Figure 6.34: Mean delay (cycles) vs. mean arrival rate σ (packet/cycle) for different ratios $\{0\%, 50\%, 100\%\}$ of short multicast packets (20% multicast traffic, fixed).

simulation resulting in a larger mean aggregate receiver throughput.

Fig. 6.34 depicts the mean delay vs. mean arrival rate σ for different ratios of short and long multicast packets $\{0\%, 50\%, 100\%\}$. We observe that with an increasing number of short multicast packets the mean delay is decreased. This is because in the presence of fewer long multicast packets, receivers are more likely to be free. As a consequence, more data packets are scheduled resulting in fewer retransmissions of control packets and decreased delay.

6.4 Supplementary simulation results

In this section, we investigate by means of extensive simulations additional aspects of our AWG based network which have not been analyzed yet in the previous three sections. Prior to the following investigations we have validated the correctness of our simulator. We have run simulations making the same assumptions as in Section 6.2.1. In extensive tests we have compared the obtained results with the analytical results presented in Section 6.2.3. Throughout our tests we have achieved very good matches between simulation and analysis. In the following simulations we not only examine additional performance metrics but, more importantly, also relax some of the simplifying assumptions made in the previous analyses in order to make our investigations more realistic. Specifically, in Section 6.4.1 we consider the self–stability of our MAC protocol. We thereby drop the assumption of a constant retransmission probability p of unsuccessful control packets. Instead, p is reduced each time a given reservation request fails. Section 6.4.2 investigates several approaches to reduce packet loss. Among other approaches, packet loss can be decreased by extending the scheduling window size and/or equipping each node with a finite buffer. In doing so, we relax the assumptions of a finite window size and a single–packet buffer made in our analytical models. Finally, in Section 6.4.3 we take circuit

switching into consideration beside packet switching. Section 6.4.4 discusses the efficiency of our network and also conducts a benchmark comparison with a previously reported reservation protocol designed for a PSC based single–hop metro WDM network. In the following simulative studies we concentrate on unicast traffic.

In the simulations in this section, the network parameters take on the following values by default: Number of nodes $N = 200$, physical degree of AWG $D = 2$, number of used FSRs $R = 4$, number of slots per frame $F = 200$, number of reservation slots per frame $M = 30$, and retransmission probability $p = 1.0$. The fraction of long data packets q equals 0.25, i.e., a generated data packet is long ($L = F = 200$ slots) with probability $q = 0.25$ and short ($L = K = (F - M) = 170$ slots) with probability $(1 - q) = 0.75$. By default the size of the scheduling window is one cycle and each node's single–packet buffer is able to store either a long or a short data packet. Each cycle is assumed to have a constant length of $D \cdot F = 400$ slots. The propagation delay is assumed to be less than one cycle. A node that has made a successful reservation can send the corresponding data packet in the next cycle. In the simulation the mean arrival rate denotes the probability that a given node generates a data packet at the beginning of that frame in which the node is allowed to send control packets. In the simulation nodes do not transmit data packets to themselves. A given data packet is destined to any of the other $(N - 1)$ nodes with equal probability $1/(N - 1)$. Furthermore, in the simulation the length and destination of a given data packet are not changed if the corresponding control packet has to be retransmitted. As in Section 6.2, the mean aggregate throughput is given in packets per frame. The mean delay is equal to the time interval between the generation and the end of transmission of a given data packet and is given in cycles. Each simulation was run for $6 \cdot 10^5$ slots including a warm–up phase of $6 \cdot 10^4$ slots. The remaining $5.4 \cdot 10^5$ slots were divided into 90 batches, each comprising $6 \cdot 10^3$ slots. Using the method of batch means we calculated the 95% confidence intervals for the mean aggregate throughput, mean delay, and relative packet loss.

6.4.1 Self–stability

So far, we have not addressed *backoff* in our MAC protocol. Without backoff the retransmission probability p remains constant irrespective of how many times a given control packet has already been retransmitted. As the mean arrival rate increases more nodes are backlogged and try to make a reservation in their assigned frame. The network becomes increasingly congested and more nodes have to retransmit their unsuccessful control packets (with constant probability $p = 1$ in our case). This leads to an increased number of control packet collisions on the slotted ALOHA control channel and retransmissions. As a result, the mean aggregate throughput decreases while the mean delay increases dramatically with an increasing mean arrival rate, as depicted in Figs. 6.35 and 6.36, respectively.

By deploying backoff this instability can be alleviated. With backoff the retransmission probability p is reduced each time the reservation fails. More precisely, a given control packet which for the first time fails in making a successful reservation is retransmitted with probability $p = 1$ in the next cycle. If the reservation fails again p is reduced by 50%. Thus, the corresponding control packet is retransmitted in the next cycle with probability $p = 0.5$. With probability $(1 - p) = 0.5$ the reservation is deferred by one cycle. The control packet is then retransmitted with probability $p = 0.5$ in that next cycle and deferred with probability $(1 - p) = 0.5$ by one more cycle, and so on. Each time the reservation fails, p is further halved. In general, p is reduced by 50% at most b times until p is equal to a given minimum retransmission probability p_{min}, where $b \geq 1$ denotes the backoff limit. Once p_{min} is reached the corresponding control

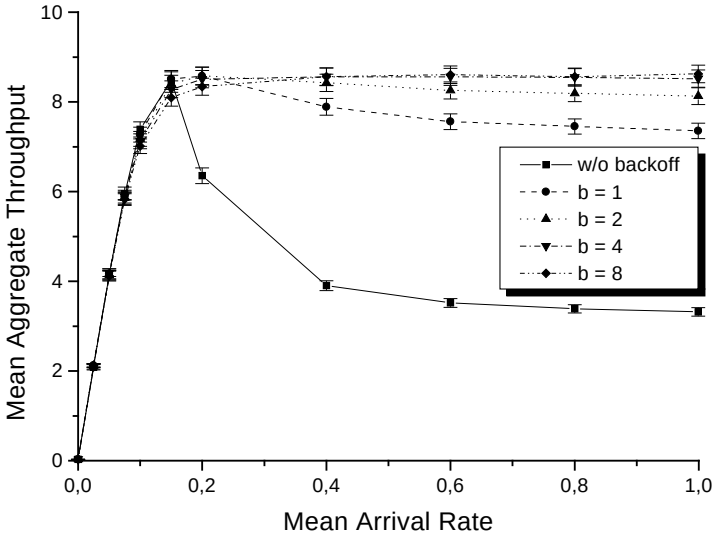

Figure 6.35: Mean aggregate throughput (packets/frame) vs. mean arrival rate for different backoff limits $b \in \{1, 2, 4, 8\}$.

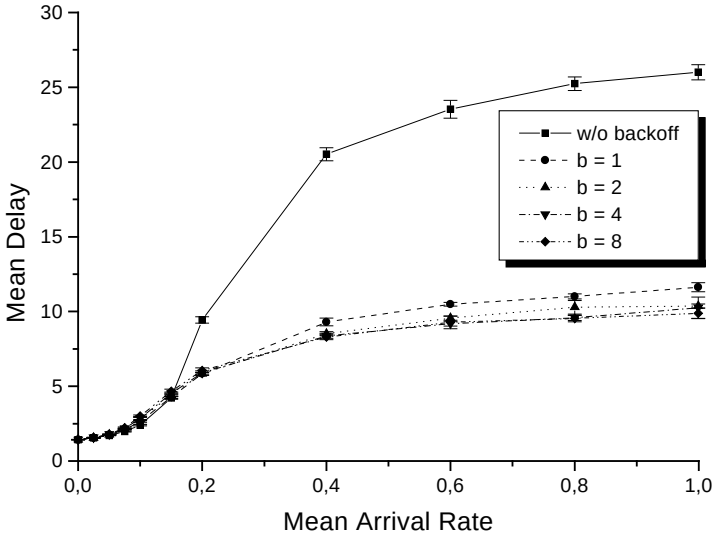

Figure 6.36: Mean delay (cycles) vs. mean arrival rate for different backoff limits $b \in \{1, 2, 4, 8\}$.

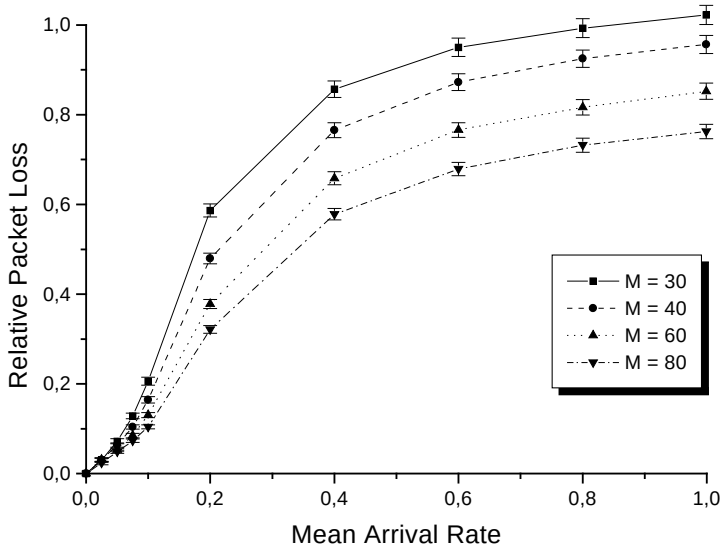

Figure 6.37: Relative packet loss vs. mean arrival rate for $b = 4$ and different $M \in \{30, 40, 60, 80\}$.

packet is retransmitted in a cycle with probability p_{min} until the reservation is successful, i.e., there is no attempt limit.

Figs. 6.35 and 6.36 show the positive impact of backoff on the throughput–delay performance of the network for $b \in \{1, 2, 4, 8\}$. We observe that already with $b = 1$, i.e., p is halved not more than once, the mean aggregate throughput is significantly increased for medium to high traffic loads while decreasing the mean delay considerably. We observe in Fig. 6.35 that with $b = 4$ the mean aggregate throughput does not decrease for an increasing mean arrival rate. This property of MAC protocols is known as *self–stability*. Further increasing b does not yield a better throughput–delay performance. Therefore, we set $b = 4$ in the subsequent simulations.

6.4.2 Packet loss

Since so far we assume that each node is equipped with a single–packet buffer new generated data packets are dropped when they find the buffer full resulting in packet loss. Fig. 6.37 depicts the relative packet loss (ratio of number of dropped data packets and number of generated data packets) as a function of the mean arrival rate for $b = 4$ and different number of reservation slots per frame $M \in \{30, 40, 60, 80\}$. Obviously, the relative packet loss increases monotonously for an increasing mean arrival rate. For increasing M the relative packet loss is decreased. Hence, choosing M as large as possible appears to be the best solution to reduce the packet loss. However, from the throughput perspective it is more advantageous to choose M neither too small nor too large. This can be seen in Fig. 6.38 where $M = 60$ achieves the largest mean aggregate throughput for a wide range of mean arrival rate. While small values of M result in more control packet collisions on slotted ALOHA and thereby a decreased mean aggregate throughput, $M = 80$ implies that more control packets are sent collisionfree but the length of the corresponding short data packets is reduced since the length of short data packets is given

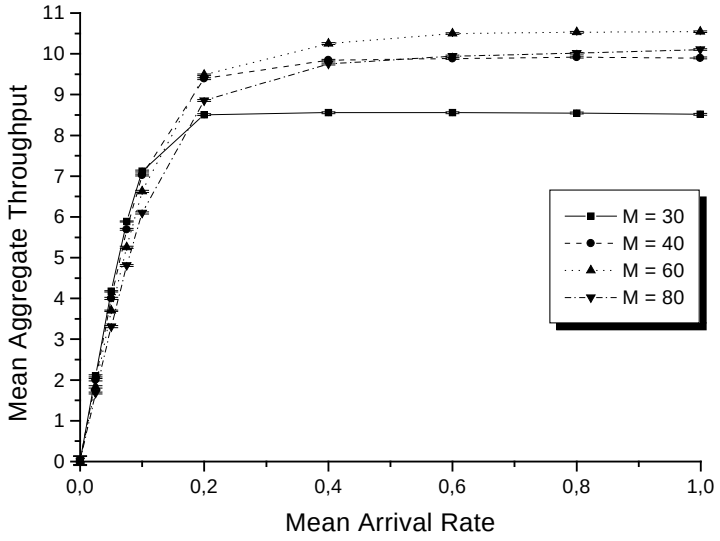

Figure 6.38: Mean aggregate throughput (packets/frame) vs. mean arrival rate for $b = 4$ and different $M \in \{30, 40, 60, 80\}$.

by $K = (F - M) = (200 - M)$ slots. Fig. 6.39 illustrates the mean delay vs. mean arrival rate. Since the mean delay for $M = 60$ and $M = 80$ do not differ significantly and $M = 60$ achieves the maximum mean aggregate throughput we set $M = 60$ in the subsequent simulations.

After considering the slotted ALOHA bottleneck we now turn our attention to two other bottlenecks of the proposed network: (*i*) Number of wavelength channels, and (*ii*) scheduling window size. Fig. 6.40 shows the relative packet loss vs. mean arrival rate for $b = 4$, $M = 60$, and different number of used FSRs $R \in \{2, 4, 6, 8, 16\}$. We observe that the packet loss is decreased significantly for increasing R. Figs. 6.41 and 6.42 show that a large R also improves the throughput–delay performance of the network. This is due to the fact that by increasing R, i.e., by using more FSRs of the underlying AWG, the degree of concurrency is increased. As a consequence, data packets are simultaneously transmitted over a larger number of wavelengths leading to a decreased packet loss, increased throughput, and decreased delay. Note that for a given transceiver tuning range $\Lambda = D \cdot R$ and a given AWG degree D increasing R implies a smaller channel spacing. Generally, AWGs with a smaller channel spacing exhibit a larger crosstalk. In order to achieve acceptable crosstalk values we set $R = 8$ in the subsequent simulations. For $D = 2$ and a typical fast transceiver tuning range of 12 nm setting $R = 8$ translates into a channel spacing of 100 GHz.

Fig. 6.43 depicts the packet loss as a function of the mean arrival rate for different scheduling window size $W \in \{2, 4, 6, 8\}$, where W is given in frames. We observe that the packet loss is decreased when the scheduling window is enlarged from two frames (i.e., one cycle for $D = 2$) to four frames (i.e., two cycles for $D = 2$). Further increasing W has no impact on the packet loss. The same observation can be made for the mean aggregate throughput and the mean delay, as shown in Figs. 6.44 and 6.45, respectively. Using a scheduling window of two cycles increases the throughput and decreases the delay. However, further increasing W does not affect

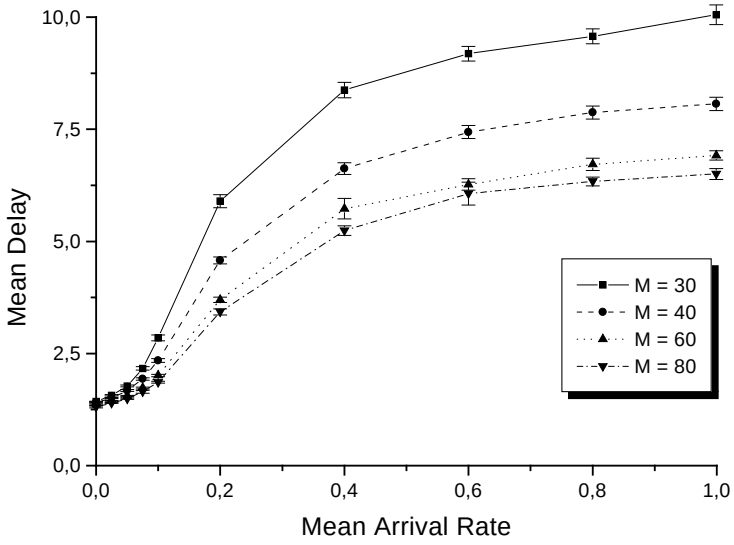

Figure 6.39: Mean delay (cycles) vs. mean arrival rate for $b = 4$ and different $M \in \{30, 40, 60, 80\}$.

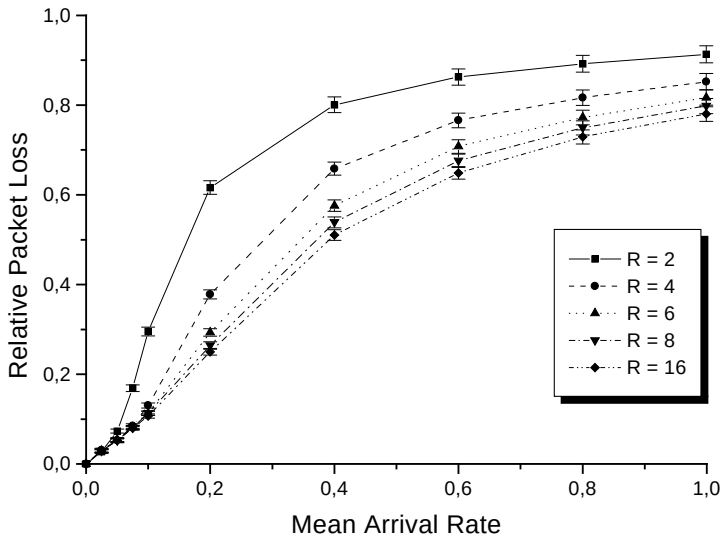

Figure 6.40: Relative packet loss vs. mean arrival rate for $b = 4$, $M = 60$, and different $R \in \{2, 4, 6, 8, 16\}$.

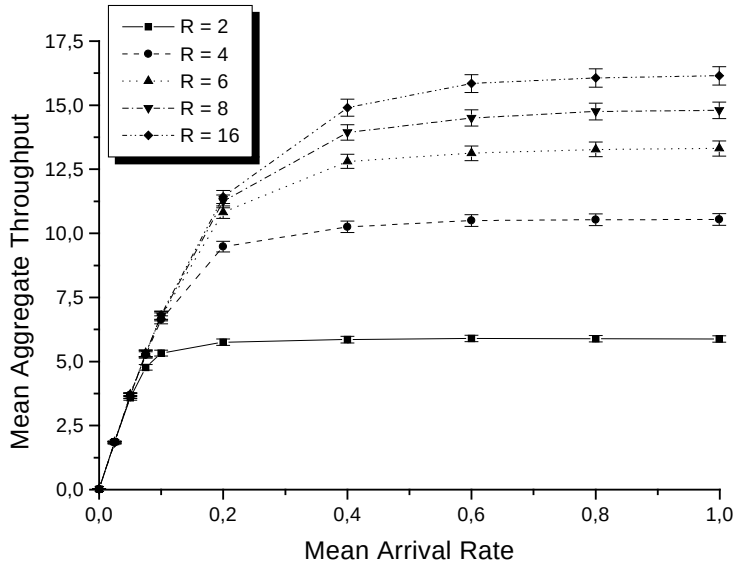

Figure 6.41: Mean aggregate throughput (packets/frame) vs. mean arrival rate for $b = 4$, $M = 60$, and different $R \in \{2, 4, 6, 8, 16\}$.

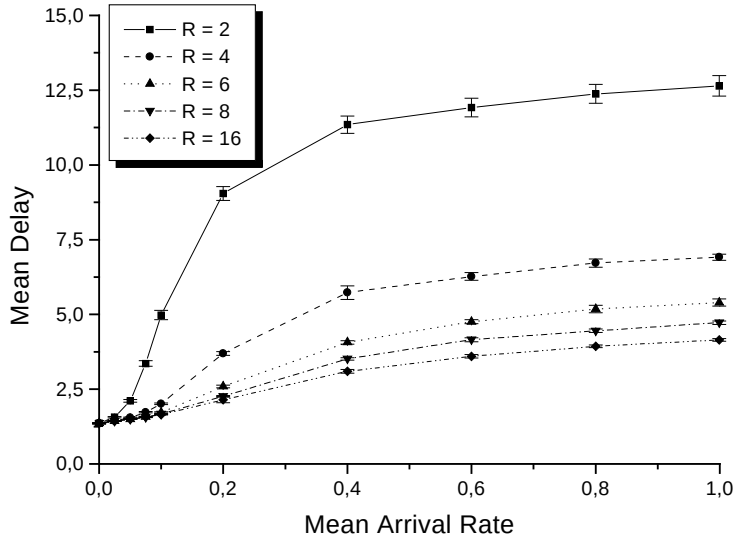

Figure 6.42: Mean delay (cycles) vs. mean arrival rate for $b = 4$, $M = 60$, and different $R \in \{2, 4, 6, 8, 16\}$.

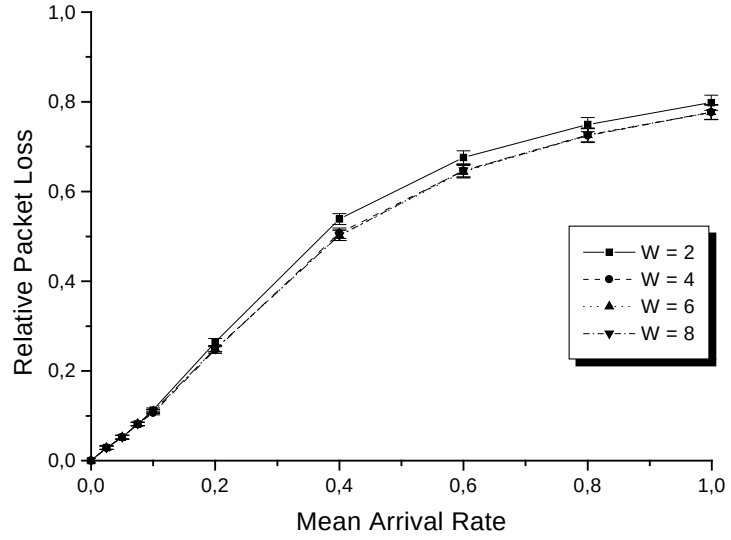

Figure 6.43: Relative packet loss vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, and different $W \in \{2, 4, 6, 8\}$.

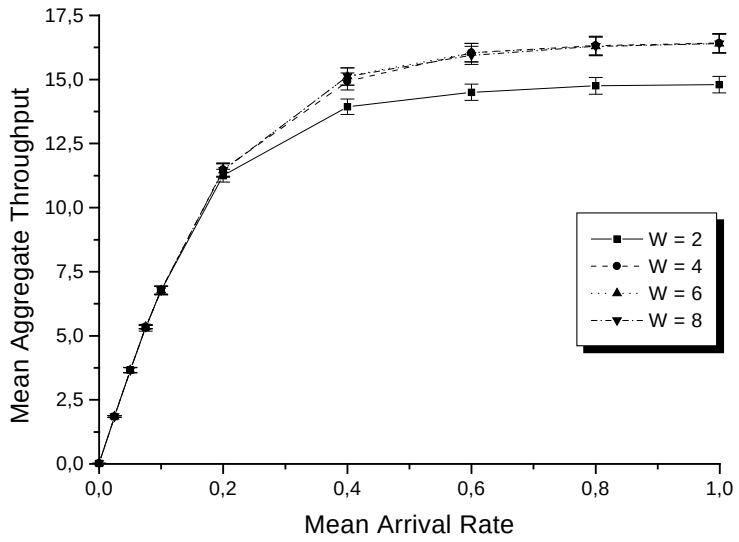

Figure 6.44: Mean aggregate throughput (packets/frame) vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, and different $W \in \{2, 4, 6, 8\}$.

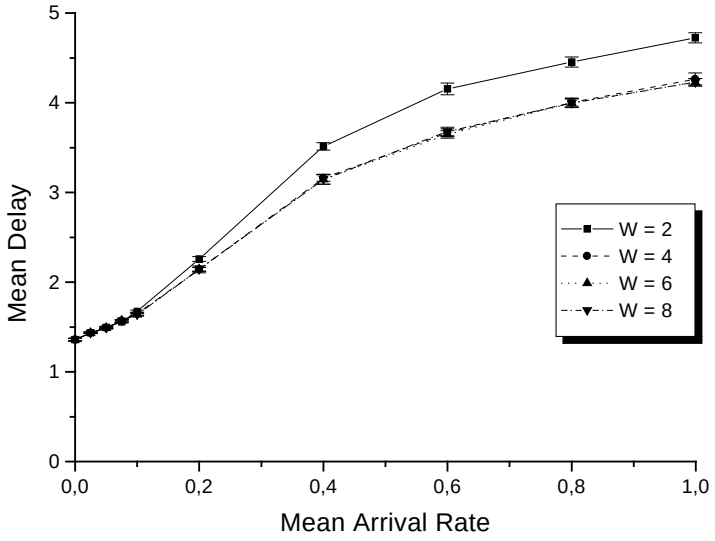

Figure 6.45: Mean delay (cycles) vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, and different $W \in \{2, 4, 6, 8\}$.

the throughput–delay performance of the network. This is due to the fact that for $W \geq 4$ the network resources are almost fully utilized. As a consequence, no additional data packets can be transmitted for increasing W which in turn leads to a stagnating packet loss, throughput, and delay.

Note that all curves are rather close to each other. The packet loss, throughput, and delay do not change very much for increasing W. Therefore, the assumption of a scheduling window size equal to one cycle as made in the analyses so far is reasonable. For the subsequent simulations we set $W = 8$. This is because we will consider also the case $D = 4$ below; in order to benefit from a scheduling window of two cycles the parameter W must be equal to eight for $D = 4$.

Another approach to reduce the packet loss is to increase the buffer at each node. The single–packet buffer is replaced with a buffer that is able to store up to B long data packets, where $B \geq 1$. (The first–in–first–out (FIFO) queueing discipline is used, which is reasonable for the considered very high–speed network). Fig. 6.46 illustrates the positive impact of larger buffer sizes on the packet loss. With increasing B more arriving data packets can be stored and do not have to be dropped resulting in a decreased packet loss. In addition, each node is less likely to be idle which leads to an increased mean aggregate throughput, as depicted in Fig. 6.47. However, Fig. 6.48 shows that the mean delay significantly increases for larger B. This is because in larger buffers there are packets which have to wait for a longer time interval until they can be transmitted. Clearly, there is a trade–off between packet loss and delay. To avoid large delays and provide a reasonable throughput–loss performance we set $B = 10$ in the following simulations. Fig. 6.49 shows that the packet loss can be further reduced if the $N = 200$ nodes are connected by a 4×4 AWG instead of a 2×2 one. Note that changing the AWG degree D implies also a different R, F, and B. For a given transceiver tuning range Λ, a larger D translates into a smaller R since $R = \Lambda/D$. Moreover, we have assumed a fixed cycle length of

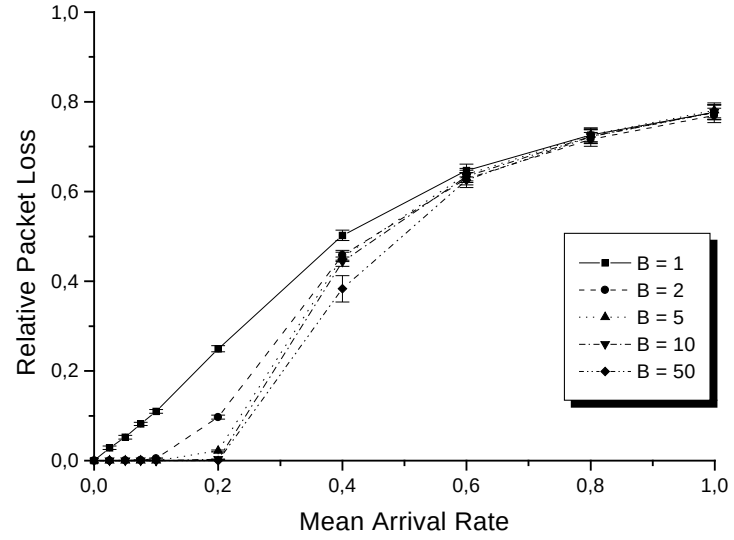

Figure 6.46: Relative packet loss vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, $W = 8$, and different $B \in \{1, 2, 5, 10, 50\}$.

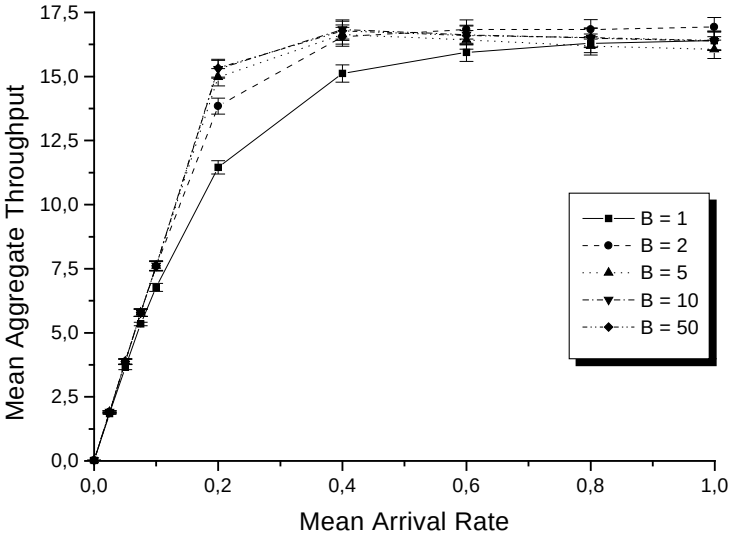

Figure 6.47: Mean aggregate throughput (packets/frame) vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, $W = 8$, and different $B \in \{1, 2, 5, 10, 50\}$.

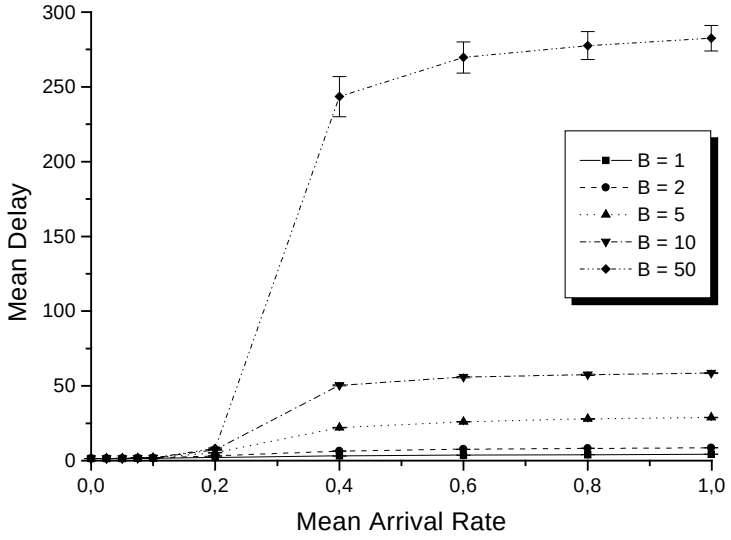

Figure 6.48: Mean delay (cycles) vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, $W = 8$, and different $B \in \{1, 2, 5, 10, 50\}$.

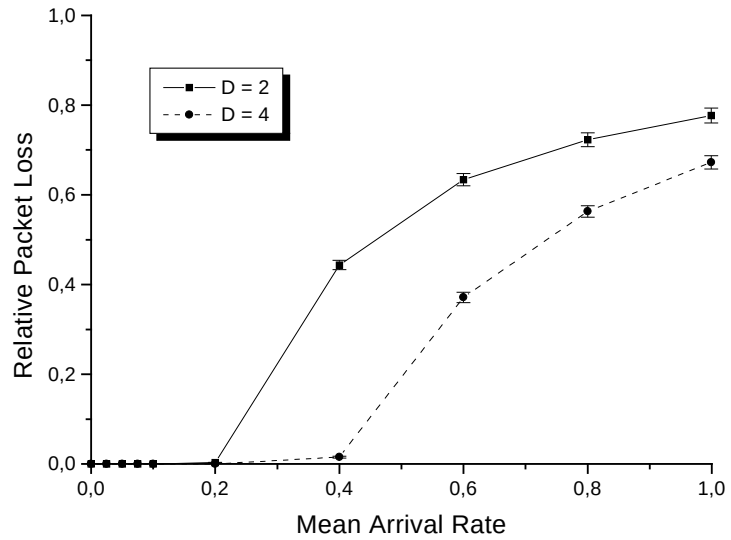

Figure 6.49: Relative packet loss vs. mean arrival rate for $b = 4$, $M = 60$, $W = 8$, and different $D \in \{2, 4\}$.

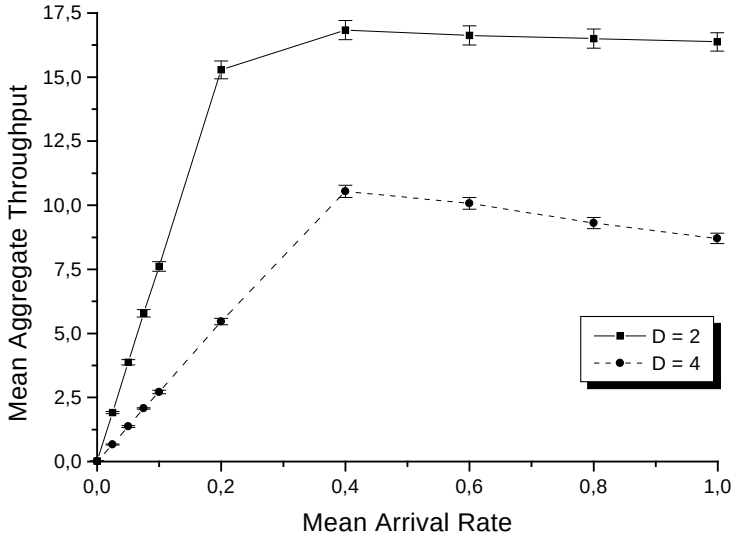

Figure 6.50: Mean aggregate throughput (packets/frame) vs. mean arrival rate for $b = 4$, $M = 60$, $W = 8$, and different $D \in \{2, 4\}$.

$D \cdot F = 400$ slots. For $D = 4$ we get $F = 100$ slots compared to $F = 200$ slots for $D = 2$. As a consequence, the length of data packets is decreased for increasing D. For $D = 4$ short data packets are $K = (F - M) = 40$ slots and long data packets are $F = 100$ slots long compared to $K = 140$ and $F = 200$ slots for $D = 2$, respectively. This results in a decreased mean aggregate throughput, as depicted in Fig. 6.50. Intuitively, we expect $D = 4$ to provide a smaller mean delay than $D = 2$ since in the former case each cycle contains as twice as many reservation slots as in the case $D = 2$ and the spatial reuse factor is doubled. We observe from Fig. 6.51 that this is true only for light to medium traffic loads. For a mean arrival rate larger than about 0.45 we can see that $D = 2$ provides smaller delays. This is due to the fact that each buffer can store up to 10 long data packets for $D = 2$ whereas for $D = 4$ each buffer is able to store up to 20 long data packets owing to the halved length of long data packets. As a consequence, for $D = 4$ at medium to high loads there are packets which have to wait for a longer time interval until they are transmitted resulting in an increased mean delay (and a decreased packet loss, as shown in Fig. 6.49). Again, we witness the trade–off between packet loss and delay. In the following simulations we set $D = 2$ for providing high throughput and low delay for a wide range of traffic loads. Note that the larger packet loss for $D = 2$ can be reduced by higher protocol layers. Protocols above the MAC layer have to make sure that packets are put in the MAC buffer such that overflow is completely avoided or kept below an acceptable level.

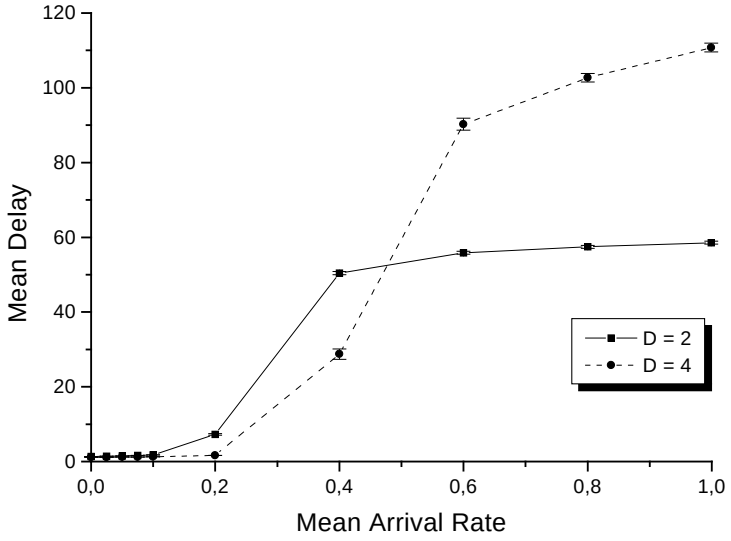

Figure 6.51: Mean delay (cycles) vs. mean arrival rate for $b = 4$, $M = 60$, $W = 8$, and different $D \in \{2, 4\}$.

6.4.3 Circuit switching

So far, we have considered only packet switching. We now take beside packet switching also circuit switching into account and show the positive impact of reservation ALOHA (R–ALOHA, see Section 5.3) on the throughput. For this purpose, we introduce two additional parameters r and ρ:

- A generated data packet belongs to a circuit with probability r, where $0 \leq r \leq 1$. Note that the case considered so far where no packet sets up a circuit corresponds to $r = 0$.
- The circuit holding time is geometrically distributed with parameter ρ, where $0 \leq \rho < 1$. The parameter ρ denotes the probability that a given circuit continues in the next cycle. Thus, the average circuit holding time is equal to $\frac{1}{1-\rho}$ cycles. Note that $\rho = 0$ corresponds to packet switching.

In the following, the mean aggregate throughput counts for both packet and circuit–switched data packets. The mean delay takes only nodes into account which are involved in the reservation process, i.e., packet switched packets and the first packet of a circuit. Nodes which have set up a circuit do not contribute to the mean delay. This is because these nodes use fixed assigned slots and do not encounter any delay caused by pretransmission coordination.

Fig. 6.52 shows the mean aggregate throughput vs. the mean arrival rate for $b = 4$, $M = 60$, $R = 8$, $W = 8$, $B = 10$, $r = 0.3$, and different $\rho \in \{0, 0.5, 0.8, 0.9, 0.95\}$. We observe that with increasing ρ the mean aggregate throughput increases. This is due to the fact that a larger ρ translates into an increased mean holding time. After making a successful setup each circuit is held for a longer time interval resulting in a higher channel utilization and mean aggregate throughput. Since each node holding an active circuit indicates this by using one dedicated

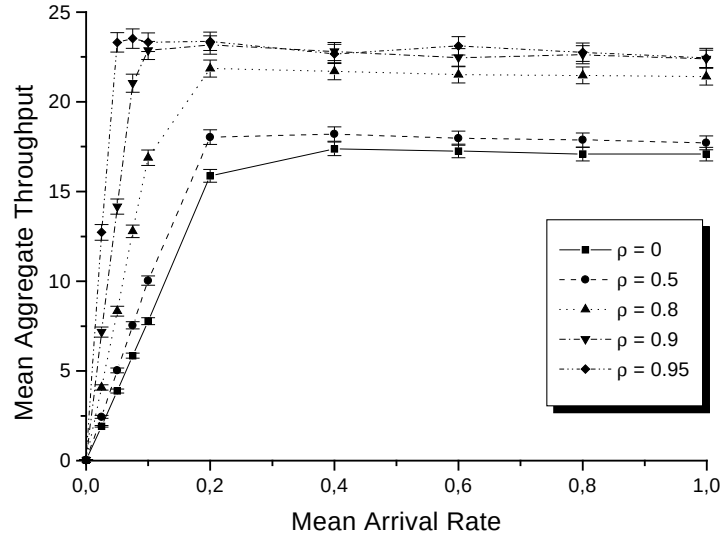

Figure 6.52: Mean aggregate throughput (packets/frame) vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, $W = 8$, $B = 10$, $r = 0.3$, and different $\rho \in \{0, 0.5, 0.8, 0.9, 0.95\}$.

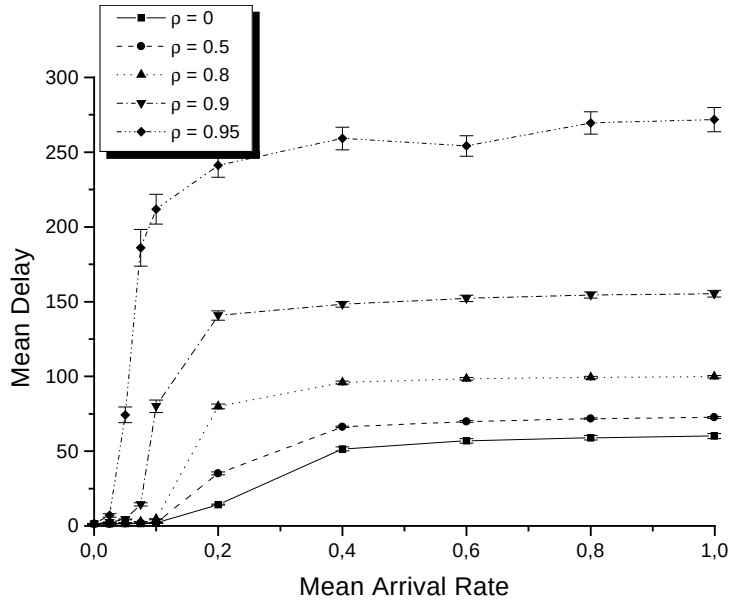

Figure 6.53: Mean delay (cycles) vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, $W = 8$, $B = 10$, $r = 0.3$, and different $\rho \in \{0, 0.5, 0.8, 0.9, 0.95\}$.

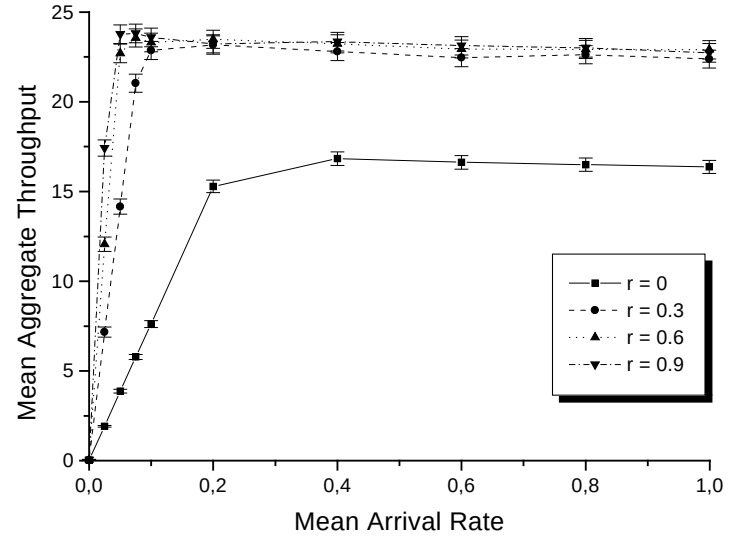

Figure 6.54: Mean aggregate throughput (packets/frame) vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, $W = 8$, $B = 10$, $\rho = 0.9$, and different $r \in \{0, 0.3, 0.6, 0.9\}$.

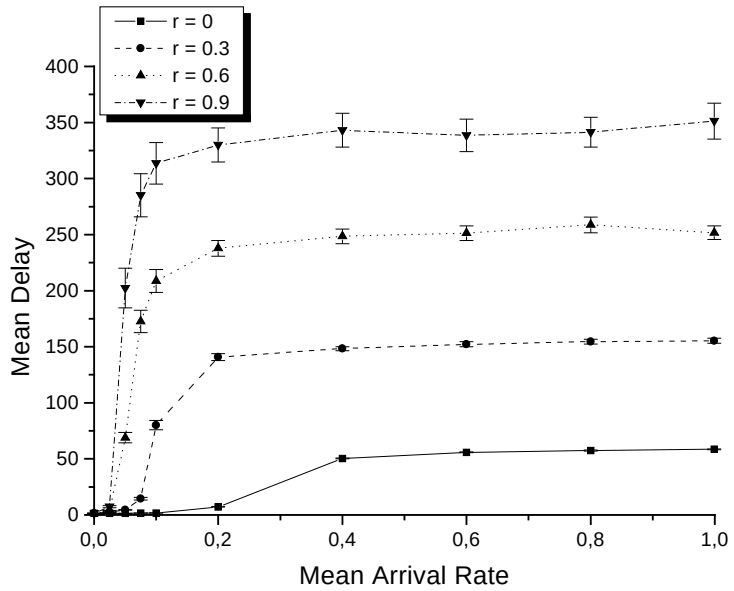

Figure 6.55: Mean delay (cycles) vs. mean arrival rate for $b = 4$, $M = 60$, $R = 8$, $W = 8$, $B = 10$, $\rho = 0.9$, and different $r \in \{0, 0.3, 0.6, 0.9\}$.

reservation slot (R–ALOHA, see Section 5.3), for an increasing ρ reservation slots are busy for a longer time interval. Consequently, for increasing ρ nodes which try to make a reservation can access fewer reservation slots resulting in more collisions of control packets. This in turn leads to more retransmissions and an increased delay encountered by these nodes, as depicted in Fig. 6.53.

Fig. 6.54 illustrates the mean aggregate throughput vs. the mean arrival rate for $b = 4$, $M = 60$, $R = 8$, $W = 8$, $B = 10$, $\rho = 0.9$, and different $r \in \{0, 0.3, 0.6, 0.9\}$. Apparently, with circuit switching the wavelengths are used more efficiently than with packet switching leading to an increased mean aggregate throughput. This is because in R–ALOHA nodes which have successfully set up a circuit can send data packets in each cycle since the corresponding control packets do not suffer from channel collisions. In Fig. 6.55 we observe that with increasing r, i.e., an increasing fraction of circuit–switched data packets, the mean delay for packet switched packets and packets setting up a circuit increases significantly. Again, this is because other nodes which try to make a reservation can access fewer reservation slots resulting in more control packet retransmissions and an increased mean delay.

6.4.4 Benchmark comparison

In this section, we discuss the efficiency of our AWG based network and MAC protocol and compare it with a previously reported single–hop metro WDM network that is based on a PSC. The nodes in the PSC network deploy the so–called *dynamic time–wavelength division multiaccess* (DT–WDMA) protocol for resolving packet collisions [CDR90]. Recall that DT–WDMA has already been discussed in Section 3.3.3, but for convenience we briefly review the underlying architecture and the main features of DT–WDMA below. We have chosen DT–WDMA since among the MAC protocols designed specifically for multiwavelength single–hop WDM networks based on a PSC, DT–WDMA has perhaps been the *most influential* one and has spurred a number of research papers on modified and improved versions [RS98]. Moreover, DT–WDMA has the following properties in common with our protocol, which allow for a reasonably fair benchmark comparison:

- For data transmission/reception each node is equipped with *one single* transceiver.
- Each node's receiver is *tunable*. Consequently, receiver collisions can potentially occur and have to be resolved by the access protocol.
- DT–WDMA belongs to the category of *reservation* protocols.
- Resources are assigned dynamically on–demand on a *per–packet* basis.
- Nodes are able to acquire and maintain *global knowledge*.
- Explicit acknowledgements are *not* required.

Next, we briefly describe the network architecture and DT–WDMA (for details the interested reader is referred to [CDR90]). DT–WDMA is proposed for a metropolitan–sized single–hop WDM network employing fixed–tuned and tunable transceivers attached to a PSC. More precisely, each node is equipped with one transceiver fixed tuned to a common control channel. For data each node deploys one transmitter fixed–tuned to a separate wavelength, i.e., each node has its own home channel for transmission, and one tunable receiver. Control information is sent over a dedicated signalling channel. Time is divided into slots (which correspond to frames in

our protocol) on each channel and slots on the control channel are further split into mini–slots (which correspond to slots in our protocol). Fixed time–division multiaccess is used within each slot on the control channel where one mini–slot is dedicated to each node. Transmitters indicate their intention to transmit a packet by transmitting the destination address during their assigned mini–slot in the control channel and then transmit their packet in the next slot on their data channel, i.e., DT–WDMA belongs to the family of tell–and–go reservation protocols. Receivers listen to the control channel and tune to the appropriate channel to receive packets addressed to them. A common distributed arbitration algorithm is used to resolve conflicts when packets from multiple transmitters contend for the same receiver. Each receiver executes the same deterministic algorithm to choose one of the contending packets. Each transmitter uses the same algorithm to determine the success or failure of its packet. This eliminates the need for explicit acknowledgements. In case of failure the corresponding node re–starts the reservation procedure.

Note that in DT–WDMA each node has its own home channel (wavelength) for transmission, i.e., the number of nodes equals the number of wavelengths, while receivers are assumed to be tunable over all these wavelengths. Thus, for large populations receivers with a large tuning range are required whose large tuning time significantly decreases the channel utilization. In our comparison we use 16 wavelengths which allows for deploying fast tunable receivers in DT–WDMA (and fast tunable transceivers in our network) whose tuning time is negligible. Furthermore, to compare our above simulation studies with DT–WDMA we set the number of nodes to $N = 200$ in both the AWG and PSC based networks. To accomodate 200 nodes in DT–WDMA we let all data wavelengths be equally shared among the nodes while we assume that each of the 200 nodes has its own mini–slot on the control channel. In our comparison we consider only packet switching and focus on the throughput performance. It was shown in [CDR90] that under the same assumptions on packet arrival process (Bernoulli) and traffic pattern (uniform traffic) as made in our above simulation studies the maximum utilization of each of the 16 wavelength channels equals 0.6 in DT–WDMA. This translates into a *maximum* aggregate throughput of $0.6 \cdot 16 = 9.6$ packets/slot, which is identical to 9.6 packets/frame for one slot in DT–WDMA corresponds to one frame in our protocol.

Fig. 6.56 depicts the mean aggregate throughput of our network vs. mean arrival rate for $D \in \{2, 4\}$ (as already shown in Fig. 6.50) and compares it with the maximum aggregate throughput of DT–WDMA. We observe that for $D = 4$ the mean aggregate throughput of our network is approximately equal to the maximum aggregate throughput of DT–WDMA at medium to high traffic loads. However, for $D = 2$ our network clearly outperforms DT–WDMA in terms of throughput. For a wide range of arrival rates our proposed network provides a mean aggregate throughput that is about 70% larger than the maximum aggregate throughput of DT–WDMA.

Next, let us consider the efficiency of DT–WDMA and our protocol. As mentioned above, the maximum *wavelength utilization* of DT–WDMA is 60%. Fig. 6.56 shows that for $D = 2$ our protocol yields a mean aggregate throughput of up to 16.8 packets/frame. Note that this results in a mean wavelength utilization of more than 100%. To see this, recall that 16 wavelengths are used in both PSC and AWG based networks. However, due to spatial wavelength reuse the 2×2 AWG provides twice as many communication channels than the PSC leading to a larger aggregate throughput. With this in mind, we take a look at the *channel utilization* of both networks. Clearly, in DT–WDMA the channel utilization and the wavelength utilization are the same, since each wavelength creates one channel. Whereas in our 2×2 AWG based network each wavelength creates two channels resulting in a mean channel utilization of approximately 53%. Recall that

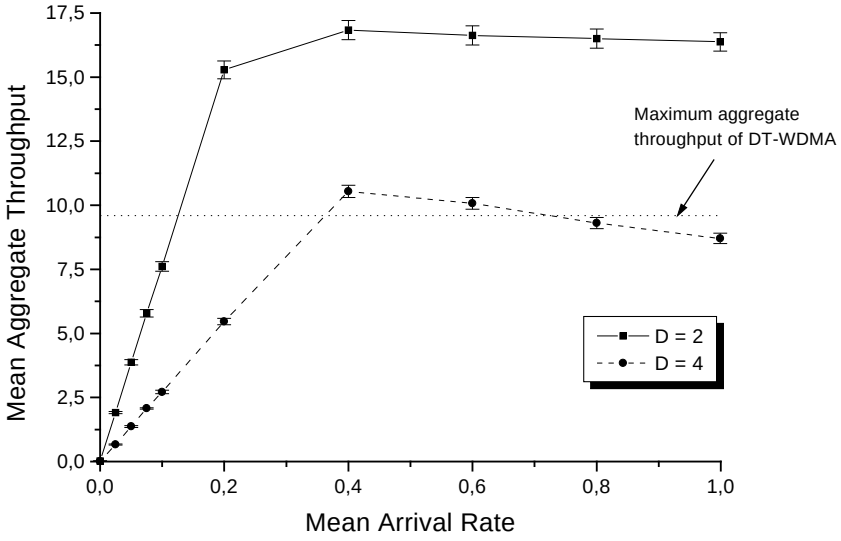

Figure 6.56: Mean aggregate throughput (packets/frame) for $b = 4$, $M = 60$, $W = 8$, and $D \in \{2, 4\}$, compared to maximum aggregate throughput (packets/frame) of DT–WDMA vs. mean arrival rate.

in our protocol wavelengths can not be spatially reused during the first M slots of each frame when nodes are involved in the reservation process. Consequently, our protocol can not fully capitalize on spatial wavelength reuse (in Chapter 9 we present a modified architecture which allows for *full* spatial wavelength reuse at any given time). However, we have seen in Section 6.3 that these M reservation slots can be used very efficiently for multicasting. Furthermore, with circuit switching the channel utilization of our network can be increased significantly, as illustrated in Figs. 6.52 and 6.54. Finally, we note that further channel utilization improvements could be achieved by replacing our simple first–come–first–served–first–fit algorithm with more efficient scheduling algorithms.

6.5 Conclusions

In the preceding sections we have investigated the impact of various network architecture and protocol parameters on the network performance by means of analysis and/or simulation. We have analyzed and verified by simulation the benefit of using multiple FSRs of the underlying AWG, spatially reusing all wavelengths at each AWG input port, and deploying the wavelength–insensitive splitters for efficiently enabling optical multicasting. Additionally, through extensive simulations we have examined self–stability, packet loss, and circuit switching while relaxing the assumptions made in our analyses.

In our investigations we have started with the special case of uniform unicast packet switching. Packets were assumed to be one frame long, i.e., spatial wavelength reuse was not considered. We have seen that using multiple FSRs of the underlying AWG significantly improves the throughput–delay performance of the network. More specifically, by exploiting three instead of one FSR the maximum mean aggregate througput is increased by approximately 88% resulting

in a channel utilization of up to 78%. Deploying multiple FSRs is useful when new nodes join the network and/or nodes are upgraded with technologically advanced transceivers with a larger tuning range.

By supporting variable–size packets the network becomes more flexible. Moreover, packets whose length is small enough are able to benefit from spatial wavelength reuse. In doing so, the utilization of wavelength resources is increased leading to an improved network efficiency and cost–effectiveness. We have seen that spatial wavelength reuse increases the mean aggregate throughput by more than 60% while reducing the mean delay by approximately 50%. It was shown that the slotted ALOHA control channel can represent a bottleneck. In particular, the parameter D, i.e., the physical degree of the AWG, has to be set carefully. Choosing $D = 4$ provides a good trade–off between spatial wavelength reuse and collisions of control packets.

Next, we have examined by means of analysis and simulation not only unicast but also multicast traffic. In our AWG based single–hop network multicasting can be done efficiently by (*i*) partitioning multicast transmissions in conjunction with spatial wavelength reuse and (*ii*) sending multicast data packets concurrently with broadcast control packets. For multicast traffic we have re–confirmed that the partitioning of multicast groups alleviates the destination conflicts but creates a channel bottleneck. We have demonstrated that the spatial wavelength reuse in the AWG based network effectively mitigates the channel bottleneck. The AWG based network achieves more than twice the transmitter throughput and roughly 30% larger receiver and multicast throughput compared to the widely studied single–hop networks based on a PSC. Furthermore, we have examined the interplay between multicast transmissions concurrently with control traffic and unicast transmissions with spatial wavelength reuse. We found that multicast transmissions concurrently with the broadcast control traffic effectively exploit the tuning of the receivers to the control slices, resulting in a significantly increased receiver utilization of almost 60% for a 50–50 mix of unicast and multicast traffic.

Extensive simulations have shown that self–stability can be easily achieved by means of back–off. The (relative) packet loss is decreased significantly by relieving the system bottlenecks — namely, reservation channel (increasing parameter M) and number of available FSRs (increasing R) — and extending the scheduling window and the buffer size at each node. Moreover, due to spatial wavelength reuse a mean wavelength utilization of more than 100% and a mean channel utilization of approximately 53% are achieved. In our benchmark comparison we have shown that for a wide range of traffic loads our AWG based network provides a mean aggregate throughput that is about 70% larger than the maximum aggregate throughput of DT–WDMA that runs on top of a PSC based single–hop WDM network. We have also seen that in presence of circuit–switched traffic the network throughput is increased. In particular for large fractions of circuit–switched traffic and long circuit holding times the throughput is improved significantly. However, with a larger number of active circuits fewer reservation slots are available to the remaining nodes resulting in an increased delay for nodes trying to make new reservations. Admission control would help solve this problem.

After gaining some insight in how the different parameters affect the network performance we are now in a position to investigate the proper setting of the parameters. We have seen that especially the physical degree of the AWG plays an important role in achieving a good overall network performance. The parameters have to be set such that the network performance is optimized in terms of throughput and/or delay under varying traffic conditions. This network dimensioning problem is addressed in greater detail in the following chapter.

Chapter 7

Network Dimensioning and Reconfiguration

In the preceding chapter we have gained some insight in how the various network parameters influence the network performance by varying them around their default values. We now go one step further in that we tackle the problem of setting the parameters properly in order to optimize the throughput–delay performance of our AWG based WDM network.

The throughput–delay optimization is motivated by efficiently enabling multiservice support, as discussed in Section 7.1. This naturally leads to a multiobjective optimization problem whose objective functions and constraints are described in Section 7.2. In Section 7.3 we develop a computationally efficient search and optimization approach which is based on genetic algorithms. Numerical results are presented in Section 7.4. Finally, we conclude in Section 7.5 and summarize the results on optimal network dimensioning and reconfiguration [**?**].

7.1 Multiservice convergence

To survive in today's competitive environment, networks have to provide multiple services in a cost–effective way in order to generate new revenue opportunities [SJ01]. The resulting multiservice networks have to carry a wide range of different traffic types, such as voice, IP, and Frame Relay. As schematically depicted in Fig. 7.1 a) – c), different types of traffic dominate during different times of the day [J. 01]. During office hours voice traffic dominates the network load. Whereas Internet (FTP, HTTP, email) and Frame Relay traffic play a major role in the evening and at night, respectively. By carrying these heterogeneous traffic types in a single converged network the utilization of the network resources can be significantly increased, as shown in Fig. 7.1 d). As a consequence, the resources of such a multiservice network are used more efficiently which is very important especially in cost–sensitive metropolitan networks.

Note that the traffic types shown in Fig. 7.1 have different throughput–delay requirements. While voice traffic is rather delay and jitter sensitive with moderate throughput requirements, data traffic is more bandwidth consuming but without stringent delay requirements. The network has to be dimensioned accordingly. More specifically, the network should be reconfigurable such that it is able to closely track changing traffic conditions for different times of the day. This rearrangeability is one of the key features of future WDM networks since it increases the utilization of network resources and thereby helps reduce network costs [LA90][LA91][LHA94]. Such a traffic–adaptive network is dynamically optimized resulting in an increased network effi-

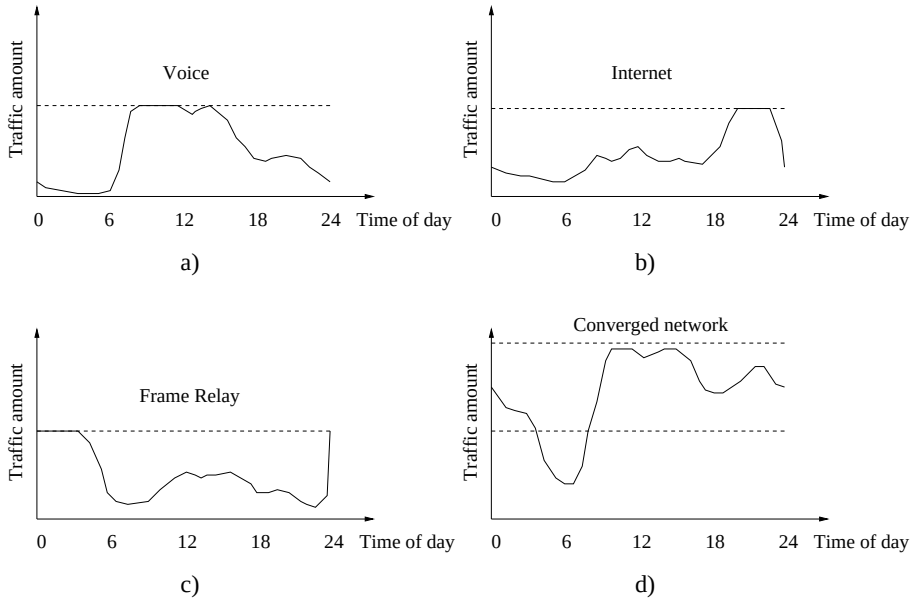

Figure 7.1: Different time–varying traffic types.

ciency [BR99a][BR01]. Under predominantly delay–sensitive traffic the network has to provide minimum delay. Whereas in situations with largely throughput–sensitive data traffic the network throughput has to be maximized. These two objectice functions naturally give rise to the so–called multiobjective optimization problem which is addressed in the next section.

7.2 Multiobjective optimization

Conventional optimization methods allow for optimizing only one single objective function subject to a given number of constraints. Thus, in our case the optimizer had to decide between throughput and delay, or a weighted combination of both. The weighing, however, requires weight factors for both throughput and delay. These factors have to be chosen properly and have a significant impact on the optimization output. Such a so–called preference–based optimization does not necessarily provide the optimal solution since the optimizer gives certain preferences by means of weight factors [Deb01]. Conversely, in the multiobjective optimization two or more objective functions can be optimized simultaneously without requiring any weight factors. Therefore, the optimizer does not influence the search of optimal solutions a priori. In our case we simultaneously optimize two objective functions — throughput and delay — which are defined more precisely in the following section.

7.2.1 Objective functions

We optimize the AWG based network for packet switching of variable–size unicast data packets. We maximize the mean aggregate throughput and minimize the mean delay. The two objective functions are given by the mean aggregate throughput and the mean delay expressions which have been derived in Section 6.2. Here we briefly review these two objective functions of our optimization.

The average throughput of the network is defined as the average number of transmitting nodes in a slot and is given by:

$$TH = D^2 \cdot \frac{F \cdot E[\mathcal{L}] + K \cdot E[\mathcal{S}]}{F \cdot D}, \tag{7.1}$$

where $E[\mathcal{L}]$ is the expected number of successfully scheduled long packets (of size F slots) from a given (fixed) AWG input port to a given (fixed) AWG output port per cycle (of length $F \cdot D$ slots), and $E[\mathcal{S}]$ is the expected number of successfully scheduled short packets (of length $K = (F - M)$ slots) from a given (fixed) AWG input port to a given (fixed) AWG output port per cycle. (We note that the throughput given by (7.1) may also be interpreted as the average number of transmitted data packets per frame; for convenience we will use this packets/frame interpretation in our numerical work in Sections 7.3 and 7.4.) $E[\mathcal{L}]$ and $E[\mathcal{S}]$ are given by:

$$E[\mathcal{L}] = \tilde{q}\left\{R - \sum_{k=0}^{\min(R,M)} P(Z=k)(R-k)\right\} := \tilde{q} \cdot \varphi(\beta) \tag{7.2}$$

and

$$\begin{aligned} E[\mathcal{S}] &= (1-\tilde{q})\left[R - \sum_{k=0}^{R}(R-k)\cdot P(Z=k)\right] + \\ &\sum_{j=1}^{M-R}\gamma_j \sum_{m=j}^{M-R}\sum_{k=m+R}^{M}\binom{k-R}{m}(1-\tilde{q})^m \tilde{q}^{k-R-m}\cdot P(Z=k) := h(\tilde{q},\ \beta), \end{aligned} \tag{7.3}$$

where

$$\tilde{q} = \frac{q \cdot \tilde{\alpha} \cdot M}{D \cdot \varphi(\beta)} \cdot \frac{\beta - \alpha}{\tilde{\alpha} - \alpha} \tag{7.4}$$

and

$$P(Z=k) = \binom{M}{k}\left(\frac{\beta e^{-\beta}}{D}\right)^k \left(1 - \frac{\beta e^{-\beta}}{D}\right)^{M-k},\ k = 0, 1, \ldots, M. \tag{7.5}$$

For $\lfloor F/K \rfloor > 1$,

$$\gamma_j = \sum_{\{m:m\le v_j\}} \binom{R}{m}\tilde{q}^m(1-\tilde{q})^{R-m}, \tag{7.6}$$

with

$$v_j = \min\left(R,\ \frac{(D-1)R\lfloor\frac{F-M}{K}\rfloor - j}{\lfloor\frac{F}{K}\rfloor - 1} + R\right). \tag{7.7}$$

In case $\lfloor F/K \rfloor = 1$, γ_j becomes

$$\gamma_j = \begin{cases} 1 & \text{if} \quad j < (D-1)R\lfloor\frac{F-M}{K}\rfloor \\ 0 & \text{otherwise.} \end{cases} \tag{7.8}$$

β is given as the solution of the equation:

$$(1-q)\cdot\frac{\tilde{\alpha}\cdot M}{D}\cdot\frac{\beta-\alpha}{\tilde{\alpha}-\alpha} = h\left(\frac{q\cdot\tilde{\alpha}\cdot M}{D\cdot\varphi(\beta)}\cdot\frac{\beta-\alpha}{\tilde{\alpha}-\alpha},\ \beta\right), \tag{7.9}$$

with $\tilde{\alpha} := S\sigma/M$ and $\alpha := Sp/M$. Equation (7.9) is solved numerically and the obtained β is inserted in (7.4) to obtain $\tilde{q}$. With β and $\tilde{q}$, $E[\mathcal{L}]$ (7.2) and $E[\mathcal{S}]$ (7.3) are calculated.

The mean packet delay is defined as the average time period in slots from the generation of the control packet corresponding to a data packet until the transmission of the data packet. The average delay in the network in slots is:

$$Delay = \left\{ \frac{S}{D \cdot (E[\mathcal{L}] + E[\mathcal{S}])} - \frac{1-\sigma}{\sigma} \right\} \cdot D \cdot F. \tag{7.10}$$

Note that the derivation in Section 6.2 considered the case $M < F$, i.e., $K > 0$. In our optimization, we allow for $M \leq F$, i.e., $K \geq 0$; the objective functions for the special case $M = F$ are derived next. In case $M = F$, the length K of a short packet degenerates to zero and hence short packets do not contribute to the throughput. There are different scenarios for evaluating the network performance for the case $M = F$. One scenario is to still consider short packets in the control packet contention. In this scenario, the packet generation probability is unchanged at σ; and all short packets that are successful in the control slot contention are successfully scheduled, i.e.,

$$E[\mathcal{S}] \quad = \quad (1-\tilde{q}) \cdot \sum_{k=0}^{M} k \cdot P(Z=k) =: h(\tilde{q},\ \beta) \tag{7.11}$$

which replaces Eqn. (7.3).

An alternative scenario is to completely ignore short packets and to consider only long packets in the control slot contention and data packet scheduling. In this alternative scenario the packet generation probability is effectively $q \cdot \sigma$, and each generated packet is long with probability one. (The network equilibrium condition is $q \cdot \frac{\sigma}{D} \cdot E[\eta] = E[\mathcal{L}]$ in this scenario.)

We consider the first scenario, where control packets are sent for short packets (of length zero), in our network optimization in this chapter. We chose the first scenario because it ensures that the packet generation probability and thus the level of control packet contention are the same both for the case $M = F$ and the case $M < F$. The alternative scenario would result in a reduced level of control packet contention in the case $M = F$ (especially when q is small) and thus an unfair performance comparison.

7.2.2 Decision variables and constraints

The throughput–delay performance of the network largely depends on the network parameter settings. There are two types of parameters in our AWG based network: (1) Architecture (hardware) parameters, e.g., AWG degree, and (2) MAC protocol (software) parameters, e.g., length of frames in timing structure. For good network performance, these parameters have to be set properly. To obtain these settings the parameters are used as decision variables in the optimization.

Let us now identify the decision variables in our optimization problem and identify the constraints on the decision variables. We select the AWG degree D as the (independent) decision variable for the network (hardware) architecture; we determine the other architecture parameters — number of used FSRs R and combiner/splitter degree S — as functions of D and the given number of nodes N and transceiver tuning range Λ (see Table 5.1), as discussed shortly. Generally, the decision variable D can take any integer value satisfying

$$D \geq 2 \qquad \text{and} \qquad D \leq \Lambda, \tag{7.12}$$

where Λ is the maximum number of wavelength channels accommodated by the fast tunable transceivers employed in the considered network. In other words, Λ is the maximum tuning range of the employed transceivers divided by the channel spacing and is thus very technology dependent. (We use the typical value $\Lambda = 8$ in our numerical investigations in Sections 7.3 and 7.4.) We also note that the number of ports of commercially available photonic devices is typically a power of two. We can easily incorporate this constraint by restricting D to the set $\{2,\ 4,\ 8, \ldots\}$.

The number of used FSRs R depends on the (independent) decision variable D and the given tuning range Λ of the transceivers. Generally, R must be an integer satisfying $R \cdot D \leq \Lambda$, i.e., $R \leq \Lambda/D$. The larger R, the more parallel channels are available between each input–output port pair of the AWG, and hence the larger the throughput. Therefore, we set R to the largest integer less than or equal to Λ/D, i.e., $R = \lfloor \Lambda/D \rfloor$. We note that the tuning range Λ and degree D are typically powers of two for commercial components. Hence, Λ/D is a power of two for practical networks, and we may write $R = \Lambda/D$. The combiner/splitter degree S depends on the decision variable D and the given number of nodes in the network N. In determining the combiner/splitter degree S, it is natural to assume that the nodes are equally distributed among the D AWG input/output ports. This arrangement minimizes the required combiner/splitter degree S, which in turn minimizes the splitting loss in the combiners/splitters. Hence, we set $S = \lceil N/D \rceil$.

We now turn to the protocol (software) parameters number of slots per frame F, number of reservation slots per frame M, and retransmission probability of unsuccessful control packets p (see Table 5.1). We identify three decision variables; these are F, M, and p. Generally, the number of slots per frame F can take any positive integer value, i.e., $F \geq 1$, while the number of control slots per frame can take any positive integer value less than or equal to F, i.e., $1 \leq M \leq F$. We note that the size of the data packets to be transported may impose additional constraints on F and M. With a given maximum packet size, F must be large enough to accommodate the maximum–size packet in a frame. If short packets have a specific size requirement, $(F - M)$ should be large enough to accommodate that packet size. For our numerical work in Sections 7.3 and 7.4 we do not impose packet size requirements. Instead, we determine the F and M values that give the optimal throughput–delay performance, subject only to $F \geq 1$ and $1 \leq M \leq F$. The packet retransmission probability p may take any real number in the interval $[0,\ 1]$. To reasonably limit the search space we restrict p to $[0,\ 0.05,\ 0.10,\ 0.15, \ldots,\ 1.0]$ in our numerical work.

7.2.3 Pareto optimality

The familiar notion of an optimal solution becomes somewhat vague when an optimization problem has more than one objective function. A solution that proves best by one criterion may rate among the poorest on another. We can find a set of optimal solutions for each objective function. However, without any further information, no solution from the set of optimal solutions can be said to be better than any other. Since a number of solutions are optimal, in a multiobjective optimization problem many such (trade–off) optimal solutions are important. This is the fundamental difference between a single–objective and a multiobjective optimization problem. In a multiobjective optimization, ideally an effort must be made in finding the set of trade–off optimal solutions by considering all objectives to be important. After a set of such trade–off solutions is found, a user can then use higher–level qualitative considerations, such as the traffic patterns illustrated in Fig. 7.1, to make a choice.

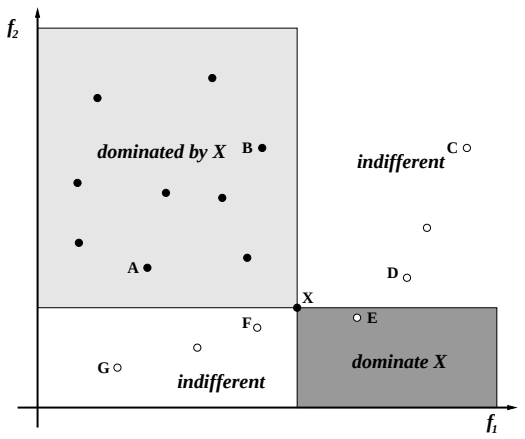

Figure 7.2: Pareto–optimal solutions.

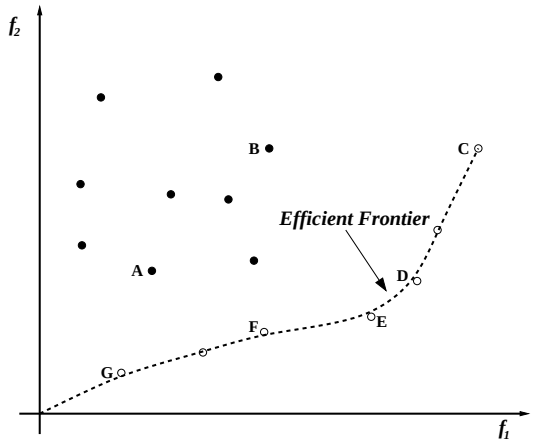

Figure 7.3: Efficient frontier.

A feasible solution to a multiobjective optimization problem is referred to as efficient point or *Pareto–optimal* solution. An illustrative example is shown in Figs. 7.2 and 7.3. We have two objectives — maximizing f_1 and minimizing f_2 — in the example. The region which is shaded in light gray is said to be dominated by the point X. All points in the region, e.g., A and B have larger values of f_2 and smaller values of f_1 than the point X. Clearly, the point X is superior to the points A and B. Thus all points in the light gray rectangle are dominated by point X. All points in the dark gray rectangle, e.g., the point E, are said to dominate the point X, since all points in the dark gray rectangle have larger values of f_1 and smaller values of f_2 than X. The point E is thus superior to the point X. Based on the concept of *Pareto dominance*, the optimality criterion for multiobjective problems can be introduced. Now, consider the points C, D, E, F, and G. The point E is unique among C, D, F, and G in that its corresponding decision vector is not dominated by any other decision vector. That means the decision vector of the point E is optimal in the sense that it cannot be improved in any objective without causing a degradation in at least one other objective. None of these can be identified as better than the others unless preference information is included. All values of the points C, D, E, F, and G are indifferent to each other. The set of these solutions is termed as Pareto–optimal solution set or *efficient frontier*. The efficient frontier corresponding to Fig. 7.2 is shown in Fig. 7.3.

The goal of multiobjective optimization is to find such a feasible efficient frontier. Classical methods for generating the Pareto–optimal solution set aggregate the objectives into a single, parameterized objective function in analogy to decision making before search. However, the parameters of this function are not set by the decision maker, but systematically varied by the optimizer [Zit99].

In our multiobjective optimization problem the two objective functions — mean aggregate throughput and mean delay — are conflicting, as we will see shortly. The solution is an optimal trade–off curve between throughput and delay. This trade–off curve gives the smallest achievable delay as a function of the desired throughput, or conversely, the largest achievable throughput as a function of the tolerable delay. The Pareto–optimal solutions are crucial for (i) the planning and provisioning of new networks, i.e., to determine the best architecture (hardware) parameter values, and (ii) the efficient operation of installed network hardware. In the following, we use the Pareto–optimal throughput–delay trade–off curve in a two–step network optimization process.

First, we optimize a new network by finding the optimal architecture (hardware) parameter values. Second, after fixing the architecture we optimize the protocol (software) parameters for an existing architecture. Specifically, with the set of Pareto–optimal solutions we are able to operate the network at different points of its throughput–delay trade–off curve according to the traffic type that dominates at a given time of the day. Thus, the network can be dynamically reconfigured in order to provide varying degree of small delay (and moderate throughput) service or large throughput (and moderate delay) service as the traffic changes with the time of the day.

To acquire detailed knowledge of the optimal throughput–delay trade–off curve we apply genetic algorithms (GAs) which belong to the family of evolutinary algorithms, as discussed in the following section.

7.3 Genetic algorithm based approach

The concept of evolutionary algorithms is loosely based on the Darwinian principles of biological evolution. One of the most striking differences to classical search and optimization algorithms is that evolutionary algorithms use a population of solutions in each iteration, instead of a single solution. Since a population of solutions is processed in each iteration, the outcome of an evolutionary algorithm is also a population of solutions. If an optimization problem has a single optimum, all evolutionary algorithm population members can be expected to converge to that optimum. However, if an optimization problem has multiple optimal solutions, an evolutionary algorithm can be used to capture multiple efficient points in its final population in one single simulation run. Evolutionary algorithms are thus well suited to find the set of multiple Pareto–optimal solutions in a computationally efficient way [Deb01]. In our optimization we deploy genetic algorithms (GAs) which are members of the family of evolutionary algorithms. The computational effort required to obtain the optimal throughput–delay trade–off curve for a given traffic load with our GA based approach depends on the parameters and the size of the exhaustive search space. The number of parameter combinations that our approach needs to evaluate to obtain the Pareto–optimal solution set is usually on the order of thousand times smaller than the exhaustive search space. In typical scenarios, our approach requires less than one day of CPU time on a 933 MHz PC to find the optimal trade–off curve, whereas the exhaustive search would require several years of CPU time. The basic operation of GAs is explained next.

7.3.1 Basic principle

The basic structure of a GA is illustrated in Fig. 7.4. GAs encode each individual's decision vector into either a binary bit string or a string of decision variables called *chromosome*, and each bit or decision variable is referred to as a *gene*. The quality of an individual in the population with respect to the objective functions is represented by a scalar value, called *fitness*. After generating the initial population (by randomly drawing the decision variables for each individual from uniform distributions over the respective ranges of the decision variables), each individual is assigned a fitness value. The population is evolved repeatedly, generation by generation, using the *crossover* operation and the *mutation* operation. The crossover and mutation operations produce offspring by manipulating the individuals in the current population that have good fitness values. The crossover operation swaps portions of the chromosomes. The mutation operation changes the value of a gene. Individuals with a better fitness value are more likely to survive and to participate in the crossover (mating) operation. After a number of generations, the population contains members with better fitness values. The best individuals in the final

```
Genetic Algorithm()
{
    t = 0;                              //start with an initial generation
    init_population P(t);
    //initialize a usually random population of individuals
    evaluate P(t);
    //evaluate fitness of all individuals of initial population
    while not terminated do {           //evolution cycle;
        t ← t + 1;                      //increase the generation counter
        P'(t) = select_parents P(t);    //select a mating pool for offspring production
        recombine P'(t);                //recombine the 'chromosome' of selected parents
        mutate P'(t);                   //perturb the mated population stochastically
        evaluate P'(t);                 //evaluate fitness of new generation
        P(t) ← P'(t);
    }
}
```

Figure 7.4: Basic structure of a genetic algorithm.

population are the outcome of the genetic algorithms. Each operation is discussed in detail in the following subsections.

7.3.2 Genetic algorithm comparison

There is a large number of GAs to choose from [Deb01]. Each GA uses a different fitness function. To find out which GA is well suited to our optimization problem we compare three common GAs and their fitness functions in the following.

The fitness function is typically one of the objective functions of the problem. However, other objective functions or a combination of objective functions can be used to increase the search performance. We evaluate three commonly used types of fitness function. We generate $G = 20$ generations, each with a population size of $P = 200$ to compare the quality of the fitness functions. We set the probability of crossover to 0.9 and the probability of mutation to 0.05, which are typical values. We compare the GA outputs with the true Pareto–optimal solutions which were found by conducting an exhaustive search over all possible combinations of the decision variables. We fix the mean arrival rate of $\sigma = 0.6$ and the fraction of long data packets $q = 0.1$ for this evaluation (see Table 5.1). All results presented in this chapter are for a channel spacing of 200 GHz, i.e., 1.6 nm at 1.55 μm. Thus, we can use 7 – 10 wavelengths at each AWG input port with fast tunable transceivers with a tuning range of $10 - 15$ nm. For all subsequent results, the number of wavelengths is fixed at eight, i.e., $\Lambda = 8$. D can take the values 2, 4, and 8. Thus, the corresponding R values are 4, 2, and 1. We fix the number of nodes in the network at $N = 200$. To reasonably limit the search space of the GA, we restrict F to be smaller than 400 slots. We note that with a large F, the considered network generally achieves larger throughput values (at large delays), however, the computational effort for evaluating a given parameter combination increases as F increases. For the exhaustive search, we therefore limit F to values less than or equal to 200 slots.

First, we evaluate the *Vector Evaluated Genetic Algorithm* (VEGA), which is easy to im-

VEGA w/o elitism	WBGA w/o elitism	RWGA w/o elitism
15	23	13
VEGA with elitism	WBGA with elitism	RWGA with elitism
55	82	115

Table 7.1: Number of Pareto–optimal solutions in final population for genetic algorithm based search with $F \leq 400$; exhaustive search for $F \leq 200$ gives 580 Pareto–optimal solutions.

plement. This algorithm divides the population into subpopulations according to the number of objective functions. Each subpopulation is assigned a fitness value based on a different objective function. In this way, each objective function is used to evaluate some members in the population. In our case, we have two objective functions. We divide the population into two subpopulations, each subpopulation uses one of the two objective functions as their fitness function. When using only one objective function for the fitness value of a subpopulation, it is likely that solutions near the optimum of an individual objective function are preferred by the selection operator in a subpopulation. Such preferences take place in parallel with other objective functions in different subpopulations. The main disadvantage of VEGA is that typically after several generations, the algorithm fails to sustain diversity among the Pareto–optimal solutions and converges near one of the individual solutions. In fact, the VEGA finds only 15 Pareto–optimal solutions compared to 580 Pareto–optimal solutions found by exhaustive search, as shown in Table 7.1. The efficient frontier spanned by these solutions is plotted in Fig. 7.5. We observe, however, that the VEGA efficient frontier is overall quite close to the true efficient frontier (found by exhaustive search).

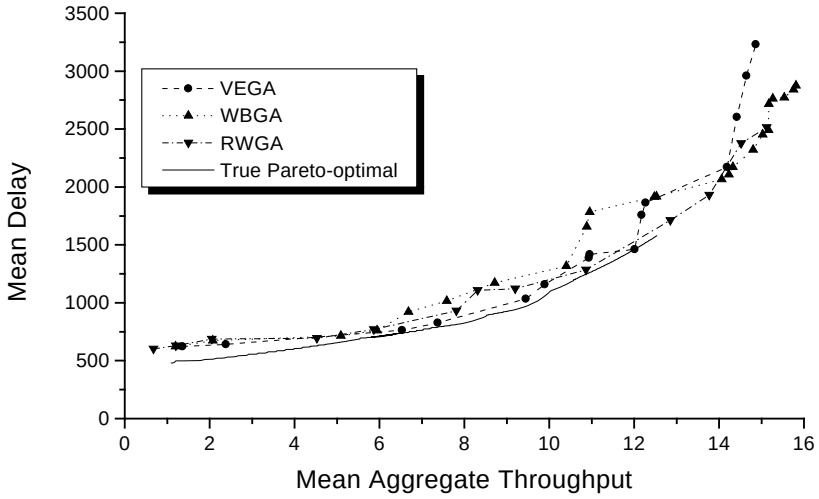

Figure 7.5: Efficient frontiers obtained with different genetic algorithms without elitism for $F \leq 400$ and with exhaustive search for $F \leq 200$.

Next, we evaluate the *Weight Based Genetic Algorithm* (WBGA) which uses the weighted sum of each objective function as fitness function, i.e.,

$$Fitness = \sum_{i=1}^{n} w_i \cdot f_i, \tag{7.13}$$

where n denotes the number of objective functions and w_i denotes the fixed weight for objective function f_i (with $\sum_{i=1}^{n} w_i = 1$). The main difficulty in WBGA is that it is hard to choose the weight factors. We use the same weight factor of $1/2$ for each objective function. Since the mean delay should be minimized in our problem, we use the negative delay as the second objective function. The fitness function used is

$$Fitness = \frac{1}{2} \cdot TH - \frac{1}{2} \cdot Delay. \tag{7.14}$$

Our goal is to maximize the mean aggregate throughput while minimizing the mean delay. Thus, with the WBGA approach, the larger the fitness value, the better. We observe from the results given in Fig. 7.5 and Table 7.1 that the WBGA finds more Pareto–optimal solutions than VEGA. However, the WBGA efficient frontier has parts (particularly in the throughput range from $9 - 13$ packets/frame) that are distant from the true efficient frontier. We note that the average network delay given in (7.10) in units of slots is on the order of thousands of slots in typical scenarios, whereas the average throughput is typically on the order of one to 16 packets per frame. To achieve a fair weighing of both throughput and delay in the fitness function, we use the delay in unit of cycles (where one cycle corresponds to $D \cdot F$ slots) in the evaluation of the fitness in (7.14) (and the following fitness definition in (7.15)); with this scaling, the delay is on the order of $1 - 20$ cycles in typical scenarios.

Finally, we evaluate the *Random Weight Genetic Algorithm* (RWGA) which weighs the objective functions randomly. The fitness function is

$$Fitness = \sum_{i=1}^{n} w_i \cdot f_i, \tag{7.15}$$

where n denotes the number of objective functions. A new independent random set of weights $w_i,\ i = 1, \ldots, n$, is drawn from a specific distribution each time an individual's fitness is calculated. We use the fitness function

$$Fitness = \varepsilon \cdot TH - (1 - \varepsilon) \cdot Delay, \tag{7.16}$$

where ε is uniformly distributed in the interval $(0, 1)$. We observe from Fig. 7.5 that the RWGA efficient frontier is relatively far from the true efficient frontier in the throughput range from $9 - 12$ packets/frame. Also, the RWGA finds only a relatively small number of Pareto–optimal solutions.

We now study the concept of *elitism*. Elitism is one of the schemes used to improve the search; with elitism the good solutions in a given generation are kept for the next generation. This prevents losing the already found good solutions in the subsequent crossover operation(s), which may turn good solutions into bad ones. For each generation we determine the Pareto–optimal solutions by comparing the throughput and delay achieved by the individuals in that generation. (Note that the thus determined Pareto–optimal solutions are not necessarily the true Pareto–optimal solutions to the optimization problem, rather they are Pareto–optimal with

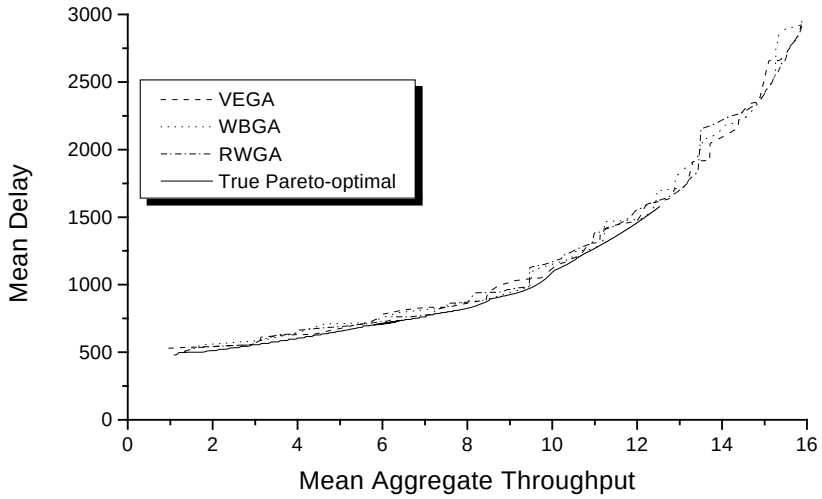

Figure 7.6: Efficient frontiers obtained with different genetic algorithms with elitism for $F \leq 400$ and with exhaustive search for $F \leq 200$.

respect to the other individuals in the considered generation.) The determined Pareto–optimal solutions are kept for the next generation; they are not subjected to the crossover operation, they are, however, subjected to the mutation operation (as explained in Section 7.3.4). If we find that a Pareto–optimal solution from a previous generation is no longer Pareto–optimal solution in a new generation, i.e., it is dominated by some other individual in the new generation, then this old Pareto–optimal solution is discarded.

The results obtained with elitism are given in Fig. 7.6 and Table 7.1. We observe that the number of Pareto–optimal solutions in the final population is dramatically larger and the efficient frontiers are closer to the true efficient frontier of the problem. From Fig. 7.6 it appears that all schemes with elitism perform quite well, with RWGA hugging the true efficient frontier most closely. This observation is corroborated by comparing the number of Pareto–optimal solutions in the final population in Table 7.1, which indicates that RWGA gives the best performance. According to the observations made in this section, we use RWGA with elitism throughout the remainder of our optimization work.

7.3.3 Algorithm parameter tuning

Prior to explaining the genetic crossover and mutation operations in Section 7.3.4, we run some numerical experiments in order to investigate the impact of the population size and the number of generations on the search and optimization performance of the deployed RWGA.

Population size

The population size trades off the time complexity (computational effort) and the number of optimal solutions. In order to accommodate all Pareto–optimal solutions, the population should

be large enough. However, as the population size grows, the time complexity for processing a generation increases (whereby the most computational effort is typically expended on evaluating the throughput and delay achieved by an individual to determine its fitness value). On the other hand, for a smaller population, the time complexity for the population decreases while the population may lose some Pareto–optimal solutions. As a result, the smallest population size that can accommodate all Pareto–optimal solutions is preferable.

For schemes that employ elitism, we categorize the population in generation t into three groups. (i) The elite group of size $P_e(t)$ which contains the Pareto–optimal solutions from the preceding generation $(t-1)$, (ii) the reproduction group of size $P_p(t)$ which is reproduced from the individuals with good fitness values in the preceding generation $(t-1)$ through crossover, and (iii) the random group of size $P_r(t)$ which is generated randomly (by drawing the decision variables from uniform distributions over their respective ranges). The random group is required to prevent the algorithm from getting stuck in local optima. The population size should accommodate these three groups appropriately. Furthermore, the size of the reproduction group and the random group need to be carefully considered. If the reproduction group is too large, the solution may get stuck in a local optimum. If the size of the random group is too large, we may spend most of the time calculating the fitness values of solutions that are very distant from the efficient frontier. However, the population size should at least be larger than the elite group. To find the proper population size, we evaluate the adopted RWGA with elitism for the population sizes $P = 150,\ 200$, and 300. We initially set the size of the reproduction group to one half of the population size, i.e., $P_p^{\text{init}} = P/2$. Once the number of Pareto–optimal solutions in a generation $(t-1)$ exceeds P_p^{init}, i.e., $P_e(t) > P_p^{\text{init}}$, we set the size of the reproduction group to $P_p(t) = P - P_e(t)$ in the next generation. Thus $P_p(t) = \min\{P_p^{\text{init}},\ P - P_e(t)\}$. If the number of Pareto–optimal solutions in a generation $(t-1)$ is less than $P - P_p^{\text{init}}$, we set the size of the random group to $P_r(t) = P - P_p^{\text{init}} - P_e(t)$ in the next generation, otherwise we set $P_r(t) = 0$; i.e., $P_r(t) = \max\{0,\ P - P_p^{\text{init}} - P_e(t)\}$. Thus, the more Pareto–optimal solutions there are in the preceding generation, the fewer randomly generated individuals are in the next generation. (If the number of Pareto–optimal solutions in a generation exceeds P_p^{init}, the succeeding generation does not contain randomly generated individuals.) For the following evaluation, the parameters $\Lambda,\ \sigma,\ q$, and the ranges of $D,\ F,\ M$, and p are set as given in Section 7.2.2. For comparison, we set the number of generations to $G = 20,\ 15$, and 10, respectively. Thus, the total number of considered individuals is $P \cdot G = 3000$ in all cases. The results are shown in Fig. 7.7. We observe from Fig. 7.7 that all three efficient frontiers hug the true Pareto–optimal frontier quite closely, with all three curves having "humps" around a throughput of 14 packets/frame. The number of Pareto–optimal solutions obtained for the population sizes $P = 150,\ 200$, and 300 are 87, 104, and 70, respectively. The population size of $P = 150$ is not very well suited for our network optimization because it typically can not accommodate all the Pareto–optimal solutions. This is because the elite group takes up almost two thirds of the population. With a population size of $P = 300$ (and only $G = 10$ generations to ensure a fair comparison) the evolution of the generations does not settle down as much as for 20 and 15 generations and therefore gives only 70 Pareto–optimal solutions (although the efficient frontier has a relatively small "hump"). Overall, we conclude that all three considered population sizes give fairly good results. We choose $P = 200$ for the following experiments in this chapter as it appears to accommodate all three population groups in a proper fashion. In Fig. 7.8 we plot the efficient frontiers obtained with different initial sizes $P_p^{\text{init}} = 50$ and 100 of the reproduction group (with $P = 200$ fixed). The number of Pareto–optimal solutions for $P_p^{\text{init}} = 50$ and 100, are 85 and 115, respectively. We observe from Fig. 7.8 that both efficient frontiers are quite close to the true Pareto–optimal

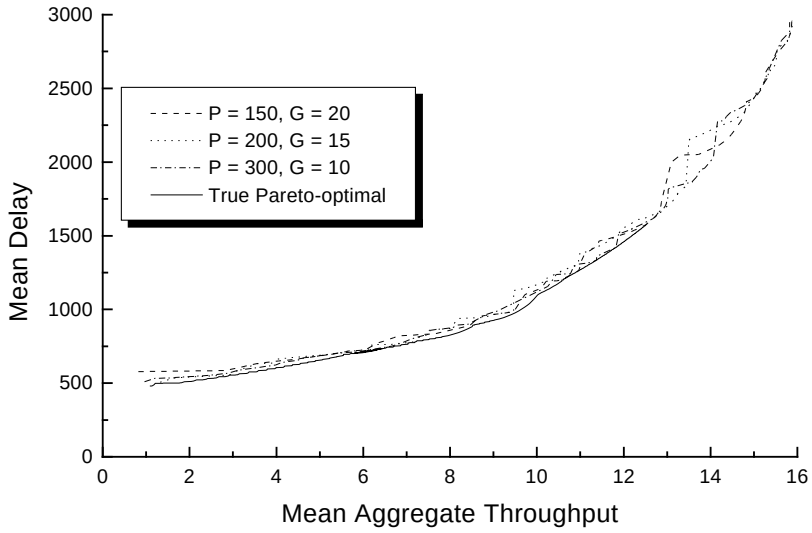

Figure 7.7: Efficient frontiers for different population sizes P with $P \cdot G = 3000$ fixed.

frontier. We set $P_p^{\text{init}} = 100$ for all the following experiments in this paper.

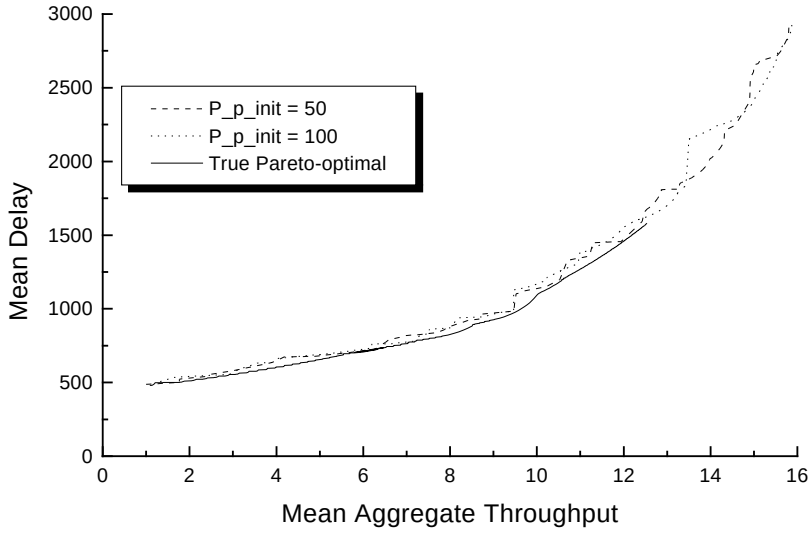

Figure 7.8: Efficient frontiers for different initial sizes P_p^{init} of the reproduction group with population size $P = 200$ fixed.

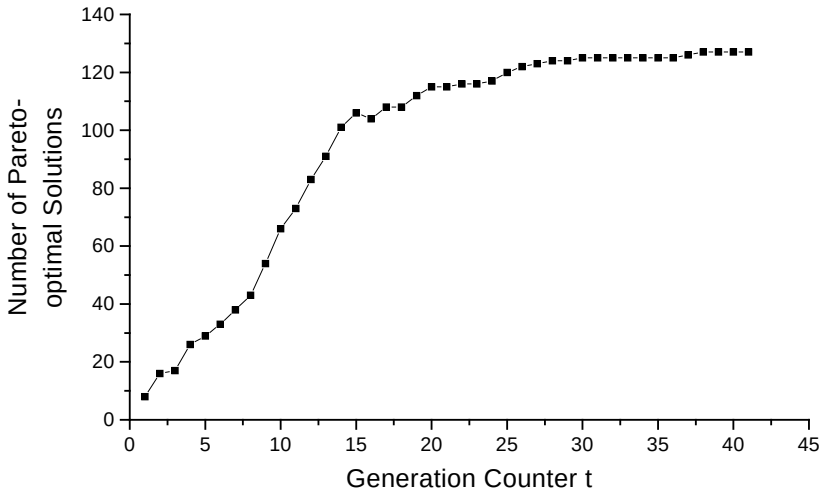

Figure 7.9: Size of elite group $P_e(t)$ as a function of generation counter t.

Number of generations

We now investigate the impact of the number of generations G. In Fig. 7.9, we plot the size of the elite group $P_e(t)$ as a function of the generation counter t. Recall that $P_e(t)$ is defined as the number of Pareto–optimal solutions in generation $(t-1)$; thus $P_e(1)$ is the number of Pareto–optimal solutions in the initial generation $t = 0$. In Fig. 7.10 we plot the sum of the fitness values of the individuals in the elite group $P_e(t)$ as a function of the generation counter. We observe from Fig. 7.9 that the number of Pareto–optimal solutions in a generation first steadily increases and then settles at a fixed value as the generations evolve. (The slight drop around the 15th generation is because we found a Pareto–optimal solution which dominates several earlier Pareto–optimal solutions.) We observe from Fig. 7.10 that the sum of the fitness values of the Pareto–optimal solutions in a generation first increases quickly, then fluctuates, and finally settles down as the generations evolve. This behavior is typical for GA based optimization and is due to the random nature of the evolution of the population. To allow for the evolution to settle down sufficiently, we set the total number of generations to $G = 40$. According to the decisions made in this section, we set the population size to $P = 200$, the number of generations to $G = 40$, and the initial size of the reproduction group to $P_p^{\text{init}} = 100$.

7.3.4 Genetic operations

We now explain the operation of GAs in general and RWGA in particular in more detail. We discuss how the two genetic operators crossover and mutation work.

Crossover operation

The crossover operation is used to produce the reproduction group in the next generation from the individuals in the current population. The main objective of the crossover operation is to

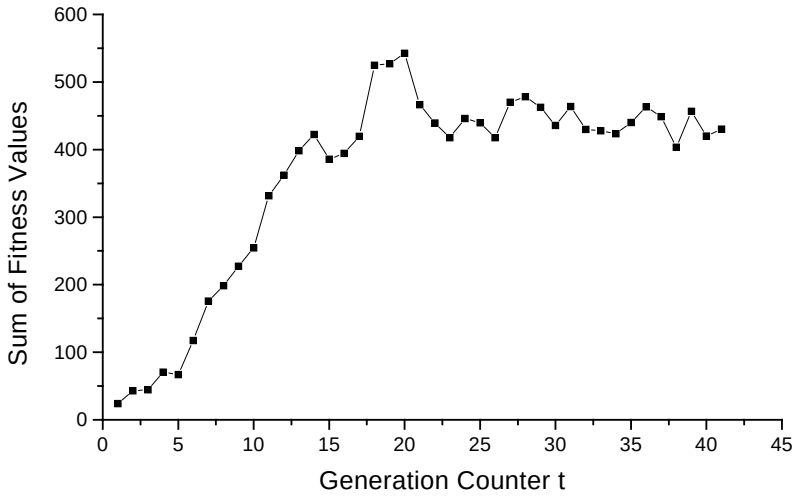

Figure 7.10: Sum of fitness values of individuals in elite group as a function of the generation counter t.

produce offspring that have large fitness values (and satisfy the problem's constraints). This is achieved by swapping parts of the chromosomes of the fittest individuals in the current generations (while eliminating the individuals with low fitness values).

The crossover operation typically proceeds as follows. First, the individuals in the generation $(t-1)$ are sorted in decreasing order of their fitness values (whereby the individuals from all three groups, i.e., elite group, reproduction group, and random group, are considered). A mating pool is formed from the first $P_p(t)$ individuals in the ordering. Parts of the chromosomes of the individuals in the mating pool are then exchanged (swapped) with a fixed crossover probability. For our network optimization, we chose to swap their M values because we have observed that M (with D, F, and p fixed) tends to explore potential solutions in the vicinity of the parents (as is also evidenced by the tables in Appendix E, which are discussed in detail in Section 7.4). Specifically, for our network optimization, the first $P_p(t)$ individuals in the ordering, i.e., the mating pool, are processed as follows. We take the first two individuals in the ordering. With the crossover probability (which we fix at the typical value 0.9), we swap their M values, i.e., we put the M value of the first individual (in the ordering) in place of the M value of the second individual, and vice versa. The other three decision values, D, F, and p, in the individuals' chromosomes remain unchanged. (Note that in our problem the swapping of M while keeping D, F, and p in place may result in a chromosome that violates the constraint $M \leq F$. If this situation arises, we discard the violating M value and randomly draw a new M from a uniform distribution over $[1, F]$.) With the complementary crossover probability (0.1), the chromosomes of the two individuals remain unchanged. The two individuals (irrespective of whether their chromosomes were swapped or not) then become members of the reproduction group in the next generation. We then move on to the third and fourth individuals in the ordering, and swap their M values with probability 0.9, move them to the reproduction group in the next generation, and so on. We note that the elite group of the next generation is formed from the

Pareto–optimal individuals in the current generation, irrespective of whether these individuals are in the mating pool of the current generation. (An individual may appear twice in the next generation if it is Pareto–optimal in the current generation and participates in the crossover operation without having the M value changed. Only one copy of such a "duplicate" individual is processed in the next generation, the other copy is discarded.)

Mutation operation

The mutation operation modifies individuals by changing small parts in their chromosomes with a given (typically small) mutation probability. Generally, bitwise operation is used. A given individual is selected with the mutation probability for the mutation operation. If an individual is selected for the mutation, then one bit in the individual's chromosome is flipped. The location of the bit is drawn randomly from a uniform distribution over the length of the chromosome. The mutation operation keeps diversity in the population and prevent GAs from getting stuck in local optima. In our network optimization, we implement the mutation operation by randomly drawing an M value from a uniform distribution over $[1, F]$. All individuals in the elite group, the reproduction group, and the random group are mutated with a mutation probability of 0.05 (a typical value). We note that we do not use bitwise mutation, because bitwise mutation would frequently produce offspring that are distant from the parents. We chose to randomly mutate M in the interval $[1, F]$, as this operation does not result in constraint violations, yet keeps the population sufficiently diverse.

After the mutation operation, we evaluate the average throughput and mean delay achieved by the individuals (in all three groups, i.e., elite group, reproduction group, and random group) in the new generation and start the next evolution cycle; as illustrated in Fig. 7.4. In this new evolution cycle, we select again the individuals with the largest fitness values for the crossover operation, which gives the reproduction group of the next generation. We also determine again the Pareto–optimal individuals to form the elite group in the next generation.

7.4 Results

In this section, we employ the genetic algorithm based approach developed in the preceding section to optimize the AWG–based single–hop WDM network. We determine the settings of the network architecture parameter D and the protocol parameters F, M, and p that give Pareto–optimal throughput–delay performance. We use the random weight genetic algorithm (RWGA) with elitism with the parameter settings found in the preceding section, i.e., a population size of $P = 200$, $G = 40$ generations, crossover probability 0.9, and mutation probability 0.05. Data packets can have one of two lengths. A data packet is F slots long with probability q, and $K = (F - M)$ slots long with probability $(1 - q)$. To reasonably limit the search space we restrict F to be no larger than 400 slots. The number of nodes in the network is set to $N = 200$ and the transceiver tuning range is fixed at $\Lambda = 8$ wavelengths.

In the first set of optimizations, we determine the Pareto–optimal performance for different (but fixed) combinations of traffic load σ and fraction of long packet traffic q. Specifically, we optimize the network for a light traffic scenario with $\sigma = 0.1$, a medium traffic scenario with $\sigma = 0.3$, and heavy load scenarios with $\sigma = 0.6$ and $\sigma = 0.8$. For each traffic load level, we consider the fractions $q = 0.1,\ 0.5$, and 0.9 of long packet traffic. In these optimizations we determine the free decision variables D, F, M, and p that give the Pareto–optimal solutions.

To put the optimizations for fixed σ and q in perspective, we also conduct an optimization

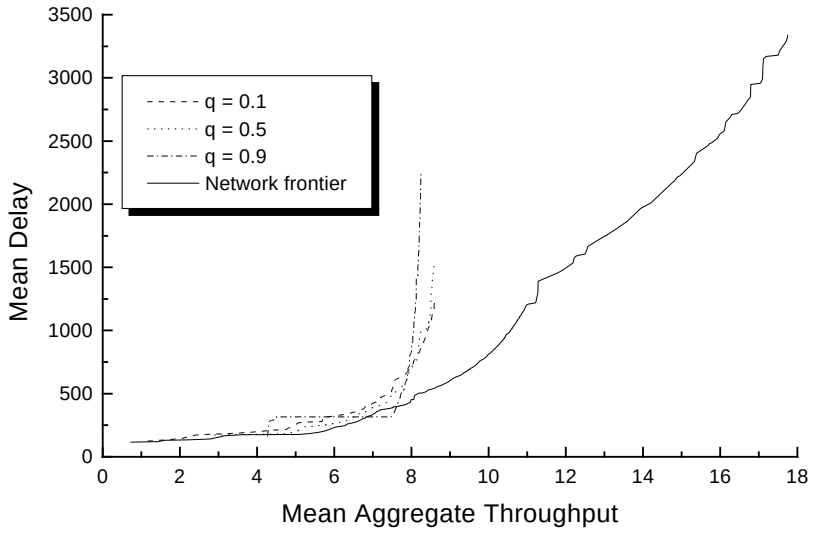

Figure 7.11: Efficient frontiers for light traffic load $\sigma = 0.1$ for different fractions q of long packet traffic and network frontier (with σ and q as free decision variables).

where the traffic load σ and the fraction q of long packet traffic are free decision variables (in addition to $D,\ F,\ M$, and p). This optimization gives the best achievable network performance, which we refer to as network frontier. Loosely speaking, the network frontier gives the Pareto–optimal performance when the network is "fed optimally" with traffic. (To find the network frontier, we exchange (swap) σ as well as M in the crossover operation and use a population size of $P = 400$ rather than $P = 200$ to accommodate the larger chromosome.) The detailed solutions for the network frontier are given in Table E.1 in Appendix E.

7.4.1 Pareto–optimal performance for light traffic load

Fig. 7.11 shows the Pareto–optimal throughput–delay frontier for a light traffic load of $\sigma = 0.1$ for $q = 0.1,\ 0.5$, and 0.9 (along with network frontier). Tables E.2, E.3, and E.4 in Appendix E give the individual Pareto–optimal solutions. The numbers of Pareto–optimal solutions with each $D = 2,\ 4$, and 8 are shown in Table 7.2. We observe from Fig. 7.11 that for a small fraction q of long packets the network is able to achieve relatively small delays (of less than 1500 slots) even for large throughputs (of 8 packets/frame and more). When the fraction q of long packet traffic is large, however, the smallest achievable delays become very large (up to 2250 slots) for large throughputs. This is because the considered network allows for the scheduling of at most $R\ (= \Lambda/D)$ long packets in a cycle (consisting of D frames) at each of the D AWG input ports. (There are also $(D-1) \cdot R$ transmission slots exclusively for short packets in a cycle at each AWG input port; in addition short packets can fill up the R long packet transmission slots.) With a larger fraction of long packets, the probability increases that a data packet fails in the scheduling and requires retransmission of the corresponding control packet, resulting in larger delays.

We also observe that the light traffic scenario is able to achieve the small delay (and small

	$\sigma = 0.1$			$\sigma = 0.3$			$\sigma = 0.6$			$\sigma = 0.8$		
q	0.1	0.5	0.9	0.1	0.5	0.9	0.1	0.5	0.9	0.1	0.5	0.9
$D = 2$	148	132	133	108	84	158	31	102	121	23	105	135
$D = 4$	0	1	8	2	65	4	86	46	5	102	46	3
$D = 8$	0	0	0	0	2	2	1	4	1	0	4	1
Total	148	133	141	110	151	164	118	152	127	125	155	139

Table 7.2: Number of Pareto–Optimal Solutions with $D = 2,\ 4$, and 8.

throughput) part of the network frontier. This is because a small number M of control slots is sufficient to ensure reasonably large success probabilities in the control packets contention when the probability σ of an idle node generating a new packet at the beginning of a cycle is small. The small M in turn allows for small frame length F, and thus short cycle length $D \cdot F$, which results in small delays.

We observe that there are some instances where the Pareto–optimal frontier for $q = 0.9$ dominates the network frontier, e.g., around a throughput of 7.7 packets/frame. This is due to the stochastic nature of the genetic algorithm, which finds a very close approximation of the true optimal frontier in a computationally efficient manner. By definition, the true network frontier can not be dominated by the true frontier for a fixed σ or q; finding these true frontiers, however, is computationally prohibitive.

We observe from Tables 7.2, E.2, E.3, and E.4 that for the considered light traffic load $\sigma = 0.1$, most of the Pareto–optimal solutions have $D = 2$. However, for a larger fraction q of long packet traffic the number of Pareto–optimal solutions with $D = 4$ increases. We observe from the Table E.4 that $D = 4$ is the best choice to achieve low delay service. This is because the long packets are more difficult to schedule and therefore tend to require more retransmissions of control packets, resulting in an increased mean delay. Recall that a control packet is discarded if the corresponding data packet cannot be scheduled. This makes the control packet contention a bottleneck when the packet scheduling becomes difficult. With larger D, fewer nodes $S = N/D$ contend for the M control slots available to them every Dth frame. This increases the probability of successful control slot contention, thus relieving the control packet contention bottleneck. Note that the control packet contention bottleneck could also be relieved by reducing the retransmission probability p. However, we see from the results in Tables E.3 and E.4 that this strategy is not selected (except in the 9th row of Table E.4 when the transition from $D = 4$ to $D = 2$ occurs). The reason for this is that the smaller p would result in a relative large increase in the mean delay, making it preferable to increase D and keep p large (the first eight rows of Table E.4).

Generally, we observe from Tables E.2, E.3, and E.4 that the Pareto–optimal solutions with larger throughput are achieved for larger F. The Pareto–optimal M values, on the other hand, remain in the range $30 - 60$ for $q = 0.1$ and $q = 0.5$ and are typically $30 - 80$ for $q = 0.9$, even for very large F. Upon close inspection we discover an interesting underlying trend in the F and M solutions as we move along the efficient frontier from small to large throughput values. The frame length F typically makes a jump to a new value (e.g., from $F = 44$ to 59 in the 4th row of Table E.2) and stays around the new value for a few solutions. For F (almost) fixed, several distinct Pareto–optimal solutions are obtained for decreasing M values (from $M = 49$ to 30 for F around 59 in Table E.2). Once F makes a jump (to values around 100 in line 20), M is reset to a larger value (of 50 in line 20). The explanation of this behavior is as follows.

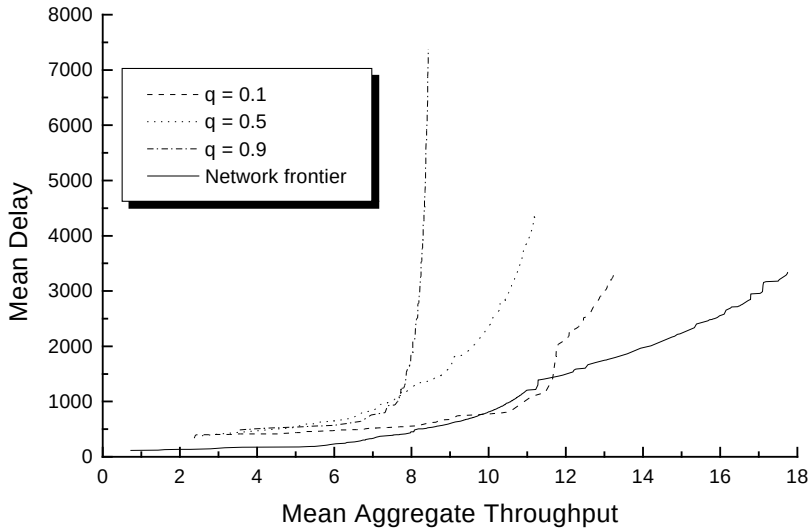

Figure 7.12: Efficient frontiers for medium traffic load $\sigma = 0.3$ for different fractions q of long packet traffic and network frontier (with σ and q as free decision variables).

For large M, the probability of successful control packet contention is large, and the probability of control packet retransmission is small, giving small delays. However, for large M, the length $K = (F - M)$ of a short packet is small, resulting in a small contribution of a short packet to the throughput (Eqn. (7.1)). Now as M decreases (for F fixed), control packet retransmission becomes more likely, increasing the mean delay, while the contribution of a short packet to the throughput increases. We also observe from the tables that for optimal network operation the retransmission probability p should be in the range from 0.75 to 1.0.

7.4.2 Pareto–optimal performance for medium traffic load

Fig. 7.12 shows the Pareto–optimal solutions for a medium traffic load of $\sigma = 0.3$. The numbers of Pareto–optimal solutions with $D = 2,\ 4$, and 8 are shown in Table 7.2 and the individual Pareto–optimal solutions are given in Tables E.5, E.6, and E.7 in Appendix E. We observe from Fig. 7.12 that the differences in performance for the different fractions q of long packet traffic are more pronounced for the larger traffic load $\sigma = 0.3$, compared to the light traffic load $\sigma = 0.1$ shown in Fig. 7.11. For $\sigma = 0.1$, the efficient frontiers for $q = 0.1$ and $q = 0.5$ roughly overlap and give both a smallest achievable delay of roughly 715 slots for a throughput of 8 packets/frame. For $\sigma = 0.3$, on the other hand, the efficient frontier for $q = 0.1$ clearly dominates, giving a smallest achievable delay of roughly 555 slots for a throughput of 8 packets/frame, whereas the corresponding smallest achievable delay for $q = 0.5$ is more than twice as large. This increasing gap in performance is again due to the fact that long packets are more difficult to schedule and thus tend to cause larger delays. The smaller delay of 555 slots for $\sigma = 0.3$, compared to 715 slots for $\sigma = 0.1$ is achievable because with the larger σ, the throughput level of 8 transmitting nodes per slot is reached with smaller sized packets (i.e., smaller F and smaller $K = (F - M)$), thus reducing the cycle length and in turn the delay. We observe from Tables E.5, E.6, and

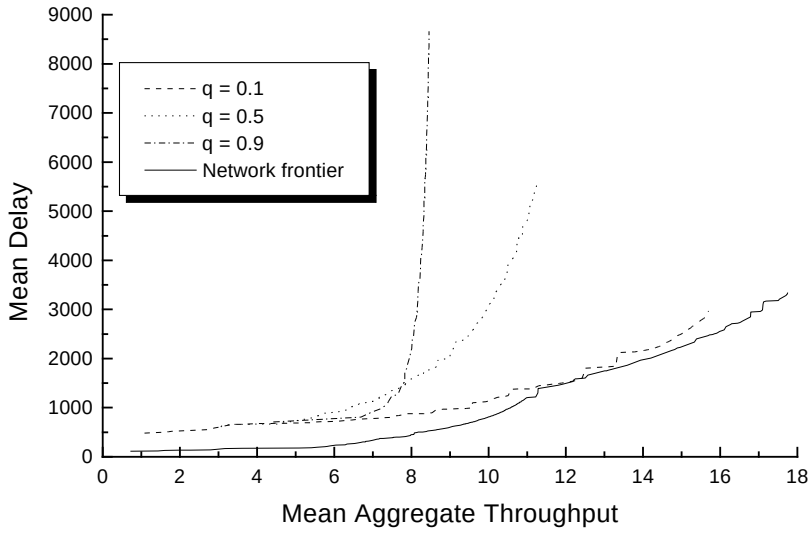

Figure 7.13: Efficient frontiers for heavy traffic load $\sigma = 0.6$ for different fractions q of long packet traffic and network frontier (with σ and q as free decision variables).

E.7 that small delays are again achieved for large D values. For $q = 0.5$ and $q = 0.9$, the first few Pareto–optimal solutions at the top of the tables have $D = 8$, then $D = 4$ is optimal as we go down the tables to larger delays. As in the case of $\sigma = 0.1$, this behavior is due to the control packet contention and data packet scheduling bottlenecks. From Table 7.2 we observe that there is no clear trend in the number of solutions with $D = 2$ and $D = 4$. This appears to be due to the stochastic nature of the genetic algorithm approach, which finds a large total number of solutions for $q = 0.5$, with many solutions being tightly spaced in the region where $D = 4$ is optimal. As before, larger throughput is optimally achieved for large F. The optimal settings of M are typically in the range from $60 - 80$. The optimal settings of p are mostly 0.95 for $q = 0.1$ and $q = 0.5$. For $q = 0.9$, the optimal p settings are typically 0.7. This smaller p setting for a medium load of predominantly long packet traffic is better as it somewhat abates the control packet contention bottleneck at the expenses of slightly larger delays, as discussed above.

7.4.3 Pareto–optimal performance for heavy traffic load

Figs. 7.13 and 7.14 show the Pareto–optimal solutions for a heavy traffic load of $\sigma = 0.6$ and $\sigma = 0.8$, respectively. The number of Pareto–optimal solutions with $D = 2,\ 4$, and 8 are given in Table 7.2. The complete parameter vectors corresponding to the Pareto–optimal solutions are given in Tables E.8 – E.13. We observe from the figures and the tables that both considered heavy load scenarios give similar results with the $\sigma = 0.8,\ q = 0.1$ scenario attaining the larger throughput region of the network frontier. We notice that with an increasing fraction q of long packet traffic, the number of Pareto–optimal solutions with $D = 2$ increases, while the number of solutions with $D = 4$ decreases. There are two primary effects at work here. On the one hand, a larger D allows for a larger throughput. To see this, note that the considered network

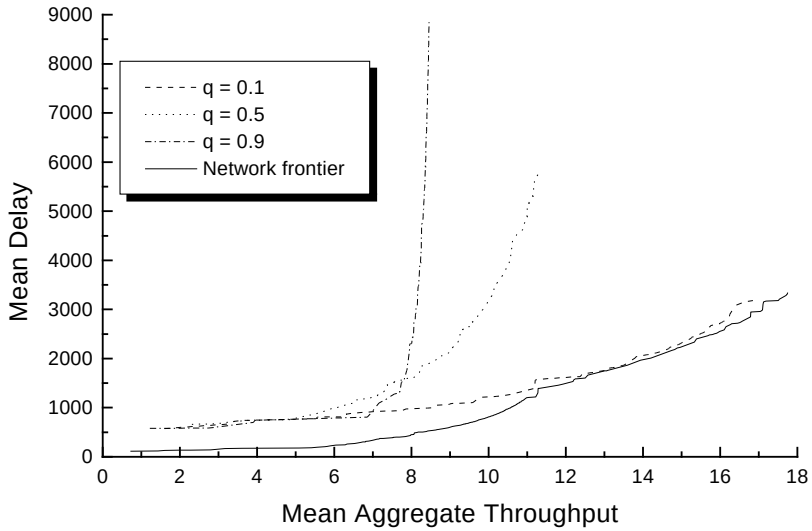

Figure 7.14: Efficient frontiers for heavy traffic load $\sigma = 0.8$ for different fractions q of long packet traffic and network frontier (with σ and q as free decision variables).

allows for the scheduling of at most R $(= \Lambda/D)$ long packets at each of the D AWG input ports within one cycle (consisting of D frames); for a total of at most $D \cdot R = \Lambda$ scheduled long packets per cycle in the entire network. The network also allows for the scheduling of at most $(D-1) \cdot R$ short packets at each of the D AWG input ports within one cycle; for a total of at most $D \cdot (D-1) \cdot R = \Lambda \cdot (D-1)$ scheduled short packets per cycle in the network (in addition short packets may take up long packet transmission slots). Thus, for a larger D the network allows for the scheduling of more short packets and thus for an overall larger throughput; this is a result of the spatial reuse of all Λ wavelengths at all D AWG ports. On the other hand, a larger D increases the delay in the network (provided the frame length F is constant). This is because a larger cycle length $D \cdot F$ increases the delay incurred by the control packet pretransmission coordination and retransmissions, which operate on a cycle basis. These throughput and delay effects combine to make $D = 2$ the better choice when long packets dominate (i.e., when q is large), since short packets make only a small contribution to the throughput. We also observe from Tables E.8 and E.11 that even when q is small, $D = 2$ is a good choice for delay–sensitive traffic. Although we see that some Pareto–optimal solutions in the small delay range have $D = 4$. This indicates that both a 2×2 AWG and a 4×4 AWG based network can achieve small delays for traffic consisting mostly of short packets, provided the protocol parameters F, M, and p are set properly. On the other hand, only a 4×4 AWG based network achieves the large throughputs on the efficient frontier for small q (i.e., predominantly short packet traffic). As before, we observe that the Pareto–optimal solutions with larger throughput values have larger frame lengths F. Also, as before, the Pareto–optimal solutions have typically between $M = 60$ and 110 control slots per frame. We note, however, some differences in the optimal setting of the retransmission probability p in this heavy traffic load scenario compared to the light/medium load scenario. As before for $q = 0.1$ the optimal p setting is typically in the range of $0.9 - 1.0$.

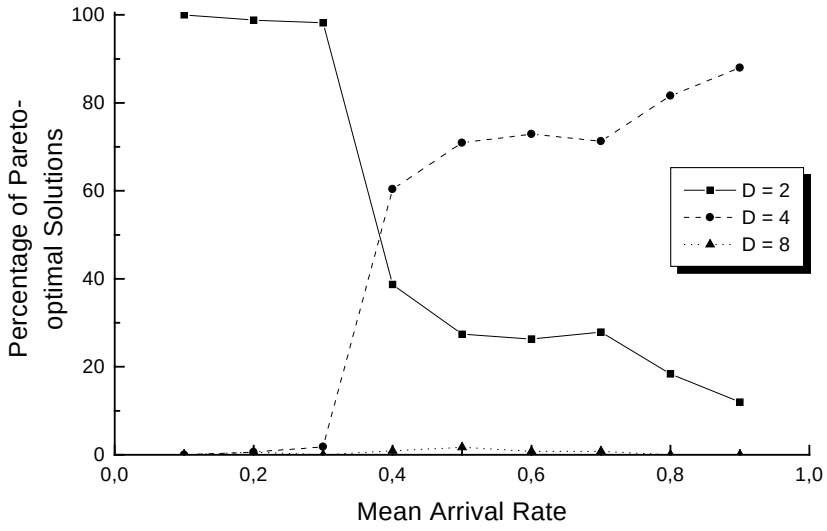

Figure 7.15: Percentage of Pareto–optimal solutions with $D = 2,\ 4$, and 8 as a function of the traffic load σ with $q = 0.1$.

For $q = 0.5$ and $q = 0.9$, however, the optimal p is now typically in the range from 0.6 to 0.95.

7.4.4 Pareto–optimal architecture planning

We now study the proper setting of the AWG degree D in detail. The setting of this network architecture (hardware) parameter has a profound impact on the network performance, as the results discussed so far illustrate. Importantly, once the network hardware for a particular D value has been installed, it is very difficult and costly to change D; whereas the protocol parameters $F,\ M$, and p can relatively easily be changed by modifying the network protocol (software). For this reason, the proper setting of D warrants special attention. We have observed so far that for predominantly long packet traffic (i.e., large q), $D = 2$ is the best choice for all levels of traffic load σ. For predominantly short packet traffic (i.e., small q), on the other hand, the choice is not so clear. For light traffic loads, $D = 2$ is the best choice, whereas for heavy traffic loads, $D = 4$ turns out to be the best choice.

To explore the optimal setting of D as a function of the traffic load σ, we plot in Figs. 7.15 and 7.16 the percentage of Pareto–optimal solutions with $D = 2,\ 4$, and 8 for $q = 0.1$ and $q = 0.9$, respectively. We observe from Fig. 7.15 that for σ less than 0.4, most Pareto–optimal solutions have $D = 2$, whereas for σ larger than 0.4, most Pareto–optimal solutions have $D = 4$. The explanation of this behavior is as follows. For light traffic loads, $D = 2$ is preferred as it achieves smaller delays while at the same time providing sufficient resources for control packet contention and data packet scheduling. (Recall that $S = N/D$ nodes at an AWG input port content for the M control slots available to them in one frame (out of the D frames in a cycle), and that spatial wavelength reuse provides for $\Lambda \cdot (D-1)$ transmission slots for short packets.) As the traffic load increases, however, the control packet contention and data packet scheduling become increasingly bottlenecks which are relieved for larger D.

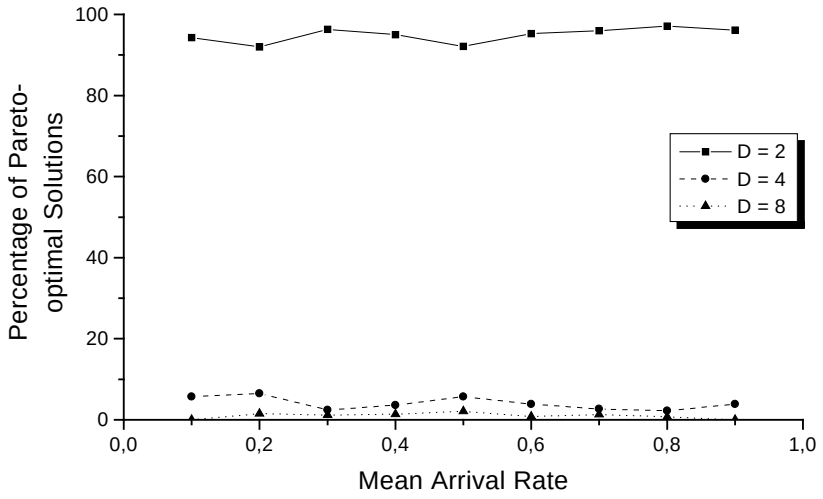

Figure 7.16: Percentage of Pareto–optimal solutions with $D = 2,\ 4$, and 8 as a function of the traffic load σ with $q = 0.9$.

7.4.5 Pareto–optimal network operation

Next, we fix the AWG degree D at $D = 2$ and $D = 4$, and allow only the protocol parameters $F,\ M$, and p to vary (i.e., only $F,\ M$, and p are decision variables, D is fixed). We employ our genetic algorithm based approach to obtain the efficient throughput–delay frontiers in these settings; we refer to these efficient frontiers as the 2×2 network frontier and the 4×4 network frontier, respectively. We compare the thus obtained efficient frontiers with the efficient frontier obtained when both the hardware parameter D and the software parameters $F,\ M$, and p are decision variables, which we refer to as optimal frontier. We compare the 2×2 frontier and the 4×4 frontier with the optimal frontier in Fig. 7.17 $(a) - (h)$. The corresponding Pareto–optimal solutions are tabulated in Tables E.14 – E.29. A number of observations can be made. First, as expected the 2×2 frontier approximately coincides with the optimal frontier for light to medium loads of predominantly short packet traffic, and all load levels of predominantly long packet traffic. For heavy loads of predominantly short packet traffic, on the other hand, the 4×4 network frontier achieves the optimal frontier, as we expect from our earlier results. We also observe that there are some instances where the optimal frontier is dominated by the 2×2 network frontier or the 4×4 network frontier, e.g., in Fig. 7.17 (c) around a throughput of 11.5 packets/frame. These instances are again due to the stochastic nature of the employed genetic algorithms. By definition, the 2×2 network frontier and the 4×4 network frontier can not dominate the true optimal frontier, which however could only be found by a computationally prohibitive exhaustive search. The genetic algorithm approach finds a very close approximation of the true optimal frontier in a computationally efficient manner.

Figs. 7.17 $(a) - (h)$ give also a number of surprising results, which we would not expect, based on our earlier observations. First, the 4×4 network is able to come close to the optimal frontier for medium and heavy loads of predominantly long packet traffic, which is a surprise

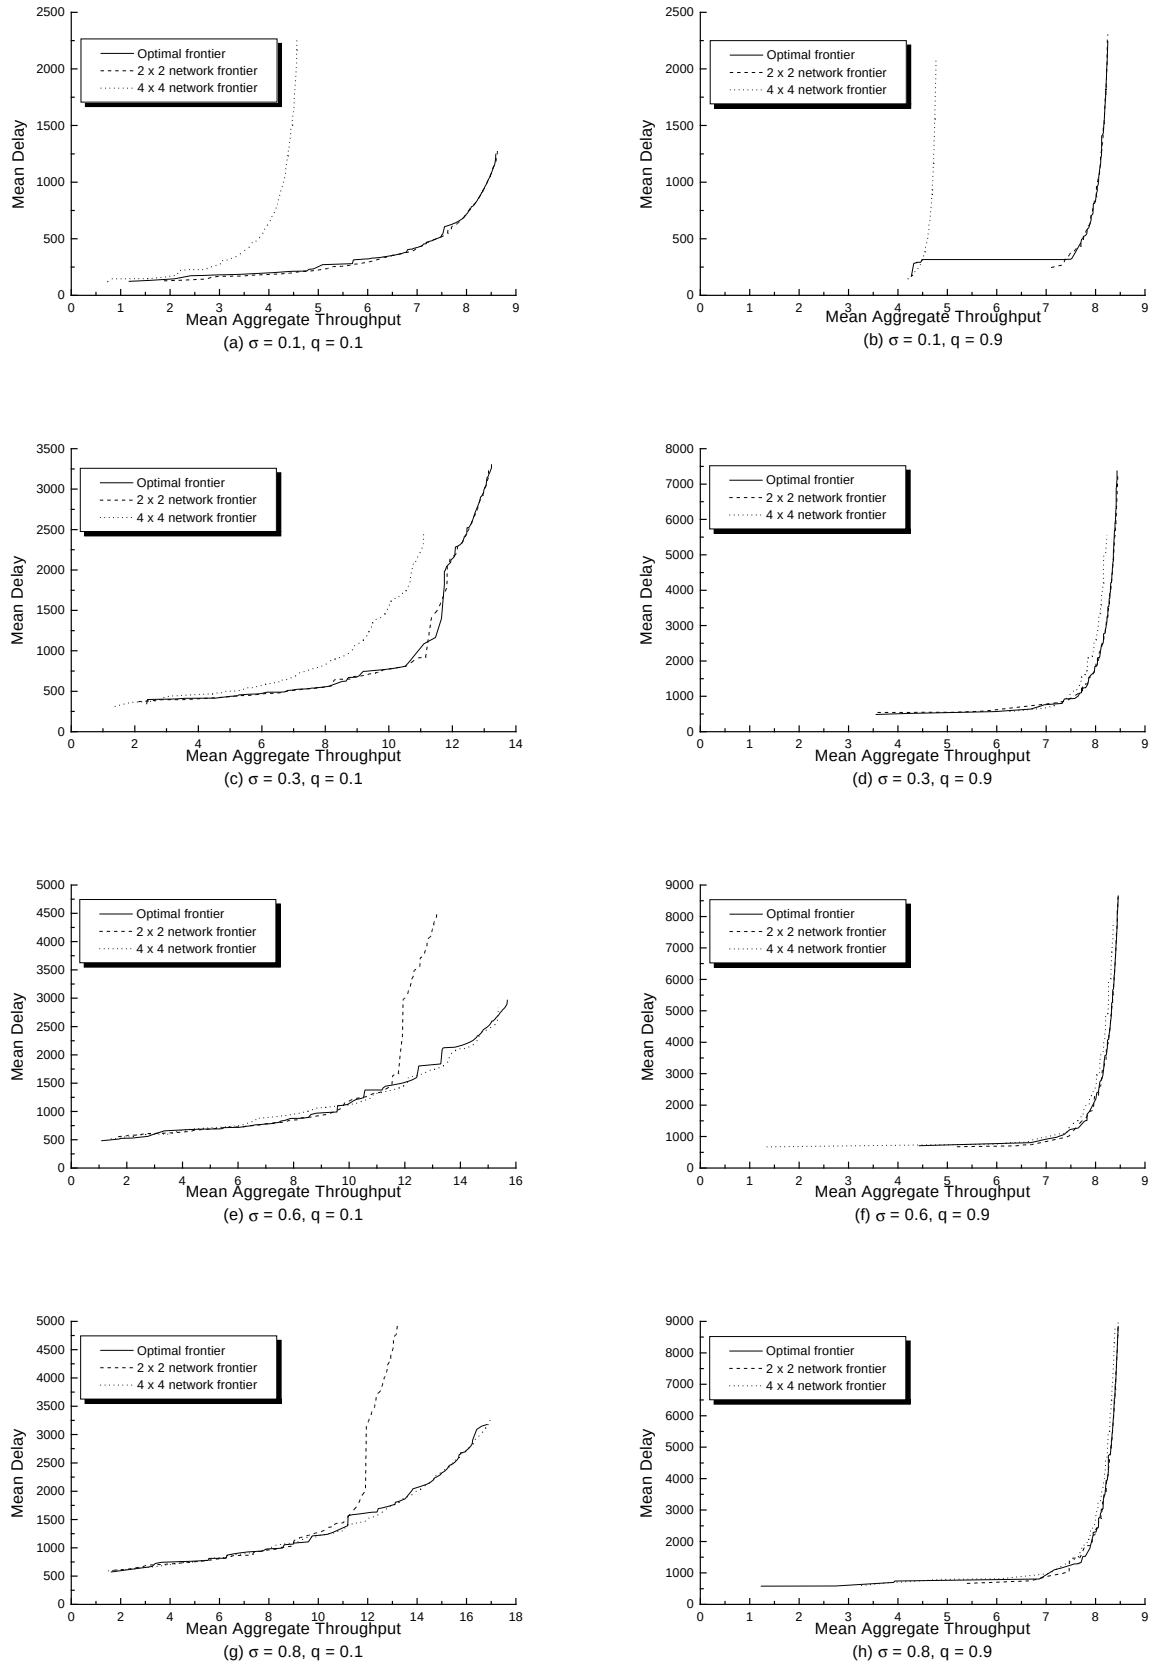

Figure 7.17: Optimal frontier (with D a free decision variable), 2×2 network frontier (with $D = 2$, fixed), and 4×4 network frontier (with $D = 4$, fixed) for different (fixed) traffic loads σ and fractions q of long packet traffic.

given the results in Table 7.2 and Fig. 7.16. The 4×4 network achieves this by properly tuning its three protocol parameters, F, M, and p, as detailed in Tables E.17, E.21, E.25, and E.29. Overall, the 4×4 network shows some flexibility in achieving good performance close to the optimal frontier for medium to heavy loads of both short and long packet traffic by properly tuning the protocol parameters (in software). For light traffic loads, however, the 4×4 network is not able to come close to the optimal frontier. The 2×2 network, on the other hand, appears to be more flexible than the 4×4 network. By properly tuning its protocol parameters, the 2×2 network is able to come fairly close to the optimal frontier even for heavy loads of short packet traffic (see Figs. 7.17 (e) and (g)). Overall, our results indicate that the 2×2 network is the best choice for achieving efficient multiservice convergence in a metro WDM network. The 2×2 network frontier approximately coincides with the optimal frontier for all load levels of long packet traffic and for light to medium loads of short packet traffic. For heavy loads of short packet traffic, the 4×4 network attains the optimal frontier. But the 2×2 network is able to come fairly close to the optimal frontier, simply by adjusting its protocol parameters in software.

7.5 Conclusions

We have discussed the dimensioning of our AWG based WDM network architecture and MAC protocol. In order to efficiently provide multiservice support the network has to be dynamically optimized in terms of throughput and delay according to varying traffic demands. This gives rise to a multiobjective optimization problem. Since in our case throughput and delay are conflicting objective functions the outcome of the optimization is a set of Pareto–optimal solutions. They are listed in Appendix E for a large number of different traffic loads and data packet size distributions. The optimal throughput–delay trade–off curve (also known as efficient frontier) allows for optimal network planning and reconfiguration according to varying traffic demands such that network resources are used efficiently and costs are reduced.

We have obtained the Pareto–optimal solutions by deploying computationally efficient genetic algorithms (GAs). By means of numerical experiments we have compared three commonly used GAs and decided to use RWGA. Subsequently, we have determined the proper setting of the RWGA parameters for an improved search and optimization. We find that a network based on a 2×2 AWG is most flexible in efficiently providing different transport services under a wide range of traffic loads and packet size distributions. In addition, using an AWG with a smaller degree reduces the number of required EDFAs which have to be placed at each AWG input and output port in order to compensate for the splitting loss induced by the combiners and splitters. This results in significant cost savings which is an important issue in cost–sensitive metro WDM networks. For a fixed network hardware the different transport services are achieved by optimally tuning the MAC protocol parameters (software) according to the found Pareto–optimal solutions. In particular, small frame lengths in the timing structure of the AWG network's MAC protocol give Pareto–optimal performance with small delay (and moderate throughput), while large frame lengths achieve optimal performance with large throughput (and moderate delays). The optimal number of control packet contention slots per frame is typically in the range from $30 - 80$. The optimal control packet retransmission probabilities are close to one for light traffic loads and in the range from $0.6 - 0.75$ for heavy loads.

Chapter 8

Feasibility Issues

After specifying, evaluating, and optimizing our network architecture and protocol in the previous chapters, we now investigate the feasibility of the network. For the signalling of reservation requests we adopt a previously reported realization of spectrum spreading. In Section 8.1, we examine the impact of the transmission impairments of Section 2.2 on the transmission limitations of the network. In Section 8.2, we investigate by means of simulations the throughput–delay and packet loss performance of the network using the aforementioned signalling approach. Unlike the analyses and supplementary simulations of Chapter 6, we do not assume Bernoulli traffic in the following simulations. Instead, we use packet header trace files which were recorded at the metro level. The obtained results are discussed in Section 8.3. Section 8.4 provides some conclusions.

8.1 Transmission limitations

A preliminary study on the transmission limitations of an AWG based single–hop WDM network in terms of transmission distance and AWG degree was conducted in [YOH97]. The simulation study takes fiber nonlinearities (FWM, SRS), ASE noise accumulation, crosstalk, and attenuation into account, but ignores the signal distortion caused by chromatic dispersion. Two types of single–mode fiber are considered: A conventional fiber with a dispersion of 17 ps/nm·km and a dispersion–shifted fiber (DSF) with a dispersion of 2 ps/nm·km. It was shown that the maximum transmission distance can reach up to 1000 km for a system based on a 400×400 AWG with 400 wavelength channels used at each AWG port at 2.5 Gb/s if conventional fiber is used with a 50 km amplifier spacing, an AWG crosstalk better than -55 dB, and a channel spacing of $\sim$10 GHz. The requirement of better than -55 dB AWG crosstalk with 400 ports and $\sim$10 GHz channel spacing should be feasible with bulk optics technology. The transmission performance of an AWG based single–hop network was investigated more in detail in [BNS99]. In addition to the aforementioned impairments also fiber chromatic dispersion, noise contributions from the receiver, SPM, and XPM were taken into account. It was shown that when the fiber chromatic dispersion is 17 ps/nm·km, the system performance is not acceptable (BER $\approx 10^{-2}$) due to the distortion induced by chromatic dispersion. While if the fiber chromatic dispersion is 2 ps/nm·km, the transmission of 400 wavelength channels via a 400×400 AWG is possible at a BER $< 10^{-9}$.

To verify the transmission performance of our AWG based metro WDM network we have run some numeric simulations and developed an analytical model in [Her02]. Due to the large computational overhead of the simulation we consider only $N = D \cdot S = 64$ nodes connected by

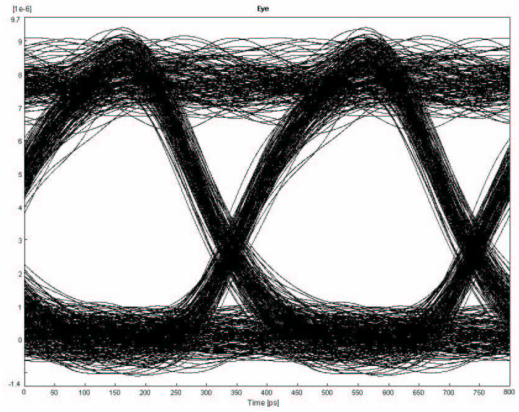

Figure 8.1: Eye diagram at 2.5 Gb/s for a network diameter of 175 km with chromatic dispersion and nonlinearities.

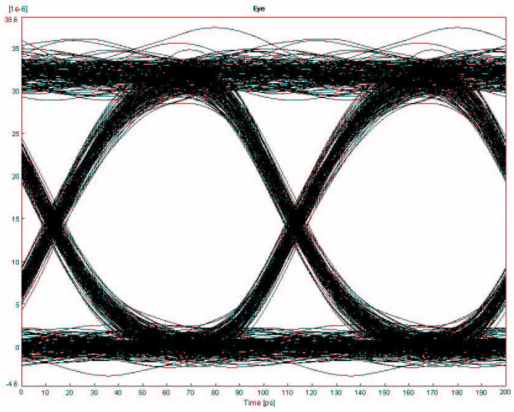

Figure 8.2: Eye diagram at 10 Gb/s for a network diameter of 200 km without chromatic dispersion and nonlinearities.

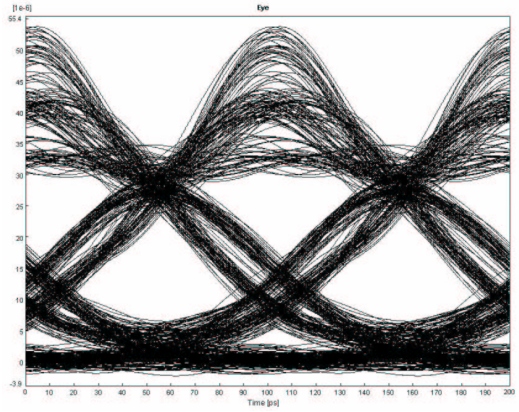

Figure 8.3: Eye diagram at 10 Gb/s for a network diameter of 200 km with chromatic dispersion, but without nonlinearities.

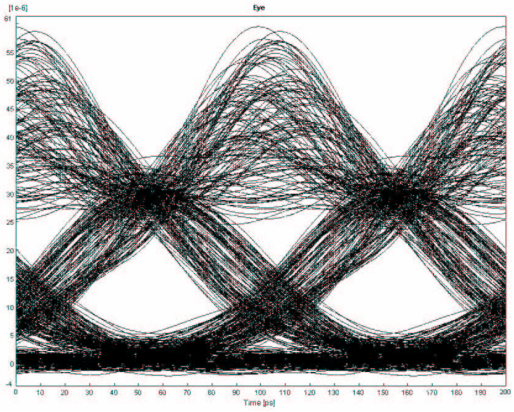

Figure 8.4: Eye diagram at 10 Gb/s for a network diameter of 200 km with chromatic dispersion and nonlinearities.

a $D \times D$ AWG with S nodes attached to each AWG port via a pair of combiner and splitter.

8.1.1 Eye diagrams

Let us first focus on data transmission and some eye diagrams measured at the receiver output without providing any dispersion compensation. To compensate for the combining, splitting, and fiber loss an Erbium doped fiber amplifier (EDFA) is used at each combiner output port and at each splitter input port. Hence, in a $D \times D$ AWG based network $2D$ EDFAs are deployed. The channel spacing is set to 100 GHz, $D = 4$, and $S = 16$. We deploy a nonreturn–to–zero (NRZ) intensity modulation (IM) with direct detection (DD). Fig. 8.1 shows the eye diagram at a line rate of 2.5 Gb/s and a network diameter of 175 km for a BER $< 10^{-9}$. The laser is directly modulated. By using the MetroCor fiber with a chromatic dispersion of -8 ps/nm·km no extra dispersion compensation is needed. We observe from Fig. 8.1 that the eyes are wide open.

The eye diagrams at a line rate of 10 Gb/s and a network diameter of 200 km for a BER $< 10^{-9}$ are depicted in Figs. 8.2 through 8.4. The figures illustrate the impact of chromatic dispersion and nonlinearities on the transmission performance of the network. In this 10 Gb/s system each laser diode is externally modulated. We use the TeraLight fiber with a chromatic dispersion of 8 ps/nm·km. Dispersion compensation is not provided. Fig. 8.2 shows that the eye openings are very large if neither chromatic dispersion nor nonlinearities are taken into account. The eye openings decrease if chromatic dispersion is accounted for, as shown in Fig. 8.3. Fig. 8.4 depicts that nonlinearities further deteriorate the eye diagram.

8.1.2 Q factor

So far, we have considered only data transmission without paying attention to the control signal. To evaluate the transmission performance in presence of the control signal we have to fix ideas about the control rate, the spreading factor, i.e., the number of chips per control bit, and the type of broadband light source. To verify the feasibility of our network we adopt the realization proposal of sending spread control by using a cost–effective off–the–shelf light emitting diode (LED) as broadband light source [GGHR98][GGR99]. Accordingly, the control rate is set to ∼9.8 kb/s, the chip length is equal to 1023 resulting in a chip rate of 10 Mchip/s, at which the LED is modulated. The data rate is set to 2.5 Gb/s. As performance measure we use the so–called Q factor which is given by

$$Q = \frac{I_1 - I_0}{i_1 + i_0}, \tag{8.1}$$

where $I_1 - I_0$ is the excess of average current available at the receiver for distinguishing bit 1 from bit 0 and the sum $i_1 + i_0$ is the root mean square (RMS) value of current induced by noise at both the 1 and 0 electric levels [MS00]. Note that the BER can be easily obtained by

$$\text{BER} = \frac{1}{2}\, erfc\left(\frac{Q}{\sqrt{2}}\right), \tag{8.2}$$

where $erfc(\cdot)$ denotes the complementary error function.

Fig. 8.5 depicts the parameter Q as a function of network diameter for $N = 64$ and different combiner/splitter degree $S \in \{8, 16, 32\}$, i.e., $D \in \{8, 4, 2\}$. Apparently, Q decreases with increasing network diameter. To achieve an acceptable transmission performance it must hold $Q > 6$ which according to Eqn.(8.2) translates into a BER $< 10^{-9}$. We observe that for $S = 8, 16, 32$ the network diameter must not be larger than $200, 175, 150$ km, respectively, in order to guarantee a BER better than 10^{-9} for both control and data.

Fig. 8.6 depicts the maximum possible combiner/splitter degree S_{max} vs. the network diameter for $D = 2$ while providing $Q = 6$, i.e., BER $= 10^{-9}$. Clearly, with increasing network diameter the maximum number of nodes attached to the same AWG port decreases. The figure illustrates that the network is able to accomodate significantly more than $N = 64$ nodes within metropolitan regions. Fig. 8.6 and also Fig. 8.5 show that our AWG based single–hop WDM network is able to cover typical metropolitan areas with a diameter in the range of $75 - 200$ km at an acceptable BER performance.

8.2 Packet traces

The applied signalling approach offers a relatively small control rate of approximately 10 kb/s. Given that data packets are transmitted at significantly higher rates of a few Gb/s such a control

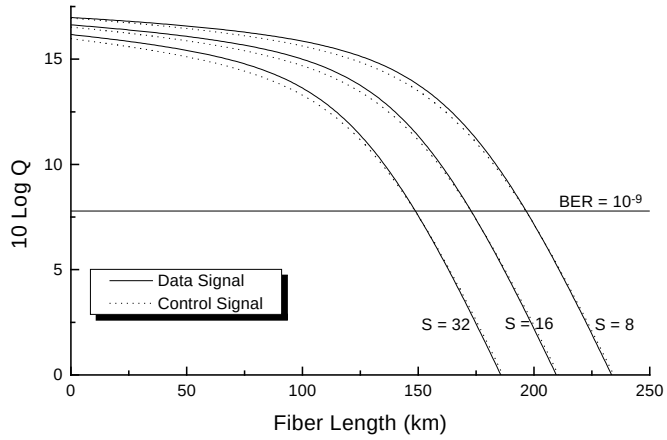

Figure 8.5: Parameter Q vs. network diameter for $N = 64$ and different combiner/splitter degree $S \in \{8, 16, 32\}$.

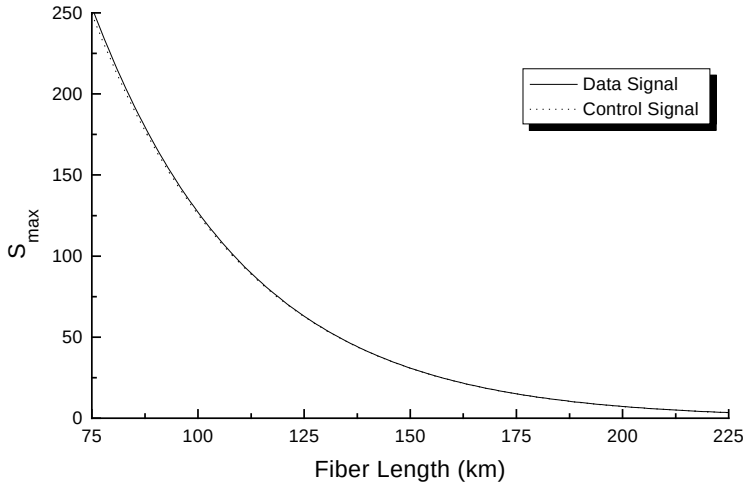

Figure 8.6: Maximum posssible combiner/splitter degree S vs. network diameter for $Q = 6$.

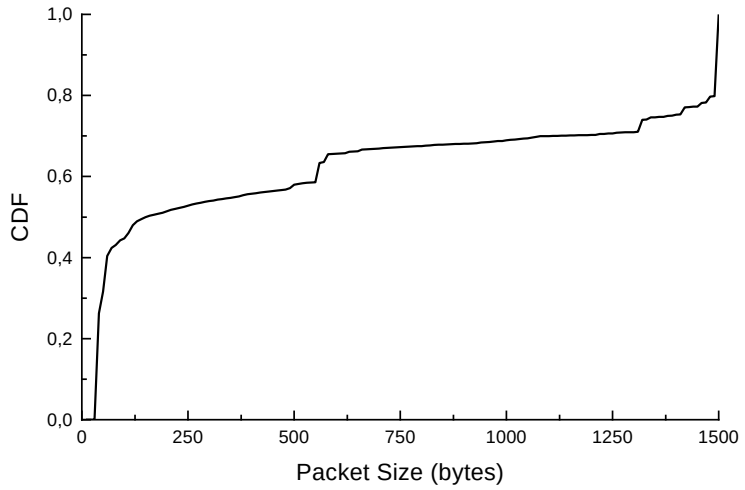

Figure 8.7: Cumulative distribution function (CDF) of packet size (bytes).

rate is not large enough for efficiently making reservations on a per–packet basis, resulting in a severe signalling bottleneck. To alleviate this, we introduce the concepts of *wormhole scheduling* [GYZ94] and *gated service* [LGK95]. The basic idea of wormhole scheduling is to schedule multiple successive packets (worm) which are destined to the same receiver at one time, i.e., one single control packet tries to make a reservation for several data packets rather than a single data packet. This implies that each node stores arriving data packets according to their destination, which is typically done in virtual output queues (VOQs) in the electronic domain at each source node; there is one VOQ for each destination. Gated service is used to determine the length of the packet worm. Recall that the reservation service in our AWG based network is gated, i.e., a given node is allowed to send control packets only in one frame per cycle. At the time the given node makes a reservation the corresponding control packet counts for all data packets that have arrived in the respective VOQ before the node sends the control packet. (The choice of the VOQ depends on the used selection/arbiter scheme.)

To see whether wormhole scheduling in conjunction with gated service is able to compensate for the slow control channel we evaluate the throughput–delay and packet loss performance of the network by means of simulation. Packet arrivals at each node are driven by a packet header trace file whose cumulative distribution function (CDF) of packet size in bytes is shown in Fig. 8.7. This AIX trace file was recorded 17 July 2002 at 11:25 am at the metro level and a data rate of 622 Mb/s (OC12) [AIX]. In the following simulations we scaled the traces up to 2.5 Gb/s (OC48) by dividing each packet time stamp by 4 similar to [XQ01]. The control rate is set to 10 kb/s. We set the backoff limit $b = 4$ and the scheduling window size $W = 2$ frames, i.e., one cycle. Due to the very long simulation run times we consider only $N = 8$ nodes interconnected by a 2×2 AWG ($D = 2$, $S = 4$). The number of used FSRs of the AWG $R = 8$. Each node is equipped with seven VOQs, one for each destination. Each VOQ is eight slots long, i.e., each VOQ can store packets whose transmission time takes up to eight slots. We assume uniform traffic, i.e., a given data packet is equally destined to all nodes except the source node. The length of a control packet is equal to eight bits which is sufficiently long to accomodate the destination addresses of 8 nodes and other fields. The network diameter equals 150 km which

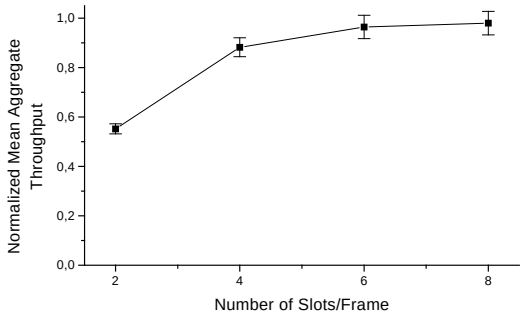

Figure 8.8: Normalized mean aggregate throughput vs. number of slots F per frame.

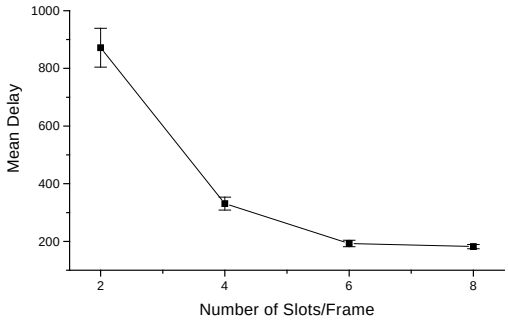

Figure 8.9: Mean delay (slots) vs. number of slots F per frame.

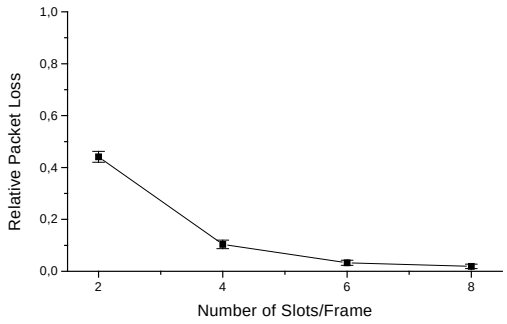

Figure 8.10: Relative packet loss vs. number of slots F per frame.

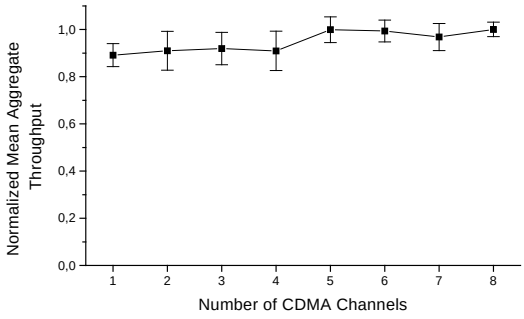

Figure 8.11: Normalized mean aggregate throughput vs. number of CDMA channels ($F = 8$ slots, fixed).

translates into an end–to–end propagation delay of approximately one slot. Each simulation was run for 10^4 slots including a warmup phase of 10^3 slots. Using the method of batch means we also determined the 95% confidence intervals.

We consider two different VOQ selection schemes. In Figs. 8.8 through 8.10 a given node ready to send a control packet selects the VOQ with the largest occupancy. If the reservation is successful the data aggregate is transmitted and the node proceeds to make a reservation for the second longest VOQ. Otherwise, the control packet is retransmitted according to the backoff mechanism. To increase the efficiency of wormhole scheduling the control packet takes also data packets into account which have arrived to the given VOQ since the last reservation attempt. Thus, the length of the data aggregate is dynamically increased, provided that new data packets have arrived and the length of the data aggregate does not exceed a given limit. Figs. 8.8 through 8.10 illustrate the normalized mean aggregate throughput, mean delay, and relative packet loss for different maximum size of data aggregates which is given by the frame length F. In these figures we set the number of reservation slots per frame to $M = F/2$. In Fig. 8.8 we have normalized the mean aggregate throughput by the offered network load. We observe that for increasing F the normalized mean aggregate throughput increases and asymptotically

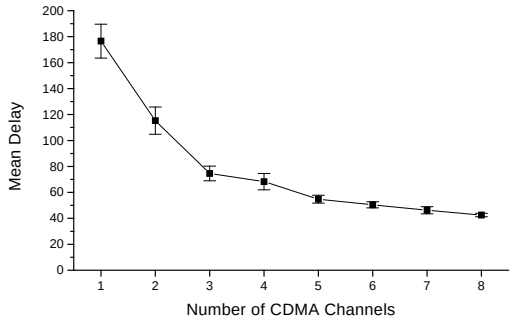

Figure 8.12: Mean delay (slots) vs. number of CDMA channels ($F = 8$ slots, fixed).

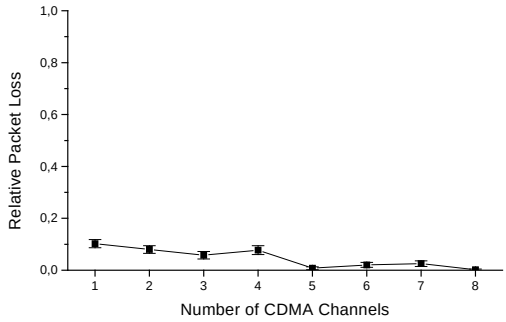

Figure 8.13: Relative packet loss vs. number of CDMA channels ($F = 8$ slots, fixed).

approaches 100%. The mean delay and packet loss performance are reduced for increasing F, as shown in Figs. 8.9 and 8.10, respectively. These results show that wormhole scheduling is able to achieve a very good performance if F is chosen sufficiently large. A large frame size F allows for large data aggregates and also provides enough reservation slots in each frame. Moreover, with a large F the time period between two successive reservation phases for a given node increases which allows nodes to collect more newly arriving data packets for forming larger data aggregates.

In Figs. 8.11 through 8.13 we set $F = 8$ and $M = F/2 = 4$. Moreover, we deploy a slightly modified VOQ selection scheme and apply up to eight different CDMA codes. CDMA enables the simultaneous transmission of multiple control packets simultaneously without collision. A given node ready to send a control packet for a data aggregate randomly selects one of the remaining CDMA codes which it has not yet used for other VOQs and transmits the corresponding control packet in a randomly chosen reservation slot. Note that this approach guarantees that a given node uses each CDMA code not more than once. However, it can happen that a given node sends more than one control packet within the same reservation slot. In this case, the multiple spread control packets are superimposed in the electronical domain before modulating the LED and do not collide since each control packet makes use of a different CDMA code. If two or more nodes simultaneously use the same CDMA code within the same reservation slot the corresponding control packets collide and are discarded by the receivers. The VOQ selection scheme is modified in that a given node ready to send a control packet always chooses the longest VOQ no matter whether previous reservation attempts have been successful or not. If a given node sticks with the same VOQ, i.e., no other VOQ has become longer than the considered one in the meantime, it resends the control packet using the same CDMA code and backing off as described in Section 6.4.1. If, however, in the meantime another VOQ has become larger than the original one the given node stops trying to make a reservation for the latter one and proceeds to the longest VOQ by randomly choosing a CDMA code and transmitting a control packet with reset backoff. Comparing Figs. 8.8 through 8.10 and 8.11 through 8.13, respectively, shows that for $F = 8$ and one single CDMA code both VOQ selection schemes yield a similar throughput–delay–loss performance. Figs. 8.11 and 8.13 illustrate that deploying more than one CDMA code improves the normalized mean aggregate throughput and relative packet loss. In particular, the mean delay is significantly decreased by using CDMA, as depicted in Fig. 8.12. This is because

each CDMA code creates a separate control channel. Control packets are equally distributed among the resultant control channels, resulting in more successfully transmitted control packets. Thus, more data aggregates can be successfully scheduled, the arriving data packets have to be buffered in the VOQs for a shorter period of time, and the corresponding control packets have to be retransmitted fewer times, resulting in a significantly decreased mean delay. Note that with CDMA more data aggregates are sent but each aggregate is shorter in length due to the smaller aggregation time compared to the case where only one spreading code is used. As a consequence, the normalized mean aggregate throughput and relative packet loss are improved only slightly.

8.3 Discussion

In the adopted signalling approach the control rate of 10 kb/s is too small for making reservations on a per–packet basis. To compensate for this signalling bottleneck we have deployed wormhole scheduling where each control packet tries to make a reservation for data aggregates. The concept of wormhole scheduling, however, is not only reasonable in single–hop WDM networks with a slow control channel. In the following chapter, we will address the survivability of our network where the control traffic is transmitted over a separate PSC at a line rate equal to that of data. With such a high control rate, scheduling on a per–packet basis would be feasible. However, it was shown in [GYZ94] that despite the fast control channel wormhole scheduling achieves a significantly improved throughput–delay performance than its per–packet counterpart. Actually, independent of the control channel rate, packet switching in slotted single–hop WDM networks can be done only very inefficiently due to chromatic dispersion [SH93]. To see this, note that in single–hop WDM networks each node's transmitter and/or receiver have to be tunable. Due to chromatic dispersion wavelengths travel at different velocities. Now, suppose that a given node is currently listening to a "slow" wavelength but has to receive a data packet on a "fast" wavelength in the subsequent slot. Due to the different travel velocities the packets on the two wavelengths are shifted in time. To compensate for this each slot has to be enlarged by an additional padding field either at the beginning or end of the slot such that the "fast" data packet is delayed sufficiently. Through the padding field in each slot the channel utilization is decreased. For instance, for a dispersion of 17 ps/nm·km, a network diameter of 150 km, and a transceiver tuning range of 35 nm the time difference between the "slowest" and "fastest" wavelength is equal to 89.25 ns which corresponds to approximately 28 bytes at 2.5 Gb/s. Accordingly, each slot has to carry a 28 byte long padding field. We observe from Fig. 8.7 that about 40% of the data packets are only 40 bytes long. Thus, for about 40% of the data packets the padding would result in an overhead of approximately 70%. Note that scheduling on an aggregate basis helps decrease the overhead and increase the channel utilization significantly. Furthermore, switching of longer data aggregates would also allow for deploying transceivers with a larger tuning time but a wider tuning range. Finally, the transfer rate of a TCP connection is upper bounded by

$$Rate \leq \frac{MSS}{RTT} \frac{1}{\sqrt{pl}} \tag{8.3}$$

where MSS denotes the maximum segment size, RTT denotes the round trip time, and pl denotes the packet loss rate [MSMO97]. Hence, for the TCP throughput performance it is advantageous to send large segments. With TCP on top of our MAC protocol it appears to be beneficial to transport large packets.

8.4 Conclusions

For investigating the feasibility of our network we have considered one possible signalling realization proposal made in the literature. By using a cost–effective off–the–shelf LED, control packets are sent at a rate of approximately 10 kb/s. We have shown that our network based on a $D \times D$ AWG is able to accomodate up to 250 nodes within metropolitan regions at an acceptable BER $< 10^{-9}$ by making use of $2D$ EDFAs. Furthermore, we have introduced wormhole scheduling which achieves a very good performance despite the slow control channel. By means of trace file driven simulations we have shown that a normalized mean aggregate throughput of asymptotically 100% and small packet loss can be achieved. Furthermore, with CDMA multiple control channels are created, one for each CDMA code. Thus, CDMA alleviates the signalling bottleneck resulting in an increased number of successfully transmitted control packets and an improved throughput–loss and in particular delay performance of the network. Wormhole scheduling also improves the channel utilization in slotted single–hop WDM networks where each slot has to have a padding field due to chromatic dispersion. Furthermore, since in wormhole scheduling transceivers are switched on an aggregate rather than packet basis *tuning times of the transceivers do not necessarily have to be in the nanosecond range.* Such transceivers with a slower tuning speed allow for transceivers with a wider tuning range. As a consequence, more FSRs of the underlying AWG can be exploited resulting in an improved throughput–delay performance of the network, as shown in Section 6.1.

Chapter 9

Protection

The central hub of our AWG based single–hop network forms a *single point of failure*. In this chapter, we address the *survivability* of our AWG based network. There are different types of network failures comprising node, fiber, and hub failure. While node and fiber failures have only a local effect in our network, the central AWG represents a single point of failure. That is, if the AWG goes down the network connectivity is entirely lost. In the following, we concentrate on the protection of this single point of failure. Under normal operation the proposed protection scheme not only avoids the single point of failure but also enables spatial wavelength reuse at all AWG ports at any given time and allows for transmitting control at significantly higher line rates by replacing the broadband light source with a laser diode [FMR03].

In Section 9.1 we introduce the novel concept of *heterogeneous* protection. The protection architecture is described in Section 9.2. Section 9.3 outlines the MAC protocols for the normal and various backup operating modes of the network, whose throughput–delay performance is evaluated in Section 9.4. Section 9.5 summarizes the results and merits of our proposed heterogeneous protection scheme.

9.1 Heterogeneous protection

Every single–hop network, including our AWG based network, suffers from a single point of failure. As opposed to multihop networks, the operation of single–hop networks is immune from node failures since nodes do not have to forward traffic. However, if the central hub — be it a PSC or an AWG — fails the network connectivity is entirely lost due to missing alternate paths. Therefore, protection of single–hop networks is essential to ensure survivability, in particular in the face of an ever increasing amount of traffic owing to higher line rates, larger wavelength counts, and spatial wavelength reuse as provided in our AWG based network.

Clearly, two protection schemes which come to mind are conventional $1:1$ and $1+1$ protection. In these cases, the network would consist of two AWGs in parallel. An example for $1:1$ protection is given in [SNY$^+$01], where the traffic is routed over the backup AWG by activating optical space switches if the working AWG goes down. A generalization of this approach is studied in [HBP$^+$98], where r AWGs working in parallel are protected by n identical standby AWGs. Again, these standby AWGs are used only in case of failure, thus implementing a conventional $n:r$ protection. Note that conventional $1+1$, $1:1$, and $n:r$ protection schemes are rather inefficient: While in the $1+1$ protection the backup device is used to carry duplicate traffic, in the $1:1$ (and also $n:r$) protection the backup device is not used at all during normal operation.

Given its totally passive nature, our AWG based network is rather reliable, albeit not tolerant against hub failure. As a result, normal network operation is provided most of the time leading to wasted backup resources in case of 1 : 1 protection. The network efficiency can be significantly improved by using the backup device for data and/or control also during normal operation, i.e., when the primary device is functional. This idea has been realized in two single–hop networks which are based on two PSCs working in parallel. While in [ACG+88] one PSC is used for data transmission and the other one for data reception, in [AGKV88] the primary PSC carries only data and the secondary PSC only control. Note that these dual–star networks provide a higher degree of concurrency but still suffer from a single point of failure. This is because if one PSC goes down normal network operation becomes impossible due to missing transmission or reception of data or control.

To enable network survivability in an efficient manner we propose a novel protection scheme which can be applied not only in our AWG based network but also in all PSC based single–hop networks. The resultant single–hop network consists of one AWG and one PSC in parallel, which we call the AWG||PSC network and describe in Section 9.2. During normal operation, i.e., both AWG and PSC are functional, the AWG||PSC network combines the respective strengths of both devices: Control is sent over the PSC exploiting its inherent broadcast capability, while data is transmitted over the AWG. Note that this out–of–band signalling allows for spatially reusing all wavelengths at each AWG port *continuously*, i.e., reuse is not interrupted by periodic reservation phases anymore. In case either device fails, AWG and PSC mutually protect each other (as described in Section 9.3). We refer to this proposed protection scheme as *heterogeneous* protection as opposed to the conventional homogeneous protection where both primary and secondary devices are identical, i.e., two AWGs or two PSCs.

9.2 AWG||PSC architecture

Fig. 9.1 depicts the architecture of the proposed AWG||PSC network. The AWG and PSC work in parallel. Each of the N nodes is connected to the AWG||PSC network by two pairs of fibers, one pair for the AWG and one pair for the PSC. Each pair of fibers contains one fiber for transmission and one fiber for reception. By attaching each node to the AWG||PSC network via two pairs of fibers one–to–one fiber backup for an improved fiber protection and survivability is provided.

The AWG||PSC network and node architecture is shown in greater detail in Fig. 9.2. The

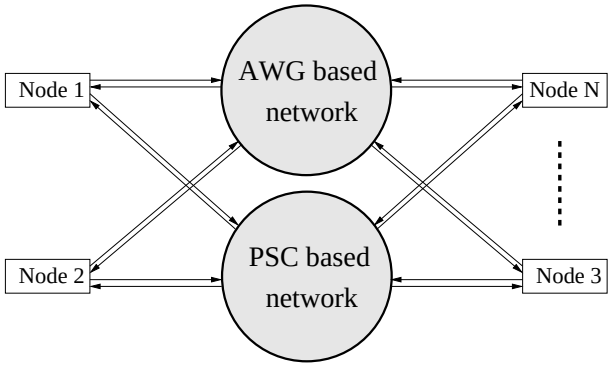

Figure 9.1: AWG||PSC network architecture.

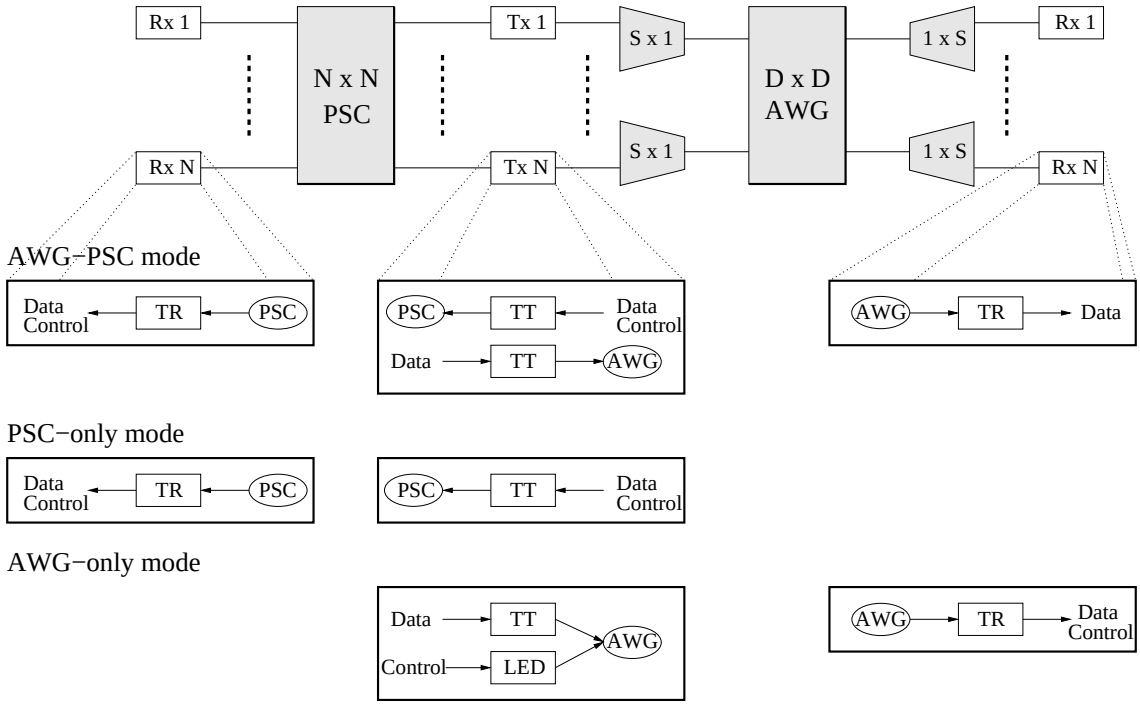

Figure 9.2: Detailed AWG||PSC network and node architecture.

AWG based part is our familiar network consisting of a $D \times D$ AWG with an $S \times 1$ combiner and a $1 \times S$ splitter attached to each AWG input and output port, respectively. The PSC based part consists of an $N \times N$ PSC. Each node's transceiver (Tx, Rx) is attached to a different PSC port. Every node is equipped with two tunable transmitters (TTs), two tunable receivers (TRs), and one broadband LED light source. The transceivers are tunable over a range of $D \cdot R$ wavelengths, where R again denotes the number of used FSRs of the underlying $D \times D$ AWG. One TT–TR pair and the LED are deployed in our AWG based network. The second TT–TR pair is used in the PSC based part of the AWG||PSC network. Fig. 9.2 shows how the transceivers are used in the various operating modes of the AWG||PSC network. In the AWG–PSC mode, i.e., both AWG and PSC are functional, only data is sent over the AWG while the PSC carries control and additional "overflow" data which can not be accomodated in the AWG based part of the AWG||PSC network. Thus, both TT–TR pairs are used simultaneously at each node. In the PSC–only mode, i.e., the AWG fails and only the PSC is functional, only one TT–TR pair is used to send/receive not only control but also data over the PSC. The second TT–TR pair attached to the AWG remains unused. In the AWG–only mode, i.e., the PSC fails and only the AWG is functional, the AWG||PSC network reduces to our familiar AWG based network. Accordingly, the LED is used for broadcasting spread control by means of spectral slicing, the TT sends data, and the TR is used to retrieve both data and despread control. The second TT–TR pair attached to the PSC remains unused. Recall from Section 3.3.5 that by deploying a TT–TR instead of a TT–FR (FR = fixed–tuned receiver) node structure the throughput–delay performance of the PSC based network can be improved due to the increased flexibility. Moreover, deploying two identical tunable transceivers at each node allows for nodal backup in the AWG||PSC network.

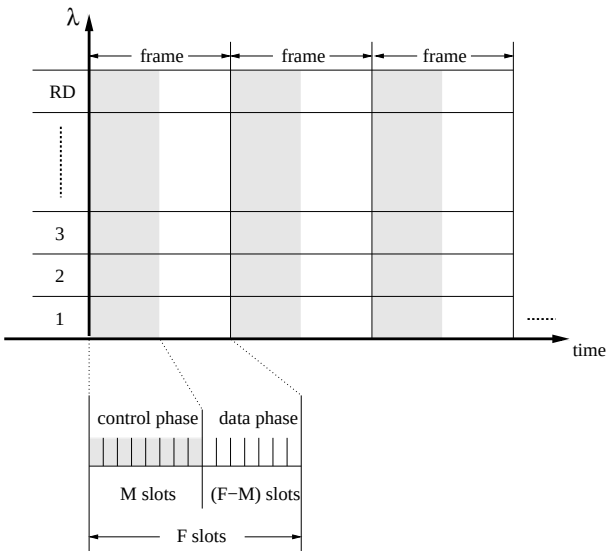

Figure 9.3: Wavelength assignment in the PSC based part of the AWG||PSC network.

9.3 MAC protocols

In this section we describe the MAC protocols for the normal AWG–PSC operating mode as well as the two backup PSC–only and AWG–only modes.

9.3.1 AWG–PSC mode

In the AWG–PSC mode time is divided into frames which are repeated periodically. Each frame comprises F slots with a slot length equal to the transmission time of a control packet. The format of a control packet is the same as described in Section 5.3. In the AWG based part of the network each frame is entirely used for data traffic without any reservation windows. On the other hand, in the PSC based part each frame is subdivided into a control phase and a data phase, as shown in Fig. 9.3. The control phase consists of the first M slots of a given frame. The remaining $(F - M)$ slots form the data phase of that frame. Pretransmission coordination takes place during the control phase, whereas data is sent during the data phase of each frame. Nodes ready to send data packets apply the same reservation protocol as outlined in Section 5.3 by sending a control packet in one of the M slots of the control phase on a prespecified wavelength. During the data phase all wavelengths, including the prespecified control wavelength, are used for data transmission. Uncollided control packets are scheduled in a first–come–first–served manner. More specifically, data packets belonging to uncollided control packets are assigned wavelength channels starting with the earliest available frame within the scheduling window on the lowest FSR of the AWG. Once all FSRs of the AWG are allocated for that frame, the wavelength assignment starts on the PSC beginning with the lowest wavelength channel. Note that a given source node can use only one wavelength per AWG FSR but any arbitrary wavelength on the PSC in order to reach the corresponding destination node due to the broadcast nature of the PSC. Once all AWG FSRs and PSC wavelengths are assigned in the given frame, the assignment continues in the next frame, again starting with the lowest FSR of the AWG, and so on. This procedure is repeated until the end of the scheduling window is reached. Control packets which have either collided or whose corresponding data packets could not be scheduled

are retransmitted with probability p in the next frame in the PSC based part of the AWG||PSC network.

9.3.2 PSC–only mode

The network operates in the PSC–only mode when the AWG goes down. Nodes learn about the failure of the AWG if scheduled data packets do not arrive after the corresponding delay. A node which has learnt about the AWG failure informs all other nodes by sending a control packet in the following frame. After receiving this control packet, all nodes change from the AWG–PSC to the PSC–only mode.

In the PSC–only mode, nodes make their reservations during the control phase of each frame. Successfully scheduled data packets are sent only during the data phase in the PSC based part of the network.

9.3.3 AWG–only mode

The network operates in the AWG–only mode when the PSC goes down. Nodes learn about the failure of the PSC when they do not receive any signal from the PSC, i.e., if a loss of signal (LOS) occurs. This missing signal can be periodic synchronization signals which are no longer broadcast over the failed PSC.

In the AWG–only mode, the AWG||PSC network reduces to our familiar AWG based single–hop network of Chapter 5. As a consequence, each node sends not only data but also reservation requests over the AWG. Recall that control packets are signalled in–band by means of spectrally slicing the spread broadband (LED) signal.

9.4 Results

The throughput–delay performance of the above mentioned MAC protocols can be analyzed in a similar manner as outlined in Section 6.2. Due to space constraints we do not include the analysis here. For details the interested reader is referred to [FMR03]. We only note that in the analysis we consider uniform unicast traffic and a fixed data packet size of $(F - M)$ slots. Furthermore, we assume that the scheduling window size is equal to one frame in the AWG–PSC and PSC–only modes, and one cycle, i.e., D frames, in the AWG–only mode. In the following numerical results we do not take the propagation delay into account. The number of nodes N, the AWG degree D, and the number of used FSRs R of the underlying AWG are set to the following default values: $N = 200$, $D = 4$, and $R = 2$. Each node's transceiver is assumed to be tunable over a fixed range of $D \cdot R = 8$ wavelengths within a negligible tuning time. The analytical results are verified by simulations. Each simulation was run for 10^6 frames including a warm–up phase of 10^5 frames; the 99% confidence intervals are less than 1% of the corresponding sample means. Throughout the following numerical results, simulations are run for the mean arrival rate σ values of 0.01, 0.05, 0.10, 0.15, 0.20, 0.40, 0.60, and 0.80.

Fig. 9.4 depicts the mean delay (given in frames) vs. the mean aggregate throughput (given in packets/frame) for a mean arrival rate $\sigma \in [0.01, 1]$ and different $D \in \{2, 4, 8\}$ and correspondingly different $R \in \{4, 2, 1\}$. We set the number of slots per frame to $F = 200$, the number of reservation slots per frame to $M = 80$, and the retransmission probability of unsuccessful control packets to $p = 0.2$. At small traffic loads $D = 2$ achieves the smallest delay. However, at medium to high traffic loads $D = 2$ exhibits a larger delay and a smaller

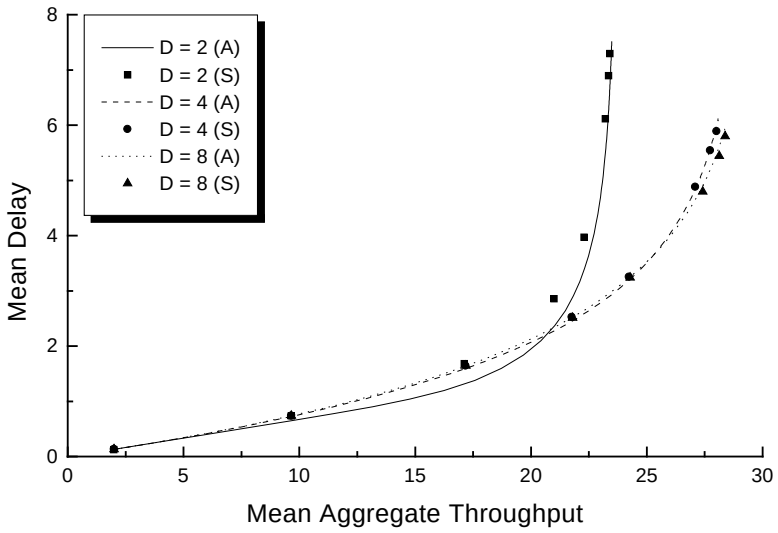

Figure 9.4: Mean delay (frames) vs. mean aggregate throughput (packets/frame) for different $D \in \{2, 4, 8\}$, $F = 200$, and $M = 80$.

throughput than $D = 4$ and $D = 8$. This is due to the fact that the number of available wavelength channels in the network is limited by $D \cdot R + D \cdot D \cdot R$, where the first term accounts for the channels available in the PSC based part and the second term accounts for the spatially reused wavelength channels available in the AWG based part of the AWG||PSC network. Hence, for $D = 2$ a total of 24 wavelength channels can be used. As shown in Fig. 9.4, at high traffic loads almost all wavelength channels are used and further control packets can not be scheduled resulting in an increased delay due to the retransmission of unsuccessful control packets. By increasing D the number of available wavelength channels can be increased. For $D = 4$ ($R = 2$) and $D = 8$ ($R = 1$) we obtain 40 and 72 wavelength channels, respectively. Fig. 9.4 illustrates that $D = 4$ and $D = 8$ achieve approximately the same improvement of the throughput–delay performance compared to $D = 2$. At medium to high traffic loads the delay increases owing to the retransmission of control packets which have been collided in the random access reservation channel.

To reduce the number of channel collisions of control packets in the reservation channel we increase the number of reservation slots per frame and set $M = 170$. The retransmission probability of unsuccessful control packets is set $p = 0.85$. Moreover, we set $F = 2M = 340$. In doing so, each frame on the AWG based part of the AWG||PSC network can accomodate two data packets in each frame, one during the control phase and another one during the data phase of the PSC based part of the AWG||PSC network. This increases the degree of concurrency since data packets can be sent not only during the data phase of each frame on both PSC and AWG but also during the control phase on the AWG while control packets are transmitted on the PSC. Fig. 9.5 shows that by setting $M = 170$ and $F = 2M = 340$ the throughput–delay performance of the AWG||PSC network is improved dramatically.

Fig. 9.6 compares the throughput–delay performance of the AWG–PSC, PSC–only, and AWG–only network operating modes. In the PSC–only mode the aggregate throughput is upper

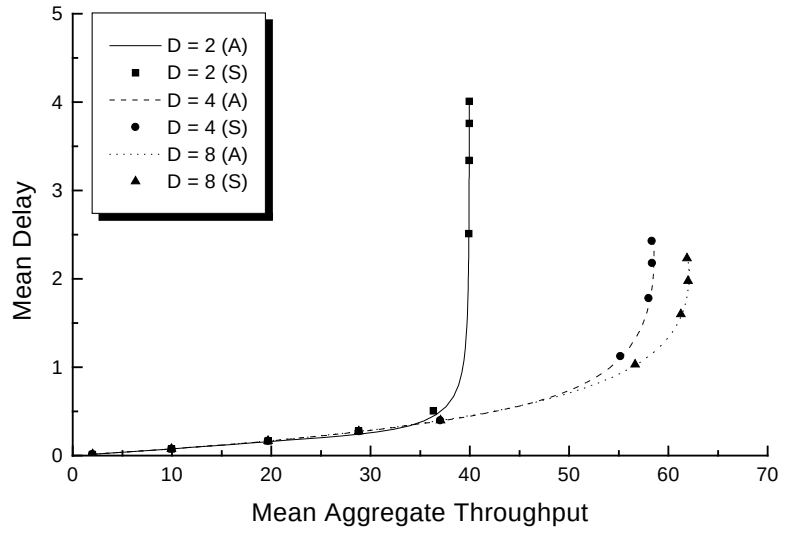

Figure 9.5: Mean delay (frames) vs. mean aggregate throughput (packets/frame) for different $D \in \{2, 4, 8\}$ and $F = 2M = 340$.

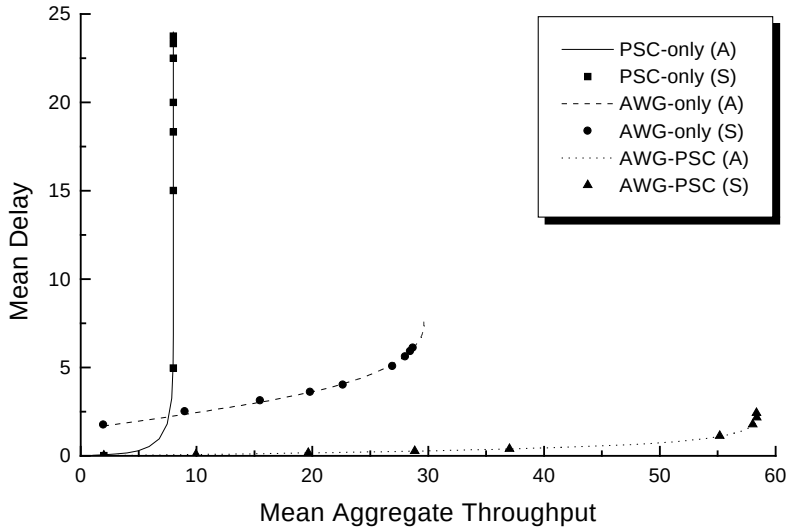

Figure 9.6: Mean delay (frames) vs. mean aggregate throughput (packets/frame) for different operating modes of the AWG||PSC network.

bounded by 8 packets/frame since only 8 wavelengths are available. Due to this rather small network capacity many control packets have to be retransmitted leading to a large delay at medium to high traffic loads. This capacity bottleneck is removed in the AWG–only mode where 32 wavelength channels are available by means of spatial wavelength reuse. As a result, the throughput–delay performance is significantly better in the AWG–only mode than in the PSC–only mode. (This comes as no surprise. We have observed the same result in Section 4.2 where we have compared PSC and AWG based single–hop networks.) Note, however, that at small traffic loads the mean delay in the AWG–only mode is larger than in the PSC–only mode. This is due to the cyclic reservation timing structure of the AWG–only mode where nodes ready to send control packets have to wait for the corresponding frame as opposed to the PSC–only mode where nodes can make a reservation in any arbitrary frame. Fig. 9.6 shows that during normal operation, i.e., in the AWG–PSC mode, the AWG||PSC network provides by far the best throughput–delay performance. In the AWG–PSC mode the AWG||PSC network uniquely combines the strengths of PSC and AWG such that for all traffic loads the delay is very small while achieving a quite large throughput. To emphasize this point, observe from Fig. 9.6 that in the AWG–PSC mode the delay of the AWG||PSC network does not exceed 3 frames and the maximum throughput is equal to 59 packets/frame. At the same level of delay, the throughput of the PSC–only network and the AWG–only network are approximately 8 and 12 packets/frame, respectively. Thus, during normal operation the AWG||PSC network effectively *triples* the combined throughput of a stand–alone PSC and a stand–alone AWG based network.

Next, let us compare the throughput–delay performance of the proposed heterogeneous and conventional homogeneous protection schemes. Beside the AWG||PSC network we consider two other two–device single–hop networks which consist of either two PSCs or two AWGs in parallel. We call them PSC||PSC and AWG||AWG network, respectively. Since in the PSC based network the number of available wavelengths is rather small we impose the control phase only on one PSC of the PSC||PSC network. This allows for sending two data packets in each frame on the second PSC of the PSC||PSC network. Consequently, with 8 wavelengths available, the PSC||PSC network has a capacity of 24 packets per frame. As shown in Fig. 9.7, the capacity of the PSC||PSC is not large enough to accomodate all reservation requests at medium to high traffic loads. As a result, the delay increases due to the retransmission of unsuccessful control packets. Conversely, in the AWG based network the number of available wavelength channels is rather large due to spatial wavelength reuse. At high traffic loads, however, the cyclic reservation timing structure becomes a bottleneck since nodes are allowed to send control packets only in one frame per cycle. To alleviate this bottleneck we impose a control phase on both AWGs of the AWG||AWG network. Thus, the AWG||AWG network provides a capacity of 64 packets/frame. A given node ready to send a control packet randomly selects either AWG with the same probability of 0.5. Fig. 9.7 depicts that owing to the larger number of available wavelength channels the AWG||AWG outperforms the PSC||PSC network in terms of throughput and delay at medium to high traffic loads. Again, we observe that at low traffic loads the AWG based network exhibits a larger delay than the PSC||PSC network due to the cyclic reservation timing structure. Note that the maximum throughput of the AWG||AWG network is clearly below its capacity. This is because the offered network load is not large enough. To see this, remember that only nodes attached to the same combiner are allowed to transmit control packets in the control phase of a given frame. With $N = 200$ and $D = 4$ the number of control packets per frame is upper bounded by $N/D = 50$ which is not enough to fully capitalize on the bandwidth of the AWG||AWG network. Fig. 9.7 demonstrates that the proposed heterogeneous protection in the AWG||PSC network significantly outperforms its homogeneous counterparts

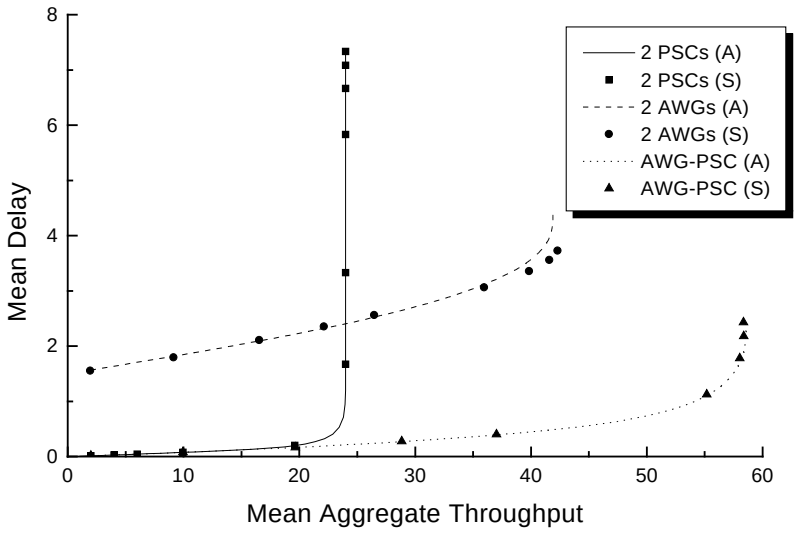

Figure 9.7: Mean delay (frames) vs. mean aggregate throughput (packets/frame) for different protection schemes: PSC||PSC (2 PSCs), AWG||AWG (2 AWGs), and AWG||PSC networks.

in terms of throughput and delay.

9.5 Conclusions

We have described and evaluated the AWG||PSC network which consists of one AWG and one PSC in parallel. The AWG||PSC network improves the network survivability by removing the single point of failure and remains functional when either the AWG or the PSC fails. Given the passive nature of both AWG and PSC the AWG||PSC network is rather reliable. As a consequence, both devices are available for communications most of the time. To increase the network efficiency both AWG and PSC are used simultaneously. During normal operation the AWG||PSC network uniquely combines the respective strengths of AWG and PSC. Control is efficiently sent to all nodes by exploiting the inherent broadcast nature of the PSC. The resultant out–of–band signalling is done with a laser diode instead of a broadband LED light source. As a consequence, control packets can be sent at much higher rates relieving the bottleneck of the original AWG based network. In addition, wavelengths can be spatially reused at all AWG ports *continuously*, i.e., reuse is not interrupted by periodic reservation phases any more. By means of analysis and verifying simulations we have found that the throughput of the AWG||PSC network is significantly larger than the total throughput obtained by combining the throughput of a stand–alone AWG based network with the throughput of a stand–alone PSC based network. We have also found that the AWG||PSC network gives over a wide operating range a significantly better throughput–delay performance than a network consisting of either two load sharing PSCs in parallel or two load sharing AWGs in parallel. This result proves the superiority of our proposed heterogeneous protection scheme to its conventional homogeneous counterparts which consist of two identical working and protecting devices. By equipping each node with two identical transceivers and connecting each node to the single–hop network via two pairs of fibers

the AWG||PSC network is also able to provide transceiver and fiber backup.

Concluding, we point out that the proposed protection scheme is generally applicable not only to AWG based but also to all PSC based single–hop networks reviewed in Chapter 3. Furthermore, note that all logical topologies that are embedded on a physical star topology also suffer from a single point of failure. The proposed protection approach helps improve the survivability of logical multihop WDM networks as well.

Chapter 10

Conclusions

Optical WDM networks can be found throughout the network hierarchy. Due to various inefficiencies current SONET/SDH metro ring networks create the so–called *metro gap* which prevents high–speed packet switched LANs, access technologies, and service providers from tapping into the vast amounts of bandwidth available in the backbone. This bandwidth bottleneck at the metro level is anticipated to become more severe in the face of the ever increasing number of users and bandwidth–hungry applications and the steadily growing amount of local intra–MAN traffic.

For the design of optical WDM networks several optical devices are readily available. In particular, the AWG appears to be a very promising component which allows for realizing high–efficiency metro WDM networks. Athermal, polarization–independent AWGs with an insertion loss of $\sim$5 dB, crosstalk of ~ -30 dB, uniform loss across all channels, and a flat broadened spectral response can be fabricated as integrated components. Many applications exist that are based on an AWG. In this work we used the AWG as a full–interconnection wavelength router. As opposed to the broadcast–and–select PSC the AWG is a *wavelength–selective* device. As a consequence, using one FSR an $N \times N$ AWG is able to simultaneously support a total of N^2 collisionfree transmissions compared to only N in case of an $N \times N$ PSC. In other words, unlike the PSC the AWG allows for *spatial wavelength reuse*. Broadcasting in AWG based wavelength–addressed networks can be achieved by spectrally slicing the optical output signal of a broadband light source, e.g., an LED. We have seen in Chapter 2 that fast tunable transceivers have begun to emerge from research laboratories. Multisection injection–current lasers and receivers deploying EOTFs exhibit a tuning range of $15 - 30$ nm and a tuning time of $\sim$10 ns.

Recently reported packet switched metro WDM network testbeds have either a physical ring or star topology. Most of them adopt the concept of a *home channel* where each node has its own dedicated wavelength for transmission or reception which can not be accessed by other nodes during idle periods. Furthermore, the applied MAC protocols are mostly targeted towards the switching of *fixed–size* packets.

In this work we focused on metro WDM networks with a *physical star* topology. Compared to ring (and bus) topologies, star networks offer a better optical power budget. Furthermore, star configurations are easy to install, configure, manage, and troubleshoot. Two types of *logical* topologies with different nodal structure requirements can be embedded on this physical star network: Single–hop and multihop topologies. While in multihop networks each node is equipped with fixed–tuned transceivers, transmitters and/or receivers have to be *tunable* in single–hop networks in order to enable any–to–any connectivity in one single hop (alternatively to the tunable transceivers each node could be equipped with an array of fixed–tuned transceivers,

one for each wavelength in the system. We ruled out this node structure due to its higher power consumption, more complex performance monitoring and management. Moreover, for survivability reasons each fixed–tuned transceiver would require an identical back–up device which makes such a node structure rather complex and prohibitively expensive, albeit feasible with current technology.). To keep the network costs low we considered single–hop networks where each node has only one single tunable transceiver. Wavelength–agile transceivers suffer from a tuning latency whose impact on the channel utilization largely depends on the transceiver type in use. Given the fact that the tuning times of the various transceiver technologies differ by several orders of magnitude, fast tunable transceivers, e.g., electro–optic transceivers, with a negligible tuning penalty of a few nanoseconds should be deployed in order to enable efficient packet switching in optical high–speed single–hop WDM networks.

We have seen in Chapter 4 that an AWG based multihop network, using one FSR of the underlying AWG, requires at least four fixed–tuned transceivers per node to outperform its single–hop counterpart with one fast tunable transceiver at each node in terms of network capacity. Therefore, we decided to concentrate on AWG based *single–hop* WDM networks which unlike multihop networks provide additional benefits as follows: (*i*) *Minimum mean hop distance* (unity), (*ii*) inherent *transparency* with respect to bit rate, modulation format, and protocol, i.e., no costly conversion from other protocols to the supported standard is needed, (*iii*) *future–proofness* in the sense that owing to the aforementioned transparency not only different legacy protocols but also future, not yet defined, protocols can easily be supported, (*iv*) easy *upgradability* since technologically advanced transceivers with a larger tuning range and a smaller tuning time can be placed at the network periphery without requiring any changes in the network itself, (*v*) *simplified management* since bits/packets are interpreted/processed only at the border of the passive network and node failures do not affect the network operation since they are not involved in packet forwarding, and (*vi*) an *improved throughput–delay performance* since no bandwidth is wasted due to forwarding and packets have to pass the central hub only once as opposed to multihop networks where especially in a high–speed environment the propagation delay might become the dominant delay factor owing to the multiple hops.

Unlike the PSC, the AWG allows for *spatial wavelength reuse* at each input port. Consequently, in AWG based single–hop networks the number of required wavelengths can be kept small which in turn allows for deploying tunable transceivers with a limited tuning range whose tuning time is in the range of few nanoseconds. We have seen in Chapter 4 that due to spatial wavelength reuse, AWG based single–hop WDM networks clearly outperform their PSC based counterparts in terms of throughput, delay, and packet loss, which have been the most common type of single–hop WDM network so far.

The proposed architecture combines the best of optics (photonics) and electronics in that transmission and switching/wavelength tuning are done in the optical (photonic) domain while buffering and logical operations are done in the electronic domain. The optical part of the proposed network combines the merits and mitigates the drawbacks of wavelength–routing and wavelength–insensitive components. It consists of an AWG with combiners (splitters) attached to each AWG input (output) port. Due to its passive nature the network is *reliable* and no wavelength stabilization is needed if an athermal AWG is deployed. While the splitters enable *optical multicasting* the AWG allows for extensive *spatial wavelength reuse*.

Each node is equipped with a *cost–effective* off–the–shelf broadband light source, e.g., LED, for control and for cost and management reasons only one *single* pair of transmitter and receiver for data which are tunable over one or more FSRs of the underlying AWG. As opposed to networks with fixed–tuned transmitters and/or receivers each node in our proposed system has

access to all wavelengths. This flexibility improves the channel utilization and throughput–delay performance and allows for *efficient multicasting* and *load balancing* especially of nonuniform traffic over all wavelengths. Control information is spread prior to modulating the broadband signal which is subsequently spectrally sliced by the AWG and thus broadcast to all receivers. Through spreading, control and data can be transmitted simultaneously without requiring any separate control channel and additional receiver at each node (*in–band signalling*) and improves the network *security*.

To efficiently support bursty traffic all wavelengths are assigned *dynamically* on demand by means of a reservation MAC protocol. Organizing the reservation process in a round–robin TDMA fashion avoids receiver collisions (destination conflicts) of control packets. The reservation slots are not fixed allocated but can be equally accessed by nodes ready to send a control packet. As a consequence, no reservation slots are wasted by idle nodes and new nodes can easily join the reservation process without any reconfiguration of the MAC protocol. For the control traffic we use a modified ALOHA protocol since due to their independence of propagation delay ALOHA based protocols can be upgraded to higher line rates without deteriorating the throughput. To increase the throughput of ALOHA we adopted reservation ALOHA (R–ALOHA) to indicate the existence of circuits in a collisionfree manner and CDMA where each spreading sequence creates a separate control channel which is used for balancing the control traffic load and providing an improved throughput–delay performance. ALOHA and CDMA help improve the network *scalability*. Uncollided control packets are scheduled in a deterministic distributed fashion based on the *global knowledge* available at all nodes. To acquire and maintain global knowledge each node processes the control traffic of all nodes. Due to the rather moderate number of nodes in metro networks we expect the total amount of control traffic to be acceptable. To avoid any computational bottleneck each node executes a simple identical first–come–first–served (FCFS) and first–fit data packet scheduling algorithm. Since no explicit acknowledgements are required the delay is reduced and the bandwidth utilization is increased. The arbitration algorithm is able to schedule *variable–size* packets and *completely prevents both channel and receiver collisions of data packets*. Furthermore, it supports hybrid switching, i.e., both packet and circuit switching, where the latter one is used for providing *guaranteed QoS*. Due to the round–robin reservation cycles, the random access of reservation slots, and the FCFS scheduling policy the MAC protocol provides *fairness* among nodes wishing to send data. The MAC protocol accomplishes traffic grooming in a distributed fashion in our network.

In summary, the proposed network and protocol are highly efficient. The degree of concurrency is significantly increased by *simultaneously transmitting control and data* by means of spreading techniques and CDMA without requiring any additional control channel and receiver, *spatially reusing all wavelengths* at each AWG port, and deploying *multiple FSRs* of the underlying AWG. The protocol *completely avoids* receiver collisions of control packets and both channel and receiver collisions of data packets, resulting in an increased wavelength channel utilization. The delay–utilization performance is further improved since no explicit acknowledgements are needed. In addition, the presented single–hop network *significantly reduces the protocol stack complexity*. The network layer is practically absent since routing is not needed. Instead of the conventional store–and–forward switching paradigm we introduced the concept of *wavelength tuning* where a given source node is able to reach different destinations by simply changing the wavelength. Moreover, since packets have to pass only one single passive all–optical hop bit errors are unlikely and can be handled by the transport layer, making the data link layer superfluous.

We have extensively investigated the proposed network architecture and MAC protocol by

means of analysis and/or simulation. Using three FSRs instead of one of the underlying AWG increases the maximum throughput by 88%. New nodes and/or transceivers with a wider tuning range can benefit from deploying multiple FSRs, thus making the network *scalable* and *upgradable*. Spatial wavelength reuse increases the maximum throughput by more than 60% and reduces the mean delay by approximately 50%. As a result, the network *efficiency*, *utilization* of resources (wavelengths, transceivers), and *cost–effectiveness* are significantly improved. *Multicasting* can be done in our AWG based network very efficiently. Sending multicast packets during the reservation phase of each frame dramatically increases the receiver utilization. In addition, the throughput–delay performance of the network is significantly increased by *partitioning* multicast transmissions while spatially reusing all wavelengths. For a 50–50 mix of unicast and multicast traffic a receiver utilization of almost 60% can be achieved. *Self–stability* is achieved by means of backoff. The packet loss is significantly reduced by increasing the number of reservation slots per frame, the number of used FSRs (which implies a smaller channel spacing for a given transceiver tuning range), and the scheduling window size. Due to spatial wavelength reuse our network achieves a mean *wavelength utilization* of more than 100% and a mean *channel utilization* of approximately 53%. In our benchmark comparison we have shown that for a wide range of traffic loads our AWG based network provides a mean aggregate throughput that is about 70% larger than the *maximum* aggregate throughput of DT–WDMA that runs on top of a PSC based single–hop metro WDM network. Circuits benefit from R–ALOHA resulting in a significantly increased channel utilization at the expense of the throughput–delay performance of packet switched traffic. We have seen that for different traffic loads different network configurations are needed to achieve the best throughput–delay performance.

To generate new revenue opportunities in today's competitive environment we let our network provide multiple services. Such a *multiservice* network is able to carry heterogeneous traffic with different throughput–delay requirements resulting in an increased resource utilization. To enable efficient multiservice convergence, architecture and protocol parameters have to be optimized and dynamically reconfigured according to the current requirements. For this purpose we applied powerful *multiobjective optimization* techniques instead of conventional single–objective approaches. We developed a genetic algorithm based methodology which finds for different traffic loads and packet size distributions so–called *Pareto–optimal* solutions for the architecture planning and network operation in a computationally efficient manner. In our numerical investigations we have taken the limited transceiver tuning range into account and examined the trade–off between spatial wavelength reuse and using multiple FSRs of the underlying AWG. Our results show that an AWG with a small number of ports is the best choice for achieving efficient multiservice convergence. Different traffic loads and packet size distributions are accomodated by adjusting the protocol parameters in software. A small AWG degree has also the positive side–effect that fewer EDFAs are needed which further reduces the network costs.

Using a cost–effective off–the–shelf LED as broadband light source for broadcasting control packets we have proven the feasibility of our network. The proposed network is able to accomodate up to 250 nodes within metropolitan regions for a BER $< 10^{-9}$ by deploying one EDFA between each combiner and AWG input port and each AWG output port and splitter, respectively. Capitalizing on the cyclic timing structure of our MAC protocol we have introduced the concept of *wormhole scheduling* where each control packet tries to make a reservation for data aggregates rather than on a per–packet basis. Our trace file driven simulation results have shown that with wormhole scheduling a normalized mean aggregate throughput of asymptotically 100% and low packet loss can be achieved. Furthermore, we have shown that the throughput–loss and in particular the delay performance of the network can be improved

by deploying *CDMA* where each CDMA code creates an *additional* control channel and thereby increases the number of successfully transmitted control packets. Wormhole scheduling also allows for deploying transceivers with a wider tuning range which in turn results in an improved throughput–delay performance of the network owing to the larger number of used FSRs of the underlying AWG.

While our AWG based single–hop network is immune from node failures the central AWG forms a single point of failure. To provide survivability in an efficient manner we proposed and evaluated the novel concept of *heterogeneous* protection where the AWG works in parallel with a PSC. This approach uniquely combines the respective strengths of both AWG and PSC in that data is sent over the AWG while control (and overflow data) benefits from the inherent broadcast capabilities of the PSC. By using a separate wavelength and transceiver, control can be sent at much higher rates and wavelengths can be spatially reused at all AWG ports *continuously* without being interrupted by periodic reservation phases anymore. Both AWG and PSC are used simultaneously during normal operation and provide mutual protection in case either device fails. By means of analysis and verifying simulation we have shown that the heterogeneous protection scheme provides a dramatically increased throughput–delay performance of the network. The proposed protection approach is also applicable to all PSC based single–hop WDM networks and all logical multihop topologies embedded on a physical star network.

10.1 Future research avenues

Possible avenues for future research include the performance evaluation of the proposed AWG based single–hop WDM network for other traffic matrices, e.g., client–server traffic, and traffic models, e.g., Markov modulated Bernoulli process (MMBP) and self–similar traffic. Also examining the interplay between our MAC protocol and TCP would be an interesting issue. As we have seen, the proposed reservation MAC protocol provides a flexible framework. For instance, by putting additional fields into each control packet, e.g., packet age or priority, and deploying other scheduling algorithms, multiclass service differentiation and performance improvements are worth being examined. Another interesting direction of research is the design and evaluation of hybrid MAC protocols where nodes apply different protocols according to the current traffic load. For instance, using its buffer length as an indicator for the network load, a given node uses a tell–and–go protocol at light traffic loads and our reservation protocol only at high traffic loads. As a consequence, at light traffic loads each node does not have to go through the pretransmission coordination process resulting in a decreased medium access delay. Finally, a quantitative performance comparison between physical ring and star networks could be very helpful to better understand their respective merits and drawbacks. The thus gained insight could help in the design novel hybrid network architectures.

10.2 Standardization activities

At the time of writing IEEE 802.17 RPRWG and IETF WG IPoRPR are working on the new resilient packet ring (RPR) standard for packet switched metro ring networks which is anticipated by the end of 2003. RPR is a bidirectional ring network using one wavelength for each direction which is OEO converted at each node. RPR deploys both counter–rotating rings to route traffic using two performance enhancing schemes: Spatial reuse by means of destination stripping of packets and shortest path steering. These two schemes are well understood and it was shown

that they are able to improve the network capacity significantly [CO90][CO93]. Another building block of RPR are the so–called insertion buffers for storing in–transit ring traffic at each node. Buffer insertion rings have already been investigated in the year 1983 [HSW83]. Several access and fairness protocols are proposed within the RPR standardization process and at the time of writing it is too early to say which protocols will emerge as the winners. However, we do want to draw the reader's attention to possible future improvements and upgrades of RPR. Clearly, one way to increase the network capacity of RPR is the deployment of WDM instead of one single wavelength on each ring. Furthermore, all–optical node architectures have been proposed to avoid OEO conversions at each node [BBLN02]. Such all–optical node structures allow to add/drop only a few wavelengths at a given node while bypassing the remaining wavelengths. These optical bypasses not only reduce the processing requirements at each node but also decrease the *logical* mean hop distance and *logical* diameter of the network [GA96]. It was shown in [RH95a][RH95b] that so–called *meshed rings* are able to further increase the capacity of spatial–reuse rings. In meshed rings wavelength routers are inserted on the ring. These wavelength routers are connected with each other by fibers and provide *physical short–cuts* such that packets do not have to travel along the entire ring but can skip several nodes on the ring. Consequently, packets travel along shorter paths resulting in a significantly increased network capacity. Note that meshed rings require several wavelength routers and the mean hop distance is larger than one.

To minimize the number of required wavelength routers and the mean hop distance of meshed rings we propose an *evolutionary* upgrade of RPR where our AWG based single–hop WDM network is placed at the center of the RPR ring network. Apart from being part of the ring network each node is attached to a separate combiner/splitter port. In doing so, the multiple wavelength routers of meshed ring networks are replaced with one single wavelength router (AWG) and the mean hop distance is minimum (unity). Note that such a hybrid ring–star network is able to optimize the two performance enhancing schemes of RPR: (1) The central AWG allows for extensive spatial wavelength reuse at each port, and (2) the shortest path is minimum between any pair of nodes. In addition, such a hybrid ring–star topology provides a better resilience than bidirectional ring networks [HBP$^+$98]

Appendix A

Publications

Journals

- M. Maier and A. Wolisz
  "Demonstrating the Potential of Arrayed–Waveguide Grating Based Single–Hop WDM Networks", *Optical Networks Magazine*, vol. 2, no. 5, pp. 75–85, Sept. 2001

- M. Maier, M. Reisslein, and A. Wolisz
  "Towards Efficient Packet Switching Metro WDM Networks", *Optical Networks Magazine, Special Issue on Optical Packet Switching Networks*, vol. 3, no. 6, pp. 44–62, Nov./Dec. 2002 (**Invited Paper**)

- M. Maier, M. Reisslein, and A. Wolisz
  "A Hybrid MAC Protocol for a Metro WDM Network Using Multiple Free Spectral Ranges of an Arrayed–Waveguide Grating", *Computer Networks*, vol. 41, no. 4, pp. 407–433, March 2003

- H. Yang, M. Maier, M. Reisslein, and W. M. Carlyle
  "A Genetic Algorithm based Methodology for Optimizing Multi–Service Convergence in a Metro WDM Network", *IEEE/OSA Journal of Lightwave Technology*, vol. 21, no. 5, May 2003

- M. Scheutzow, M. Maier, M. Reisslein, and A. Wolisz
  "Wavelength Reuse for Efficient Packet–Switched Transport in an AWG–Based Metro WDM Network", *IEEE/OSA Journal of Lightwave Technology*, vol. 21, no. 6, June 2003

- C. Fan, M. Maier, and M. Reisslein
  "The AWG||PSC Network: A Performance Enhanced Single–Hop WDM Network with Heterogeneous Protection", *submitted to IEEE Journal on Selected Areas in Communications*, Dec. 2002

- M. Maier, M. Scheutzow, and M. Reisslein
  "The Arrayed–Waveguide Grating Based Single–Hop WDM Network: An Architecture for Efficient Multicasting", *IEEE Journal on Selected Areas in Communications, in revision*, June 2003

Book Chapter

- H. Woesner, M. Maier, and A. Wolisz
"Comparison of Single–Hop and Multihop AWG–based WDM Networks", *Next Generation Optical Network Design and Modelling*, pp. 51–65, Kluwer Academic Publisher, Feb. 2003

Conference Proceedings

- M. Maier and A. Wolisz
"Single–Hop WDM Network with High Spectrum Reuse Based on an Arrayed Waveguide Grating", *IEEE/ACM/SPIE Optical Network Workshop*, University of Dallas, Texas, USA, Jan./Feb. 2000

- M. Maier, M. Reisslein, and A. Wolisz
"High–performance Switchless WDM Network Using Multiple Free Spectral Ranges of an Arrayed–Waveguide Grating", *Proc., Terabit Optical Networking: Architecture, Control, and Management Issues, Part of Photonics East*, vol. 4213, pp. 101–112, Boston, USA, Nov. 2000 (Paper won the **Best Paper Award** of the conference)

- K. Dolzer, T. Fischer, C. Gauger, M. Jäger, M. Maier, E. Patzak, M. Schlosser, D. Schupke, F.–J. Westphal, and H. Woesner
"Vergleich von Architekturen für das zukünftige optische Internet" (in German), *Proc., 3. ITG Fachtagung Photonische Netze*, pp. 79–88, Leipzig, April 2002

- M. Maier, M. Scheutzow, M. Reisslein, and A. Wolisz
"Wavelength Reuse for Efficient Transport of Variable–Size Packets in a Metro WDM Network", *Proc., IEEE INFOCOM*, pp. 1432–1441, New York, NY, USA, June 2002

- C. Fan, M. Maier, and M. Reisslein
"The AWG||PSC Network: A Performance Enhanced Single–Hop WDM Network with Heterogeneous Protection", *Proc., IEEE INFOCOM*, San Francisco, USA, March/April 2003

Invited Talks

- M. Maier
"On the Design of High–Efficiency WDM Networks", *IEEE Communications and Signal Processing, Phoenix Chapter*, Arizona State University, Tempe, USA, Feb. 2001

- M. Maier
"Next–Generation Metro WDM Networks", *IEEE Computer Communications Workshop*, Charlottesville, Virginia, USA, Oct. 2001

Talks without Proceedings

- H. Woesner and M. Maier
"Routing and Resilience in AWG Based WDM Networks", *TransiNet Workshop*, Berlin, Oct. 2002

- T. Fischer, M. Maier, and H. Woesner
"Switching Optical Packets – Motivation and Prospects", *TransiNet Workshop*, Berlin, Oct. 2002

Appendix B

Acronyms

ACK	Acknowledgement
ADM	Add–Drop Multiplexer
AOTF	Acousto–Optic Tunable Filter
APS	Automatic Protection Switching
ASE	Amplified Spontaneous Emission
ASK	Amplitude Shift Keying
ATM	Asynchronous Transfer Mode
AWG	Arrayed–Waveguide Grating
BER	Bit Error Rate
BPF	Bandpass Filter
CDF	Cumulative Distribution Function
CDMA	Code Division Multiple Access
CPM	Crossphase Modulation
CSMA	Carrier Sense Multiple Access
CSMA/CA	Carrier Sense Multiple Access with Collision Avoidance
DBR	Distributed Bragg Reflector
DD	Direct Detection
DEMUX	Demultiplexer
DFB	Distributed Feedback
DSF	Dispersion–Shifted Fiber
DSSS	Direct Sequence Spread Spectrum
DQDB	Distributed Queue Dual Bus
DSL	Digital Subscriber Loop
DWDM	Dense Wavelength Division Multiplexing
EDFA	Erbium Doped Fiber Amplifier
EFM	Ethernet in the First Mile
EOTF	Electro–Optic Tunable Filter
ESCON	Enterprise System Connection
FBG	Fiber Bragg Grating
FCFS	First Come First Served
FDDI	Fibre Distributed Data Interface
FDL	Fiber Delay Line
FEC	Forward Error Correction
FIFO	First In First Out

FR	Frame Relay
FSK	Frequency Shift Keying
FSR	Free Spectral Range
FTTC	Fiber To The Curb
FTTH	Fiber To The Home
FTTx	Fiber To The x
FWHM	Full Width at Half Maximum
FWM	Four–Wave Mixing
GbE	Gigabit Ethernet
HFC	Hybrid Fiber Coax
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IM	Intensity Modulation
IP	Internet Protocol
ISI	Intersymbol Interference
LAN	Local Area Network
LC	Liquid Crystal
LD	Laser Diode
LDA	Laser Diode Amplifier
LED	Light Emitting Diode
LOS	Loss Of Signal
MAC	Medium Access Control
MAN	Metropolitan Area Network
MEF	Metro Ethernet Forum
MEMS	Micro Electro–Mechanical System
MPLS	Multi–Protocol Label Switching
MPλS	Multi–Protocol Lambda Switching
MUX	Multiplexer
MZI	Mach Zehnder Interferometer
NRZ	Nonreturn To Zero
NZ–DSF	Nonzero Dispersion Shifted Fiber
OADM	Optical Add–Drop Multiplexer
OAM	Operation Administration Maintenance
OBS	Optical Burst Switching
OC	Optical Carrier
OEO	Optical–Electronic–Optical
OFR	Optical Flow Routing
OLS	Optical Label Switching
OPS	Optical Packet Switching
OXC	Optical Cross–Connect
PD	Photodiode
PDM	Polarization Division Multiplexing
PHASAR	Phased Array
PMD	Polarization–Mode Dispersion
PON	Passive Optical Network
PoP	Point of Presence
PSC	Passive Star Coupler

PSR	Photonic Slot Routing
QoS	Quality of Service
R–ALOHA	Reservation ALOHA
RAM	Random Access Memory
RMS	Root Mean Square
RPR	Resilient Packet Ring
RPRWG	Resilient Packet Ring Working Group
SBS	Stimulated Brillouin Scattering
SCM	Subcarrier Multiplexing
SDH	Synchronous Digital Hierarchy
SG–DBR	Sampled Grating – Distributed Bragg Reflector
SMF	Single–Mode Fiber
SMSR	Side–Mode Suppression Ratio
SNR	Signal–to–Noise Ratio
SOA	Semiconductor Optical Amplifier
SONET	Synchronous Optical Network
SPM	Self–Phase Modulation
SRS	Stimulated Raman Scattering
TCP	Transmission Control Protocol
TDM	Time Division Multiplexing
TDMA	Time Division Multiple Access
TEC	Thermo–Electric Cooler
UMTS	Universal Mobile Telecommunications System
VOQ	Virtual Output Queue
WAN	Wide Area Network
WDM	Wavelength Division Multiplexing
WDMA	Wavelength Division Multiple Access
WG IPoRPR	Working Group IP over Resilient Packet Ring
WGR	Waveguide Grating Router
WLAN	Wireless Local Area Network
xDSL	x Digital Subscriber Loop
XGM	Cross–Gain Modulation
XPM	Crossphase Modulation

Appendix C

Spatial Wavelength Reuse: Refined Approximation of $P(Z = k)$

In this appendix we derive a refined approximation for the distribution of Z, the number of successful control packets destined to a given AWG output port in a given frame. This refined approximation does not approximate the binomial distributions of the random variables Y_i^n (see Eqn. (6.25)) and Y_i^o (see Eqn. (6.28)) with Poisson distributions.

C.1 Refined approximation

From (6.34) we note that $P(X_i = 1) = P(Y_i = 1)$. Recalling that $Y_i = Y_i^n + Y_i^o$, we have

$$P(Y_i = 1) = P(Y_i^n = 1,\ Y_i^o = 0) + P(Y_i^n = 0,\ Y_i^o = 1). \tag{C.1}$$

By the independence of the Y_i^n's and Y_i^o's we have

$$P(Y_i = 1) = P(Y_i^n = 1)P(Y_i^o = 0) + P(Y_i^n = 0)P(Y_i^o = 1). \tag{C.2}$$

Hence, with (6.25) and (6.28) we obtain

$$\begin{aligned}
P(Y_i = 1) &= \eta\frac{\sigma}{M}\left(1-\frac{\sigma}{M}\right)^{\eta-1}\left(1-\frac{p}{M}\right)^{S-\eta} + \left(1-\frac{\sigma}{M}\right)^{\eta}(S-\eta)\frac{p}{M}\left(1-\frac{p}{M}\right)^{S-\eta-1} && \text{(C.3)}\\
&= \frac{1}{M}\left(1-\frac{\sigma}{M}\right)^{\eta-1}\left(1-\frac{p}{M}\right)^{S-\eta-1}\left[\eta\sigma\left(1-\frac{p}{M}\right) + p(S-\eta)\left(1-\frac{\sigma}{M}\right)\right] && \text{(C.4)}\\
&\approx \frac{S}{M}\left(1-\frac{\sigma}{M}\right)^{\nu S-1}\left(1-\frac{p}{M}\right)^{S(1-\nu)-1}\cdot\\
&\quad \left[\nu\sigma\left(1-\frac{p}{M}\right) + p(1-\nu)\left(1-\frac{\sigma}{M}\right)\right] && \text{(C.5)}\\
&=: \kappa, && \text{(C.6)}
\end{aligned}$$

where (C.5) follows by approximating η/S by its expectation ν. Thus, the refined approximation of the distribution of Z is given by

$$P(Z = k) = \binom{M}{k}\left(\frac{\kappa}{D}\right)^k\left(1-\frac{\kappa}{D}\right)^{M-k},\ k = 0, 1, \ldots, M. \tag{C.7}$$

Table C.1: Mean aggregate throughput TH_{net} obtained with unrefined approximation, refined approximation, and simulation for default network parameters.

σ	0.02	0.04	0.1	0.2	0.5	1.0
TH_{net}, unrefined	0.886	1.77	4.29	7.32	8.45	8.10
TH_{net}, refined	0.888	1.77	4.29	7.37	8.52	8.16
TH_{net}, simul.	0.883	1.77	4.29	7.32	8.48	8.14

Table C.2: Mean aggregate throughput TH_{net} obtained with unrefined approximation, refined approximation, and simulation for $M = 8$.

σ	0.02	0.04	0.1	0.2	0.5	1.0
TH_{net}, unrefined	0.966	1.90	0.331	0.285	0.266	0.260
TH_{net}, refined	0.966	1.90	0.274	0.241	0.226	0.221
TH_{net}, simul.	0.966	1.89	0.272	0.242	0.224	0.223

C.2 Numerical evaluation

In this section we evaluate the refined approximation for $P(Z = k)$ and compare it with the approximation (6.36), which we henceforth refer to as "unrefined". First, note that in the unrefined approximation, the binomial distribution $BIN(\eta,\ \sigma/M)$ is approximated by the Poisson distribution with parameter $\eta\sigma/M$ (and the $BIN((S-\eta),\ p/M)$ distribution is approximated by the Poisson distribution with parameter $(S-\eta)p/M$). Clearly, this approximation is accurate when (i) S is large, (ii) σ is small (and also p), and (iii) M is large. We also note that the evaluation of the refined approximation is computationally slightly more demanding, as it involves the evaluation of the expression in (C.5) as compared to the single exponential term in (6.36).

In Table C.1 we compare the mean aggregate throughput TH_{net} obtained with unrefined approximation, refined approximation, and simulation for the typical network parameters chosen as default parameters in Section 6.2.3 (i.e., $R = 2$, $q = 0.25$, $M = 30$, $D = 4$, $N = 200$ (and thus $S = 50$), $p = 0.8$, $F = 200$, and $K = 170$). We observe that (i) the throughput obtained by both analytical approximations almost coincides, and (ii) the analytical results match very well with the simulation results. (Similar observations hold for the average delay.) In fact, we found in our extensive numerical investigations that these two observations hold for all parameter values considered in Section 6.2.3. In particular, both analytical approximations essentially coincide and match very well with the simulations for S values as small as 10. We also found that both analytical approximations essentially coincide and match very well with the simulations for M values as small as 15. Table C.2 compares the throughput obtained with unrefined approximation, refined approximation, and simulation for $M = 8$ and $K = 192$ and all other network parameters at their default values. For this small M value, we observe that unrefined approximation and refined approximation give identical throughput for small σ values ($\sigma = 0.02$, 0.04), but differ for larger σ values. We also observe that the refined approximation matches the simulation results very well. In summary, we find that the unrefined approximation gives accurate results for a wide range of network parameters. The refined approximation is more accurate when M is small and σ is large, at the expense of a slightly more demanding computation. However, small M values typically give poor network performance, as is illustrated

in Fig. 6.12, and may therefore not be desirable in practice.

Appendix D

Multicasting: Scheduling with $\Gamma \geq 1$ and $\left\lfloor \frac{M}{K} \right\rfloor + \left\lfloor \frac{F-M}{K} \right\rfloor + 1 = \left\lfloor \frac{F}{K} \right\rfloor$

In this appendix we analyze the packet scheduling with $\Gamma \geq 1$ occupied columns if

$$\left\lfloor \frac{M}{K} \right\rfloor + \left\lfloor \frac{F-M}{K} \right\rfloor + 1 = \left\lfloor \frac{F}{K} \right\rfloor . \tag{D.1}$$

If (D.1) holds, as opposed to (6.91), then there are situations where the feasibility of a scheduling pattern depends not only on the number of scheduled packets of the different types, but also the order in which these packets are scheduled. Specifically, if (D.1) holds and $\Gamma \geq 1$ then the number of (s,a) and $(s,1)$ packets that can be scheduled depends on the specific order in which these packets appear in the control slots. Consider an example with $F = 200$, $M = 30$, $K = 20$, $R = 5$ and $\Gamma = 1$. In this example, $\lfloor \frac{M}{K} \rfloor + \lfloor \frac{F-M}{K} \rfloor = 9$ and $\lfloor \frac{F}{K} \rfloor = 10$. Now suppose that for a given scheduling window of a given port we first have 15 $(s,1)$–packets to schedule and then 6 (s,a)–packets. Since we can reasonably ignore potential receiver collisions due to $(s,1)$–packets from other ports and $(s,1)$–packets from the considered port in our analytical model (see Section 6.2.3), the 15 $(s,1)$–packets are scheduled in the first 60 slots of the considered scheduling window. The 6 (s,a)–packets are scheduled in the 120 subsequent slots. Thus, there is room for one more (s,a)–packet in this scenario. Next consider a scenario where we first have the 6 (s,a)–packets to schedule and then the 15 $(s,1)$–packets. To avoid receiver collisions of the (s,a)–packets with the $(s,1)$–packets from the other ports (which occupy the first column, i.e., slots 31 through 50, in the considered scheduling window), one (s,a)–packet is scheduled in slots 1 through 20 and the other five (s,a)–packets are scheduled in slots 51 through 150. Then, following the adopted first–come–first–served and first–fit scheduling policy, five $(s,1)$–packets are scheduled in slots 21 through 40 and the remaining ten $(s,1)$–packets are scheduled in slots 151 through 190. Thus, there is no room for any additional packet. Note that if we had $F = 199$ and all other parameters as given in the example, then (6.91) would hold and we would not have the situation where the feasibility of a scheduling pattern depends on the order of the packets.

More generally, whenever (D.1) holds, $\Gamma \geq 1$, and an (s,a)–packet is scheduled in slot $M+K\cdot\Gamma+1$ and onwards, then there are slots "wasted" due to the packet ordering, as illustrated in the second scenario in the above example. Formally, we add an indicator e to the scheduling pattern to capture this effect. The indicator e is set to one if an (s,a)–packet is scheduled in slot $M + K \cdot \Gamma + 1$ and onwards; e is set to zero otherwise. The scheduling pattern $(\Gamma,\ i,\ b,\ c,\ d,\ e)$

indicates that given Γ occupied columns, the control slots up to index $i,\ i = 1, \ldots, M$, have been scanned and b $(l,1)$–packets, c (s,a)– packets, and d $(s,1)$–packets have been scheduled. Also, if $e = 1$ then an (s,a)–packet is scheduled starting in slot $M + K \cdot \Gamma + 1$, and if $e = 0$ this is not the case. For the case considered in this appendix, the first schedulability condition is given by (6.82) and the second original schedulability condition (6.83) is not considered. The third schedulability is as given by (6.84). The fourth schedulability condition is replaced by

$$a = 0,\ b = 0,\ 0 \leq c \leq \left\lfloor \frac{F}{K} \right\rfloor - \Gamma, 0 \leq d \leq \left(\left\lfloor \frac{F}{K} \right\rfloor - c \right) \cdot R + (D-1) \cdot \left\lfloor \frac{F-M}{K} \right\rfloor \cdot R, e = 0 \quad \text{(D.2)}$$

and

$$a = 0,\ b = 0,\ 1 \leq c \leq \left\lfloor \frac{M}{K} \right\rfloor + \left\lfloor \frac{F-M}{K} \right\rfloor - \Gamma,$$

$$\max\left\{0,\ \left(\left\lfloor \frac{M}{K} \right\rfloor - 1 \right) \cdot R + 1 \right\} \leq d \leq \left(\left\lfloor \frac{M}{K} \right\rfloor + \left\lfloor \frac{F-M}{K} \right\rfloor - c \right) \cdot R + (D-1) \cdot \left\lfloor \frac{F-M}{K} \right\rfloor \cdot R, e = 1. \quad \text{(D.3)}$$

We let ${}^{\Gamma}Q^{i}_{b,c,d,e}$ denote the probability that the scheduling pattern $(\Gamma,\ i,\ b,\ c,\ d,\ e)$ arises. For all feasible scheduling patterns we calculate ${}^{\Gamma}Q^{i}_{b,c,d,e}$ with the recursion

$$\begin{aligned} {}^{\Gamma}Q^{i}_{b,c,d,e} &= {}^{\Gamma}Q^{i-1}_{b,c,d,e} \cdot \left\{ \left(1 - \frac{\kappa}{D}\right) + \frac{\kappa}{D}(\hat{p}_{l,a} + \beta + \gamma + \delta) \right\} + & \text{(D.4)} \\ & {}^{\Gamma}Q^{i-1}_{b-1,c,d,e} \cdot \frac{\kappa \hat{p}_{l,1}}{D} + {}^{\Gamma}Q^{i-1}_{b,c-1,d,e} \cdot \frac{\kappa \hat{p}_{s,a}}{D} + & \text{(D.5)} \\ & {}^{\Gamma}Q^{i-1}_{b,c,d-1,e} \cdot \frac{\kappa \hat{p}_{s,1}}{D} + {}^{\Gamma}Q^{i-1}_{b,c-1,d,e-1} \cdot \frac{\kappa \hat{p}_{s,a}}{D}, & \text{(D.6)} \end{aligned}$$

where

$$\beta = \begin{cases} 0 & \text{if } (\Gamma,\ i,\ b+1,\ c,\ d,\ e) \text{ is feasible,} \\ \hat{p}_{l,1} & \text{otherwise.} \end{cases} \quad \text{(D.7)}$$

$$\gamma = \begin{cases} 0 & \text{if } (\Gamma,\ i,\ b,\ c+1,\ d,\ e) \text{ or } (\Gamma,\ i,\ b,\ c+1,\ d,\ e+1) \text{ is feasible,} \\ \hat{p}_{s,a} & \text{otherwise.} \end{cases} \quad \text{(D.8)}$$

$$\delta = \begin{cases} 0 & \text{if } (\Gamma,\ i,\ b,\ c,\ d+1,\ e) \text{ is feasible,} \\ \hat{p}_{s,1} & \text{otherwise.} \end{cases} \quad \text{(D.9)}$$

We initialize this recursion with ${}^{\Gamma}Q^{0}_{b,c,d,e} = 1$ if $b = c = d = e = 0$, and ${}^{\Gamma}Q^{0}_{b,c,d,e} = 0$ otherwise, and note that all undefined ${}^{\Gamma}Q^{i}_{b,c,d,e}$ (e.g., those with negative $b,\ c,\ d$, or e) are set to zero.

Appendix E

Network Dimensioning and Reconfiguration: Pareto–Optimal Solutions

In this appendix we provide tables which contain the Pareto–optimal solutions of network dimensioning and reconfiguration for different traffic loads and packet size distributions.

Table E.1: Network Frontier: Pareto–Optimal Solutions with σ and q as free decision variables

D	F	M	p	σ	q	TH	$Delay$	D	F	M	p	σ	q	TH	$Delay$
4	17	16	0.90	0.10	0.10	0.7145	115.8411	2	225	57	1.00	0.15	0.20	10.3502	916.6104
2	32	30	0.90	0.10	0.10	1.4367	120.0612	2	228	58	1.00	0.15	0.20	10.3625	920.9398
2	38	37	0.80	0.10	0.15	1.6093	129.9985	2	225	56	1.00	0.15	0.20	10.3716	924.8782
2	41	36	0.90	0.10	0.10	1.9752	132.7913	2	228	57	1.00	0.15	0.20	10.3848	928.8318
2	37	31	0.90	0.10	0.10	2.2737	134.4635	2	225	55	1.00	0.15	0.20	10.3913	933.6846
4	23	21	1.00	0.10	0.55	2.7933	142.1494	2	228	56	1.00	0.15	0.20	10.4055	937.2099
2	50	39	0.80	0.10	0.15	3.1624	165.6575	2	225	54	1.00	0.15	0.20	10.4092	943.0707
2	50	36	0.80	0.10	0.15	3.6121	174.1727	2	228	55	1.00	0.15	0.20	10.4245	946.1338
2	44	43	1.00	0.10	0.55	5.0891	176.7281	2	228	54	1.00	0.15	0.20	10.4417	955.6449
2	44	39	1.00	0.10	0.55	5.4073	186.3178	2	251	65	1.00	0.15	0.20	10.4474	964.9112
2	44	38	1.00	0.10	0.55	5.4824	189.3144	2	251	64	1.00	0.15	0.20	10.4730	970.8729
2	44	37	1.00	0.10	0.55	5.5552	192.6479	2	246	61	1.00	0.15	0.20	10.4924	970.8972
2	44	34	1.00	0.10	0.55	5.7556	205.2983	2	238	56	1.00	0.15	0.20	10.5124	978.3156
2	44	33	1.00	0.10	0.55	5.8142	210.7215	2	251	61	1.00	0.15	0.20	10.5442	990.6309
2	54	42	1.00	0.10	0.55	5.8919	219.4786	2	246	58	1.00	0.15	0.20	10.5563	993.6456
2	60	45	1.00	0.10	0.55	6.0414	235.9437	2	251	60	1.00	0.15	0.20	10.5658	997.9322
2	54	35	1.00	0.10	0.55	6.2799	246.1695	2	246	57	1.00	0.15	0.20	10.5748	1002.1607
2	65	44	1.00	0.10	0.55	6.3295	258.2406	2	246	56	1.00	0.15	0.20	10.5917	1011.2001
2	65	40	1.00	0.10	0.55	6.5221	271.2333	2	251	58	1.00	0.15	0.20	10.6051	1013.8417
2	65	38	1.00	0.10	0.55	6.6082	279.7183	2	246	55	1.00	0.15	0.20	10.6069	1020.8285
2	65	36	1.00	0.10	0.55	6.6843	290.1489	2	251	57	1.00	0.15	0.20	10.6227	1022.5298
2	65	34	1.00	0.10	0.55	6.7467	303.3178	2	251	56	1.00	0.15	0.20	10.6387	1031.7530
2	71	38	1.00	0.10	0.55	6.8075	305.5385	2	251	55	1.00	0.15	0.20	10.6530	1041.5771
2	80	45	1.00	0.10	0.55	6.8106	314.6119	2	264	61	1.00	0.15	0.20	10.6695	1041.9385
2	80	43	1.00	0.10	0.55	6.8866	321.4020	2	260	57	1.00	0.15	0.20	10.7043	1059.1942
2	80	42	1.00	0.10	0.55	6.9228	325.1686	2	264	58	1.00	0.15	0.20	10.7236	1066.3514
2	80	41	1.00	0.10	0.55	6.9575	329.3052	2	261	56	1.00	0.15	0.20	10.7272	1072.8587
2	80	39	1.00	0.10	0.55	7.0200	339.2137	2	278	64	1.00	0.15	0.20	10.7336	1075.3094
2	80	37	1.00	0.10	0.55	7.0753	350.7444	2	264	57	1.00	0.15	0.20	10.7388	1075.4895
2	80	36	1.00	0.10	0.55	7.0987	357.5131	2	264	56	1.00	0.15	0.20	10.7525	1085.1904
2	92	52	1.00	0.15	0.20	7.1606	365.0482	2	283	64	1.00	0.15	0.20	10.7764	1094.6495
2	98	56	1.00	0.15	0.20	7.1768	371.8966	2	278	61	1.00	0.15	0.20	10.7913	1097.1928
2	98	55	1.00	0.15	0.20	7.2659	375.7193	2	270	56	1.00	0.15	0.20	10.8013	1109.8538
2	98	54	1.00	0.15	0.20	7.3534	379.7875	2	283	61	1.00	0.15	0.20	10.8319	1116.9265
2	98	53	1.00	0.15	0.20	7.4390	384.1617	2	278	58	1.00	0.15	0.20	10.8387	1122.9003
2	98	52	1.00	0.15	0.20	7.5224	388.8557	2	278	57	1.00	0.15	0.20	10.8517	1132.5230
2	104	56	1.00	0.15	0.20	7.5254	394.6658	2	294	64	1.00	0.15	0.20	10.8655	1137.1977
2	104	55	1.00	0.15	0.20	7.6075	398.7226	2	283	58	1.00	0.15	0.20	10.8771	1143.0964
2	98	50	1.00	0.15	0.20	7.6822	399.3170	2	296	64	1.00	0.15	0.20	10.8810	1144.9337
2	104	54	1.00	0.15	0.20	7.6879	403.0398	2	296	61	1.00	0.15	0.20	10.9310	1168.2340
2	104	52	1.00	0.15	0.20	7.8425	412.6632	2	309	66	1.00	0.15	0.20	10.9420	1180.8834
2	116	58	1.00	0.15	0.20	7.9697	431.9546	2	294	58	1.00	0.15	0.20	10.9569	1187.5277
2	128	65	1.00	0.15	0.20	7.9895	451.8270	2	296	58	1.00	0.15	0.20	10.9708	1195.6061
2	128	64	1.00	0.15	0.20	8.0609	454.8187	2	315	66	1.00	0.15	0.20	10.9849	1203.8132
2	104	46	1.00	0.15	0.20	8.0777	485.1692	2	152	78	0.70	0.60	0.05	11.0368	1208.9706
2	116	54	1.00	0.15	0.20	8.0850	486.2053	2	152	77	0.70	0.60	0.05	11.1314	1214.0486
2	104	45	1.00	0.15	0.20	8.1304	494.0690	2	152	76	0.70	0.60	0.05	11.2224	1219.4886

Table E.1: *continued*

D	F	M	p	σ	q	TH	$Delay$	D	F	M	p	σ	q	TH	$Delay$
2	104	44	1.00	0.15	0.20	8.1780	503.9025	2	217	51	0.90	0.20	0.10	11.2772	1298.4444
2	118	52	1.00	0.15	0.20	8.2916	505.5001	2	364	66	1.00	0.15	0.20	11.2827	1391.0730
2	116	49	1.00	0.15	0.20	8.3889	516.3276	4	179	77	0.70	0.60	0.05	11.3135	1393.8773
2	128	56	1.00	0.15	0.20	8.4176	526.1529	4	173	90	0.70	0.85	0.05	11.4318	1408.7031
2	128	55	1.00	0.15	0.20	8.4770	531.1628	4	179	90	0.70	0.85	0.05	11.8062	1457.5599
2	128	54	1.00	0.15	0.20	8.5348	536.5024	4	173	87	0.70	0.95	0.05	11.9245	1478.9428
2	128	52	1.00	0.15	0.20	8.6443	548.3391	4	194	97	0.70	0.85	0.05	12.1860	1534.6495
2	135	56	1.00	0.15	0.20	8.6526	554.9269	4	194	101	0.70	0.95	0.05	12.2049	1565.8844
2	135	55	1.00	0.15	0.20	8.7073	560.2108	4	199	73	0.70	0.60	0.05	12.2457	1586.8121
2	128	50	1.00	0.15	0.20	8.7443	562.0659	4	199	104	0.70	0.95	0.05	12.2805	1589.9627
2	128	49	1.00	0.15	0.20	8.7903	569.7408	4	199	102	0.70	0.95	0.05	12.4317	1600.6828
2	135	50	1.00	0.15	0.20	8.9504	592.8039	4	199	101	0.70	0.95	0.05	12.5050	1606.2423
2	153	58	1.00	0.15	0.20	9.0646	617.9991	4	199	90	0.70	0.85	0.05	12.5746	1664.7491
2	156	58	1.00	0.15	0.20	9.1405	630.1167	4	194	87	0.70	0.95	0.05	12.7661	1703.6212
2	156	56	1.00	0.15	0.20	9.2311	641.2489	4	199	87	0.70	0.95	0.05	13.0042	1747.5290
2	156	55	1.00	0.15	0.20	9.2740	647.3547	4	194	81	0.70	0.95	0.05	13.0629	1751.2576
2	156	54	1.00	0.15	0.20	9.3151	653.8623	4	194	74	0.70	0.85	0.05	13.0660	1756.5210
2	161	56	1.00	0.15	0.20	9.3466	661.8017	4	199	81	0.70	0.95	0.05	13.2787	1796.3931
2	165	58	1.00	0.15	0.20	9.3516	666.4696	4	223	90	0.70	0.85	0.05	13.5941	1865.5228
2	156	52	1.00	0.15	0.20	9.3913	668.2883	4	241	97	0.70	0.85	0.05	13.9408	1965.3433
2	165	57	1.00	0.15	0.20	9.3935	672.1809	4	227	85	0.70	0.95	0.05	14.2085	2010.8314
2	165	56	1.00	0.15	0.20	9.4340	678.2440	4	266	97	0.70	0.85	0.05	14.7519	2169.2170
2	165	55	1.00	0.15	0.20	9.4727	684.7021	4	266	95	0.70	0.85	0.05	14.8194	2184.1303
2	171	58	1.00	0.15	0.20	9.4800	690.7049	4	266	102	0.70	0.95	0.05	14.8892	2215.4511
2	165	54	1.00	0.15	0.20	9.5097	691.5852	4	266	90	0.70	0.85	0.05	14.9606	2225.2425
2	171	57	1.00	0.15	0.20	9.5194	696.6239	4	266	87	0.70	0.95	0.05	15.3310	2335.8930
2	173	58	1.00	0.15	0.20	9.5208	698.7833	4	266	81	0.70	0.95	0.05	15.3874	2401.2089
2	172	57	1.00	0.15	0.20	9.5395	700.6977	4	302	97	0.70	0.85	0.05	15.6840	2462.7953
2	171	56	1.00	0.15	0.20	9.5573	702.9074	4	302	95	0.70	0.85	0.05	15.7265	2479.7269
2	173	57	1.00	0.15	0.20	9.5594	704.7715	4	302	106	0.70	0.95	0.05	15.7833	2487.2202
2	165	52	1.00	0.15	0.20	9.5776	706.8434	4	302	105	0.70	0.95	0.05	15.8159	2493.9754
2	171	55	1.00	0.15	0.20	9.5936	709.6003	4	302	104	0.70	0.95	0.05	15.8473	2500.9007
2	173	56	1.00	0.15	0.20	9.5965	711.1286	4	302	102	0.70	0.95	0.05	15.9062	2515.2866
2	172	55	1.00	0.15	0.20	9.6129	713.7500	4	302	101	0.70	0.95	0.05	15.9337	2522.7621
2	171	54	1.00	0.15	0.20	9.6281	716.7337	4	310	106	0.70	0.95	0.05	15.9870	2553.1068
2	173	55	1.00	0.15	0.20	9.6320	717.8997	4	310	104	0.70	0.95	0.05	16.0461	2567.1497
2	172	54	1.00	0.15	0.20	9.6470	720.9251	4	310	102	0.70	0.95	0.05	16.1001	2581.9167
2	173	54	1.00	0.15	0.20	9.6657	725.1166	4	302	87	0.70	0.95	0.05	16.1548	2652.0289
2	171	52	1.00	0.15	0.20	9.6909	732.5468	4	330	97	0.70	0.85	0.05	16.2684	2691.1339
2	172	52	1.00	0.15	0.20	9.7090	736.8307	4	330	95	0.70	0.85	0.05	16.2952	2709.6354
2	173	52	1.00	0.15	0.20	9.7269	741.1146	4	330	106	0.70	0.95	0.05	16.4532	2717.8234
2	171	50	1.00	0.15	0.20	9.7438	750.8849	4	330	105	0.70	0.95	0.05	16.4777	2725.2049
2	172	50	1.00	0.15	0.20	9.7611	755.2760	4	330	104	0.70	0.95	0.05	16.5011	2732.7723
2	173	50	1.00	0.15	0.20	9.7782	759.6672	4	330	102	0.70	0.95	0.05	16.5438	2748.4920
2	190	57	1.00	0.15	0.20	9.8655	774.0265	4	330	101	0.70	0.95	0.05	16.5633	2756.6606
2	190	56	1.00	0.15	0.20	9.8966	781.0082	4	343	106	0.70	0.95	0.05	16.7271	2824.8892
2	190	55	1.00	0.15	0.20	9.9260	788.4448	4	343	104	0.70	0.95	0.05	16.7683	2840.4270
2	190	54	1.00	0.15	0.20	9.9535	796.3708	4	343	103	0.70	0.95	0.05	16.7871	2848.4926
2	190	52	1.00	0.15	0.20	10.0024	813.9409	4	359	95	0.70	0.85	0.05	16.7907	2947.7549
2	197	55	1.00	0.15	0.20	10.0323	817.4928	4	359	106	0.70	0.95	0.05	17.0369	2956.6624
2	203	57	1.00	0.15	0.20	10.0651	826.9863	4	359	105	0.70	0.95	0.05	17.0544	2964.6926
2	203	56	1.00	0.15	0.20	10.0921	834.4456	4	359	104	0.70	0.95	0.05	17.0707	2972.9250
2	203	55	1.00	0.15	0.20	10.1175	842.3910	4	359	103	0.70	0.95	0.05	17.0857	2981.3669
2	210	58	1.00	0.15	0.20	10.1356	848.2341	4	359	102	0.70	0.95	0.05	17.0994	2990.0261
2	203	54	1.00	0.15	0.20	10.1411	850.8593	4	381	97	0.70	0.85	0.05	17.1121	3107.0365
2	210	57	1.00	0.15	0.20	10.1623	855.5030	4	359	87	0.70	0.95	0.05	17.1212	3152.5775
2	210	56	1.00	0.15	0.20	10.1874	863.2196	4	386	95	0.70	0.85	0.05	17.1852	3169.4523
2	211	56	1.00	0.15	0.20	10.2005	867.3302	4	386	106	0.70	0.95	0.05	17.5015	3179.0298
2	210	55	1.00	0.15	0.20	10.2109	871.4390	4	386	105	0.70	0.95	0.05	17.5134	3187.6639
2	217	58	1.00	0.15	0.20	10.2283	876.5085	4	386	104	0.70	0.95	0.05	17.5241	3196.5155
2	210	54	1.00	0.15	0.20	10.2325	880.1993	4	386	103	0.70	0.95	0.05	17.5335	3205.5923
2	217	57	1.00	0.15	0.20	10.2532	884.0198	4	386	102	0.70	0.95	0.05	17.5416	3214.9027
2	220	58	1.00	0.15	0.20	10.2662	888.6262	4	399	106	0.70	0.95	0.05	17.7028	3286.0956
2	217	56	1.00	0.15	0.20	10.2765	891.9936	4	400	106	0.70	0.95	0.05	17.7177	3294.3314
2	220	57	1.00	0.15	0.20	10.2904	896.2413	4	400	105	0.70	0.95	0.05	17.7270	3303.2787
2	217	55	1.00	0.15	0.20	10.2982	900.4870	4	400	104	0.70	0.95	0.05	17.7351	3312.4513
2	220	56	1.00	0.15	0.20	10.3130	904.3253	4	400	102	0.70	0.95	0.05	17.7474	3331.5054
2	225	58	1.00	0.15	0.20	10.3272	908.8222	4	400	101	0.70	0.95	0.05	17.7515	3341.4067
2	220	55	1.00	0.15	0.20	10.3339	912.9361								

Table E.2: Pareto–Optimal Solutions for $\sigma = 0.1$ and $q = 0.1$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
2	40	39	0.90	1.1618	123.5485	2	218	39	0.90	7.9365	685.1076
2	37	34	0.75	1.6017	133.1404	2	220	39	0.90	7.9503	691.3930
2	44	38	0.75	2.0881	146.6693	2	218	38	0.90	7.9585	694.9413
2	59	49	0.75	2.4125	173.2367	2	226	40	0.90	7.9684	700.8297
2	59	48	0.75	2.5552	174.6934	2	220	38	0.90	7.9720	701.3169
2	59	47	0.75	2.6973	176.2499	2	226	39	0.90	7.9904	710.2492
2	60	46	0.75	2.9502	180.9531	2	218	36	0.90	7.9967	717.9251
2	59	44	0.75	3.1199	181.6177	2	226	38	0.90	8.0110	720.4437
2	59	42	0.75	3.3978	185.8642	2	226	36	0.90	8.0461	744.2710
2	59	39	0.75	3.8068	193.6569	2	260	45	0.90	8.0644	763.7058
2	59	37	0.75	4.0727	200.0694	2	260	44	0.90	8.0866	770.9984
2	59	34	0.75	4.4570	212.3851	2	267	45	0.90	8.1034	784.2671
2	69	39	0.90	4.6591	213.2121	2	260	42	0.90	8.1284	787.2516
2	69	38	0.90	4.7726	216.3246	2	260	41	0.90	8.1477	796.3622
2	59	31	0.75	4.8148	229.8407	2	260	40	0.90	8.1658	806.2642
2	65	34	0.75	4.9003	233.9836	2	260	39	0.90	8.1825	817.1008
2	59	30	0.75	4.9253	237.3986	2	267	41	0.90	8.1831	817.8027
2	99	52	0.90	5.0858	270.7719	2	267	40	0.90	8.2003	827.9714
2	68	28	0.90	5.6910	280.1374	2	267	39	0.90	8.2160	839.0997
2	111	50	0.90	5.7112	313.2584	2	280	42	0.90	8.2272	847.8094
2	111	49	0.90	5.7835	315.4947	2	267	38	0.90	8.2303	851.1437
2	111	48	0.90	5.8552	317.8999	2	280	41	0.90	8.2440	857.6208
2	111	47	0.90	5.9265	320.3998	2	286	42	0.90	8.2541	865.9768
2	84	32	0.75	6.0170	322.7897	2	280	40	0.90	8.2596	868.2846
2	111	45	0.90	6.0670	326.0436	2	285	41	0.90	8.2660	872.9355
2	112	45	0.90	6.0981	328.9809	2	286	41	0.90	8.2703	875.9984
2	111	44	0.90	6.1362	329.1570	2	280	39	0.90	8.2737	879.9547
2	111	42	0.90	6.2722	336.0959	2	285	40	0.90	8.2810	883.7897
2	112	42	0.90	6.3012	339.1238	2	286	40	0.90	8.2852	886.8907
2	111	41	0.90	6.3388	339.9854	2	280	38	0.90	8.2864	892.5851
2	112	41	0.90	6.3670	343.0483	2	285	39	0.90	8.2945	895.6682
2	111	40	0.90	6.4042	344.2128	2	286	39	0.90	8.2986	898.8109
2	112	40	0.90	6.4316	347.3138	2	297	42	0.90	8.3007	899.2836
2	111	39	0.90	6.4683	348.8392	2	285	38	0.90	8.3067	908.5241
2	112	39	0.90	6.4950	351.9819	2	297	41	0.90	8.3156	909.6907
2	111	38	0.90	6.5310	353.8462	2	297	40	0.90	8.3293	921.0018
2	112	38	0.90	6.5570	357.0340	2	297	39	0.90	8.3416	933.3806
2	111	36	0.90	6.6510	365.5490	2	311	42	0.90	8.3552	941.6740
2	112	36	0.90	6.6755	368.8423	2	309	41	0.90	8.3614	946.4459
2	111	34	0.90	6.7619	380.0412	2	322	44	0.90	8.3664	954.8519
2	112	34	0.90	6.7848	383.4650	2	309	39	0.90	8.3850	971.0929
2	133	42	0.90	6.8080	402.7095	2	322	42	0.90	8.3946	974.9808
2	133	41	0.90	6.8609	407.3699	2	322	41	0.90	8.4072	986.2640
2	130	39	0.90	6.9055	408.5504	2	322	40	0.90	8.4185	998.5273
2	133	40	0.90	6.9127	412.4352	2	322	39	0.90	8.4284	1011.9480
2	112	31	0.90	6.9232	413.5080	2	344	44	0.90	8.4414	1020.0902
2	133	39	0.90	6.9631	417.9785	2	344	42	0.90	8.4661	1041.5944
2	133	38	0.90	7.0121	423.9779	2	346	42	0.90	8.4721	1047.6502
2	133	36	0.90	7.1045	438.0002	2	344	41	0.90	8.4768	1053.6485
2	148	41	0.90	7.1279	453.3139	2	346	41	0.90	8.4827	1059.7743
2	133	34	0.90	7.1876	455.3646	2	344	40	0.90	8.4863	1066.7496
2	155	42	0.90	7.1917	469.3231	2	346	40	0.90	8.4920	1072.9516
2	155	40	0.90	7.2768	480.6575	2	366	45	0.90	8.4953	1075.0627
2	141	34	0.90	7.3094	482.7550	2	358	42	0.90	8.5069	1083.9849
2	155	39	0.90	7.3174	487.1178	2	366	44	0.90	8.5074	1085.3285
2	167	39	0.95	7.4903	515.8923	2	367	44	0.90	8.5102	1088.2939
2	155	34	0.90	7.4924	530.6881	2	366	42	0.90	8.5289	1108.2080
2	155	33	0.90	7.5194	542.6893	2	379	45	0.90	8.5316	1113.2480
2	218	52	0.90	7.5566	607.1496	2	380	45	0.90	8.5342	1116.1854
2	218	51	0.90	7.5898	611.0580	2	366	41	0.90	8.5380	1121.0330
2	218	50	0.90	7.6224	615.2282	2	379	44	0.90	8.5428	1123.8784
2	218	49	0.90	7.6547	619.6202	2	380	44	0.90	8.5454	1126.8438
2	218	48	0.90	7.6863	624.3439	2	366	40	0.90	8.5459	1134.9720
2	218	47	0.90	7.7176	629.2536	2	379	42	0.90	8.5626	1147.5706
2	220	47	0.90	7.7344	635.0265	2	380	42	0.90	8.5651	1150.5985
2	218	45	0.90	7.7777	640.3379	2	379	41	0.90	8.5709	1160.8511
2	220	45	0.90	7.7938	646.2126	2	380	41	0.90	8.5733	1163.9140
2	218	44	0.90	7.8067	646.4525	2	379	40	0.90	8.5779	1175.2852
2	218	42	0.90	7.8620	660.0802	2	380	40	0.90	8.5803	1178.3862
2	220	42	0.90	7.8770	666.1360	2	379	39	0.90	8.5835	1191.0816
2	218	41	0.90	7.8881	667.7191	2	380	39	0.90	8.5858	1194.2243
2	220	41	0.90	7.9027	673.8450	2	379	38	0.90	8.5876	1208.1777
2	218	40	0.90	7.9130	676.0216	2	380	38	0.90	8.5898	1211.3655
2	220	40	0.90	7.9272	682.2236	2	380	36	0.90	8.5918	1251.4291

Table E.3: Pareto–Optimal Solutions for $\sigma = 0.1$ and $q = 0.5$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
4	24	18	0.80	2.8935	172.7559	2	157	40	0.95	7.9155	635.5594
2	49	48	0.95	4.7014	181.5074	2	157	38	0.95	7.9282	655.2609
2	46	40	0.75	5.0460	202.5229	2	157	37	0.95	7.9309	666.6763
2	35	26	0.85	5.1285	227.9409	2	202	58	0.95	7.9743	702.9611
2	63	55	0.80	5.1828	235.9143	2	202	57	0.95	7.9906	706.5663
2	63	54	0.80	5.2502	237.3885	2	202	56	0.95	8.0065	710.4594
2	63	53	0.80	5.3172	238.9075	2	202	55	0.95	8.0221	714.5020
2	63	52	0.80	5.3836	240.5442	2	202	54	0.95	8.0372	718.7521
2	63	51	0.80	5.4496	242.2687	2	202	53	0.95	8.0518	723.2725
2	63	50	0.80	5.5149	244.0861	2	202	52	0.95	8.0659	728.0405
2	63	49	0.80	5.5796	246.0167	2	202	50	0.95	8.0926	738.3690
2	63	47	0.80	5.7070	250.2833	2	202	49	0.95	8.1050	744.0235
2	63	46	0.80	5.7695	252.5996	2	202	48	0.95	8.1167	750.0286
2	63	45	0.80	5.8311	255.1112	2	202	47	0.95	8.1275	756.5036
2	63	43	0.80	5.9512	260.6752	2	202	45	0.95	8.1464	770.8353
2	63	42	0.80	6.0095	263.8060	2	202	43	0.95	8.1612	787.3646
2	63	41	0.80	6.0662	267.2139	2	202	42	0.95	8.1666	796.6705
2	65	42	0.80	6.1017	272.2189	2	202	41	0.95	8.1705	806.8091
2	63	39	0.80	6.1752	274.8608	2	202	38	0.95	8.1713	843.0745
2	63	38	0.80	6.2269	279.2356	2	225	45	0.85	8.1931	893.2008
2	63	37	0.80	6.2763	284.0409	2	225	43	0.85	8.2015	912.4978
2	63	36	0.80	6.3232	289.3329	2	225	42	0.85	8.2036	923.4136
2	65	37	0.80	6.3566	293.0580	2	225	41	0.85	8.2044	935.1091
2	63	34	0.80	6.4073	301.8711	2	260	55	0.80	8.2212	976.1365
2	63	33	0.80	6.4434	309.3252	2	259	54	0.80	8.2265	978.2901
2	69	37	0.80	6.5032	311.0924	2	260	54	0.80	8.2302	982.0673
2	85	48	0.80	6.5424	334.7668	2	260	53	0.80	8.2387	988.3947
2	85	47	0.80	6.5855	337.7496	2	295	66	0.95	8.3150	991.8486
2	69	33	0.80	6.6401	339.2893	2	295	65	0.95	8.3261	995.4873
2	71	34	0.80	6.6719	340.7103	2	263	45	0.95	8.3833	1003.6123
2	87	44	0.85	6.7911	348.2774	2	295	59	0.95	8.3866	1021.5497
2	85	41	0.80	6.8211	361.1938	2	295	58	0.95	8.3957	1026.6016
2	85	38	0.80	6.9204	377.4029	2	295	57	0.95	8.4045	1031.8667
2	85	37	0.80	6.9493	383.8631	2	295	56	0.95	8.4127	1037.5520
2	109	54	0.95	6.9792	387.8415	2	295	55	0.95	8.4206	1043.4558
2	109	53	0.95	7.0145	390.2807	2	295	54	0.95	8.4281	1049.6628
2	109	52	0.95	7.0493	392.8535	2	295	53	0.95	8.4351	1056.2642
2	109	50	0.95	7.1174	398.4268	2	295	52	0.95	8.4416	1063.2274
2	109	49	0.95	7.1505	401.4780	2	295	50	0.95	8.4529	1078.3111
2	109	47	0.95	7.2145	408.2124	2	295	49	0.95	8.4577	1086.5690
2	109	45	0.95	7.2752	415.9458	2	309	55	0.95	8.4598	1092.9758
2	109	43	0.95	7.3318	424.8651	2	309	54	0.95	8.4666	1099.4773
2	109	42	0.95	7.3582	429.8866	2	309	53	0.95	8.4728	1106.3920
2	109	41	0.95	7.3832	435.3574	2	309	52	0.95	8.4785	1113.6856
2	110	41	0.95	7.3987	439.3515	2	309	50	0.95	8.4884	1129.4852
2	109	38	0.95	7.4475	454.9263	2	309	49	0.95	8.4924	1138.1350
2	109	37	0.95	7.4644	462.8517	2	309	47	0.95	8.4980	1157.2259
2	141	55	0.95	7.4751	498.7365	2	309	45	0.95	8.5000	1179.1490
2	141	54	0.95	7.5007	501.7032	2	360	58	0.85	8.5079	1301.0082
2	141	53	0.95	7.5258	504.8585	2	361	58	0.85	8.5100	1304.6221
2	141	52	0.95	7.5505	508.1867	2	360	57	0.85	8.5132	1307.9339
2	141	50	0.95	7.5981	515.3962	2	361	57	0.85	8.5152	1311.5670
2	141	49	0.95	7.6210	519.3432	2	360	56	0.85	8.5180	1315.2462
2	141	48	0.95	7.6432	523.5348	2	361	56	0.85	8.5200	1318.8996
2	141	47	0.95	7.6646	528.0545	2	360	55	0.85	8.5225	1322.8786
2	141	46	0.95	7.6851	532.8790	2	361	55	0.85	8.5244	1326.5533
2	141	45	0.95	7.7047	538.0583	2	360	54	0.85	8.5264	1331.0818
2	141	43	0.95	7.7406	549.5961	2	361	54	0.85	8.5283	1334.7792
2	141	42	0.95	7.7567	556.0918	2	360	53	0.85	8.5300	1339.5062
2	141	41	0.95	7.7713	563.1687	2	361	53	0.85	8.5318	1343.2270
2	141	40	0.95	7.7844	570.7890	2	360	52	0.85	8.5329	1348.5740
2	157	48	0.95	7.8030	582.9430	2	361	52	0.85	8.5347	1352.3200
2	157	47	0.95	7.8208	587.9756	2	360	50	0.85	8.5372	1367.9894
2	157	45	0.95	7.8538	599.1146	2	361	50	0.85	8.5390	1371.7893
2	157	43	0.95	7.8826	611.9616	2	361	49	0.85	8.5401	1382.4794
2	157	42	0.95	7.8951	619.1944	2	397	47	0.85	8.5938	1546.3247
2	157	41	0.95	7.9060	627.0744						

Table E.4: Pareto–Optimal Solutions for $\sigma = 0.1$ and $q = 0.9$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
4	24	21	1.00	4.2673	162.3399	2	175	55	1.00	8.0653	1053.2001
4	49	48	0.80	4.3202	282.2721	2	190	58	1.00	8.0905	1133.4556
4	49	46	0.80	4.3356	284.0520	2	199	76	1.00	8.0915	1148.9174
4	49	44	0.80	4.3505	286.3995	2	199	74	1.00	8.0949	1151.8905
4	40	22	0.80	4.3782	286.7716	2	199	73	1.00	8.0965	1153.4122
4	49	40	0.80	4.3785	291.4752	2	199	70	1.00	8.1005	1158.4507
4	51	35	1.00	4.4633	292.4994	2	199	68	1.00	8.1025	1162.1935
4	55	35	1.00	4.4872	315.4406	2	199	66	1.00	8.1041	1166.2122
2	46	40	0.75	7.3458	315.5049	2	211	78	1.00	8.1064	1215.3322
2	50	43	1.00	7.4989	318.8284	2	211	76	1.00	8.1098	1218.1989
2	52	43	1.00	7.5260	331.5815	2	211	74	1.00	8.1127	1221.3512
2	72	41	1.00	7.7049	466.5047	2	211	73	1.00	8.1140	1222.9647
2	82	59	0.90	7.7184	495.9218	2	211	70	1.00	8.1173	1228.3070
2	82	58	0.90	7.7233	497.2501	2	211	68	1.00	8.1188	1232.2755
2	82	57	0.90	7.7279	498.6781	2	211	66	1.00	8.1199	1236.5365
2	82	56	0.90	7.7322	500.1356	2	211	64	1.00	8.1201	1241.3024
2	82	51	0.90	7.7479	509.0388	2	213	68	1.00	8.1214	1243.9558
2	82	50	0.90	7.7496	511.1888	2	248	95	1.00	8.1246	1407.0085
2	82	49	0.90	7.7507	513.4691	2	248	83	1.00	8.1463	1420.8232
2	82	48	0.90	7.7512	515.9489	2	248	78	1.00	8.1529	1428.4473
2	89	66	1.00	7.7604	521.5716	2	248	76	1.00	8.1550	1431.8167
2	89	64	1.00	7.7719	523.5819	2	248	74	1.00	8.1566	1435.5218
2	89	62	1.00	7.7827	525.7972	2	248	73	1.00	8.1574	1437.4182
2	89	58	1.00	7.8014	530.9339	2	248	70	1.00	8.1589	1443.6974
2	89	51	1.00	7.8215	543.3548	2	248	68	1.00	8.1592	1448.3617
2	89	49	1.00	7.8229	548.1031	2	254	76	1.00	8.1610	1466.4574
2	99	64	1.00	7.8327	582.4113	2	254	74	1.00	8.1626	1470.2522
2	100	64	1.00	7.8382	588.2943	2	254	73	1.00	8.1632	1472.1945
2	100	62	1.00	7.8468	590.7834	2	254	70	1.00	8.1645	1478.6255
2	100	58	1.00	7.8612	596.5549	2	261	78	1.00	8.1660	1503.3256
2	100	51	1.00	7.8738	610.5110	2	261	77	1.00	8.1670	1505.0477
2	106	60	1.00	7.8829	629.1163	2	261	76	1.00	8.1678	1506.8716
2	109	51	1.00	7.9088	665.4576	2	261	74	1.00	8.1691	1510.7709
2	121	70	1.00	7.9101	704.3837	2	261	73	1.00	8.1697	1512.7668
2	121	68	1.00	7.9178	706.6594	2	261	72	1.00	8.1702	1514.8259
2	121	66	1.00	7.9248	709.1030	2	261	70	1.00	8.1707	1519.3751
2	121	65	1.00	7.9281	710.5023	2	287	74	1.00	8.1907	1661.2692
2	121	64	1.00	7.9312	711.8361	2	294	78	1.00	8.1943	1693.4012
2	121	62	1.00	7.9368	714.8479	2	294	76	1.00	8.1953	1697.3956
2	121	59	1.00	7.9435	719.9596	2	294	74	1.00	8.1959	1701.7880
2	121	58	1.00	7.9452	721.8323	2	294	73	1.00	8.1961	1704.0361
2	121	56	1.00	7.9477	725.9944	2	297	77	1.00	8.1971	1712.6405
2	121	54	1.00	7.9489	730.5792	2	303	78	1.00	8.2009	1745.2400
2	140	76	1.00	7.9561	808.2825	2	303	76	1.00	8.2018	1749.3567
2	140	74	1.00	7.9632	810.3741	2	303	74	1.00	8.2022	1753.8835
2	140	73	1.00	7.9666	811.4447	2	303	73	1.00	8.2023	1756.2005
2	140	70	1.00	7.9760	814.9894	2	306	78	1.00	8.2030	1762.5196
2	140	68	1.00	7.9817	817.6235	2	306	76	1.00	8.2038	1766.6770
2	140	66	1.00	7.9869	820.4508	2	306	74	1.00	8.2042	1771.2487
2	140	65	1.00	7.9891	822.0698	2	306	73	1.00	8.2043	1773.5886
2	140	64	1.00	7.9913	823.6130	2	309	78	1.00	8.2051	1779.7992
2	140	62	1.00	7.9950	827.0978	2	309	76	1.00	8.2058	1783.9974
2	140	59	1.00	7.9988	833.0111	2	309	74	1.00	8.2062	1788.6139
2	140	58	1.00	7.9995	835.1778	2	309	73	1.00	8.2062	1790.9767
2	140	56	1.00	8.0001	839.9935	2	314	82	1.00	8.2064	1800.7522
2	153	68	1.00	8.0164	893.5457	2	314	78	1.00	8.2085	1808.5986
2	155	70	1.00	8.0167	902.3109	2	314	76	1.00	8.2091	1812.8647
2	155	68	1.00	8.0212	905.2261	2	314	73	1.00	8.2094	1819.9570
2	155	66	1.00	8.0251	908.3562	2	314	74	1.00	8.2094	1817.5559
2	155	64	1.00	8.0284	911.8572	2	332	83	1.00	8.2177	1902.0698
2	155	59	1.00	8.0328	922.2623	2	332	82	1.00	8.2183	1903.9800
2	155	58	1.00	8.0330	924.6612	2	332	78	1.00	8.2198	1912.2762
2	166	73	1.00	8.0352	962.1428	2	332	76	1.00	8.2202	1916.7868
2	166	70	1.00	8.0418	966.3458	2	333	78	1.00	8.2204	1918.0361
2	166	68	1.00	8.0456	969.4679	2	333	76	1.00	8.2208	1922.5603
2	166	66	1.00	8.0488	972.8202	2	343	70	1.00	8.2245	1996.7266
2	166	64	1.00	8.0513	976.5697	2	355	82	1.00	8.2318	2035.8822
2	166	59	1.00	8.0539	987.7132	2	357	79	1.00	8.2337	2053.8106
2	171	66	1.00	8.0585	1002.1220	2	361	82	1.00	8.2350	2070.2915
2	171	64	1.00	8.0607	1005.9844	2	390	78	1.00	8.2492	2246.3486
2	171	58	1.00	8.0622	1020.1101						

Table E.5: Pareto–Optimal Solutions for $\sigma = 0.3$ and $q = 0.1$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
4	45	40	1.00	2.3723	338.7587	2	261	62	1.00	12.0755	2180.6202
2	79	78	1.00	2.3846	369.4111	2	261	61	1.00	12.0815	2193.8380
4	49	43	0.65	2.4049	399.2489	2	261	60	1.00	12.0822	2208.5330
2	92	87	0.95	3.2832	405.2141	2	284	73	0.95	12.0953	2284.3395
2	64	57	0.90	3.5741	412.0022	2	284	72	0.95	12.1290	2289.1439
2	81	70	1.00	4.5445	414.1696	2	284	71	0.95	12.1611	2294.4007
2	93	79	0.95	5.0265	437.3718	2	284	70	0.95	12.1915	2300.1592
2	92	77	0.95	5.2183	440.6751	2	284	68	0.95	12.2460	2313.3933
2	103	87	0.95	5.2786	455.1786	2	284	64	0.95	12.3243	2348.7182
2	103	85	0.95	5.6153	463.1780	2	284	63	0.95	12.3357	2359.9178
2	103	84	0.95	5.7852	466.5696	2	295	68	0.95	12.3715	2402.9965
2	103	83	0.95	5.9531	470.0984	2	293	64	0.95	12.4206	2423.1494
2	109	87	0.95	6.1609	487.9397	2	295	64	0.95	12.4412	2439.6897
2	103	79	0.95	6.5769	489.3956	2	295	63	0.95	12.4504	2451.3231
2	103	78	0.95	6.7329	493.6560	2	295	62	0.95	12.4556	2464.1731
2	109	83	0.95	6.7896	501.7065	2	313	72	0.95	12.4612	2522.8945
2	103	77	0.95	6.7996	510.5624	2	312	70	1.00	12.5107	2524.5831
2	103	76	0.95	6.9505	514.9480	2	311	68	0.95	12.5381	2533.3285
2	118	89	0.95	7.0692	521.5902	2	313	68	0.95	12.5577	2549.6200
2	92	64	0.95	7.1915	527.3516	2	315	68	0.95	12.5771	2565.9115
2	103	73	0.95	7.3851	529.4723	2	312	64	1.00	12.6076	2579.6491
2	103	72	0.95	7.5235	534.8245	2	313	64	0.95	12.6148	2588.5521
2	109	77	0.95	7.5693	540.3039	2	313	63	0.95	12.6208	2600.8953
2	103	71	0.95	7.6582	540.4673	2	315	64	0.95	12.6328	2605.0924
2	109	76	0.95	7.7068	544.9450	2	315	63	0.95	12.6385	2617.5144
2	103	70	0.95	7.7890	546.4250	2	335	70	0.95	12.7205	2713.2160
2	109	73	0.95	8.1012	560.3153	2	338	70	0.95	12.7467	2737.5134
2	141	96	0.95	8.5295	622.2562	2	353	70	0.95	12.8708	2859.0007
2	141	95	0.95	8.6511	625.0717	2	360	72	0.95	12.8860	2901.7317
2	109	64	0.95	8.7259	669.4365	2	360	71	0.95	12.9065	2908.3953
2	159	106	0.95	8.9770	674.9582	2	360	70	0.95	12.9251	2915.6948
2	114	65	0.95	9.0673	692.1787	2	361	70	0.95	12.9327	2923.7939
2	109	59	0.95	9.1208	717.1048	2	360	68	0.95	12.9562	2932.4703
2	141	85	0.95	9.1926	745.3059	2	361	68	0.95	12.9635	2940.6161
2	141	84	0.95	9.2996	748.5159	2	364	68	0.95	12.9853	2965.0533
2	141	83	0.95	9.4051	751.8767	2	360	64	0.95	12.9862	2977.2484
2	141	79	0.95	9.8101	767.0674	2	361	64	0.95	12.9931	2985.5185
2	141	78	0.95	9.9065	771.3613	2	370	70	0.95	12.9992	2996.6863
2	141	77	0.95	10.0008	775.8835	2	364	64	0.95	13.0134	3010.3289
2	141	76	0.95	10.0928	780.6497	2	370	68	0.95	13.0279	3013.9278
2	141	73	0.95	10.3533	796.6159	2	373	70	1.00	13.0284	3018.1715
2	141	72	0.95	10.4344	802.5597	2	373	68	1.00	13.0552	3035.9763
2	141	71	0.95	10.5122	808.8675	2	373	64	1.00	13.0745	3084.0035
2	159	83	0.95	10.6047	847.8610	2	389	72	0.90	13.0856	3139.7224
2	220	114	0.90	11.1053	1087.6350	2	389	71	0.90	13.1031	3146.8529
2	240	121	0.95	11.4684	1166.2799	2	389	70	0.90	13.1188	3154.6406
2	295	148	0.95	11.6569	1399.3660	2	389	68	0.90	13.1441	3172.4436
2	389	195	0.90	11.7478	1819.3841	2	395	71	0.95	13.1533	3191.1560
2	240	64	0.95	11.7494	1984.8323	2	395	70	0.95	13.1681	3199.1651
2	240	63	0.95	11.7715	1994.2967	2	395	68	0.95	13.1913	3217.5716
2	240	59	0.95	11.8166	2043.3565	2	400	71	1.00	13.1932	3228.3818
2	252	63	1.00	11.9452	2093.9313	2	400	70	1.00	13.2071	3236.6450
2	261	68	1.00	11.9556	2124.3695	2	399	68	0.95	13.2156	3250.1546
2	261	64	1.00	12.0499	2157.9757	2	400	68	1.00	13.2282	3255.7387
2	261	63	1.00	12.0648	2168.7146	2	400	64	1.00	13.2356	3307.2424

Table E.6: Pareto–Optimal Solutions for $\sigma = 0.3$ and $q = 0.5$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
8	17	16	0.95	2.5709	382.7429	2	123	67	0.95	8.7557	1470.3953
8	23	21	0.95	2.9794	409.8701	2	123	64	0.95	8.7974	1494.7925
4	36	29	0.95	4.9584	531.2120	2	123	62	0.95	8.8147	1513.4240
4	49	43	0.65	5.2394	592.3880	2	141	77	0.95	9.0127	1616.5634
4	59	55	0.95	5.4939	596.0768	2	165	96	0.95	9.0944	1803.0229
4	63	59	0.95	5.5325	623.0541	2	167	96	0.95	9.1391	1824.8777
4	63	58	0.95	5.6006	626.1531	2	165	89	0.95	9.2716	1829.3365
4	63	57	0.95	5.6677	629.4164	2	165	87	0.95	9.3175	1838.0165
4	63	56	0.95	5.7337	632.8654	2	165	86	0.95	9.3396	1842.5371
4	63	55	0.95	5.7984	636.4888	2	165	85	0.95	9.3610	1847.2491
4	63	54	0.95	5.8619	640.2999	2	167	86	0.95	9.3790	1864.8709
4	63	53	0.95	5.9238	644.2895	2	167	85	0.95	9.3999	1869.6400
4	63	52	0.95	5.9844	648.5409	2	165	78	0.95	9.4442	1898.3002
4	63	50	0.95	6.1005	657.7897	2	167	78	0.95	9.4792	1921.3099
4	63	48	0.95	6.2091	668.2208	2	165	74	0.95	9.4992	1924.9458
4	63	47	0.95	6.2601	673.9414	2	165	71	0.95	9.5286	1948.1425
4	63	46	0.95	6.3089	680.0671	2	167	74	0.95	9.5321	1948.2784
4	63	45	0.95	6.3549	686.6102	2	165	70	0.95	9.5358	1956.5473
4	63	44	0.95	6.3982	693.6183	2	167	70	0.95	9.5666	1980.2630
4	63	43	0.95	6.4383	701.1650	2	177	78	0.95	9.6424	2036.3584
4	63	42	0.95	6.4751	709.2991	2	178	79	0.95	9.6447	2041.3628
4	63	41	0.95	6.5080	718.0637	2	183	81	0.95	9.6936	2086.0797
4	63	40	0.95	6.5369	727.6188	2	183	78	0.95	9.7317	2105.3875
4	63	39	0.95	6.5612	737.9782	2	183	71	0.95	9.7855	2160.6671
4	67	44	0.90	6.5640	745.8004	2	183	70	0.95	9.7884	2169.9888
4	63	38	0.95	6.5804	749.3191	2	193	78	0.95	9.8683	2220.4360
4	63	37	0.95	6.5940	761.7129	2	193	71	0.95	9.9076	2278.7364
4	67	42	0.90	6.6303	762.2380	2	193	70	0.95	9.9083	2288.5674
4	67	40	0.90	6.6820	781.3981	2	210	86	0.95	9.9942	2361.9468
4	67	39	0.90	6.7016	792.2349	2	210	78	0.95	10.0706	2416.0184
4	79	55	0.95	6.7081	798.1368	2	210	72	0.95	10.0894	2469.1626
4	79	54	0.95	6.7522	802.9158	2	230	94	0.95	10.1241	2541.7967
4	79	53	0.95	6.7949	807.9186	2	230	78	0.95	10.2703	2646.1154
4	79	52	0.95	6.8361	813.2507	2	232	78	0.95	10.2884	2669.1251
4	79	50	0.95	6.9134	824.8483	2	254	100	0.95	10.2973	2776.8890
4	74	42	0.95	6.9563	833.1455	2	254	96	0.95	10.3473	2796.3660
4	79	48	0.95	6.9829	837.9284	2	246	81	0.95	10.3992	2804.2383
4	79	47	0.95	7.0145	845.1018	2	254	89	0.95	10.4178	2836.6264
4	79	46	0.95	7.0435	852.7832	2	254	87	0.95	10.4332	2849.8842
4	79	45	0.95	7.0700	860.9880	2	254	86	0.95	10.4400	2856.8309
4	79	44	0.95	7.0935	869.7758	2	254	85	0.95	10.4461	2864.0117
4	79	43	0.95	7.1139	879.2392	2	254	81	0.95	10.4637	2895.4330
4	79	42	0.95	7.1307	889.4391	2	260	89	0.95	10.4689	2903.6333
4	79	41	0.95	7.1438	900.4291	2	264	85	0.95	10.5256	2976.7680
4	79	40	0.95	7.1527	912.4109	2	272	93	0.95	10.5343	3011.9094
4	79	39	0.95	7.1569	925.4063	2	272	89	0.95	10.5642	3037.6472
4	96	60	0.95	7.1706	944.8269	2	272	87	0.95	10.5758	3051.8445
4	96	59	0.95	7.2072	949.4172	2	272	86	0.95	10.5808	3059.2835
4	96	58	0.95	7.2427	954.1395	2	272	85	0.95	10.5850	3066.9731
4	96	57	0.95	7.2772	959.1121	2	284	96	0.95	10.6020	3126.6454
4	96	55	0.95	7.3424	969.8890	2	284	89	0.95	10.6515	3171.6610
4	96	54	0.95	7.3731	975.6963	2	284	87	0.95	10.6609	3186.4847
4	96	53	0.95	7.4023	981.7756	2	284	86	0.95	10.6647	3194.2518
4	96	52	0.95	7.4300	988.2540	2	284	85	0.95	10.6679	3202.2808
4	96	48	0.95	7.5225	1018.2421	2	295	96	0.95	10.6824	3247.7479
4	96	47	0.95	7.5404	1026.9687	2	295	89	0.95	10.7253	3294.5070
4	96	46	0.95	7.5558	1036.3027	2	295	87	0.95	10.7328	3309.9049
4	102	53	0.95	7.5683	1043.1366	2	295	86	0.95	10.7357	3317.9729
4	96	45	0.95	7.5685	1046.2727	2	295	85	0.95	10.7379	3326.3128
4	96	44	0.95	7.5783	1056.9511	2	295	84	0.95	10.7395	3334.9172
4	96	43	0.95	7.5849	1068.4505	2	316	85	0.95	10.8580	3563.1011
4	96	42	0.95	7.5879	1080.8448	2	354	112	0.95	10.9240	3803.8598
4	99	44	0.95	7.6466	1089.9809	2	354	100	0.95	11.0102	3870.1524
2	89	58	0.95	7.7789	1127.3006	2	352	96	0.95	11.0186	3875.2788
4	113	54	0.95	7.8071	1148.4903	2	354	96	0.95	11.0284	3897.2974
4	120	55	0.95	7.9319	1212.3771	2	354	89	0.95	11.0429	3953.4084
4	118	51	0.95	7.9587	1223.1750	2	364	100	0.95	11.0600	3979.4787
4	125	52	0.95	8.0702	1286.8042	2	373	100	0.95	11.1025	4077.8724
2	119	85	0.95	8.1154	1329.9973	2	373	96	0.95	11.1165	4106.4744
2	119	78	0.95	8.3546	1359.7707	2	381	100	0.95	11.1386	4165.3335
2	119	74	0.95	8.4775	1379.2029	2	381	96	0.95	11.1510	4194.5489
2	119	71	0.95	8.5578	1396.1073	2	381	89	0.95	11.1554	4254.9396
2	120	72	0.95	8.5631	1401.8890	2	388	96	0.95	11.1800	4271.6142
2	123	74	0.95	8.6019	1425.5627	2	388	89	0.95	11.1820	4333.1143
2	123	71	0.95	8.6761	1443.0352	2	397	89	0.95	11.2149	4433.6247
2	123	70	0.95	8.6983	1449.3770						

Table E.7: Pareto–Optimal Solutions for $\sigma = 0.3$ and $q = 0.9$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
8	15	14	0.65	3.5494	486.2434	2	190	80	0.65	8.2488	3526.0701
8	21	20	0.70	4.1795	517.2701	2	190	78	0.65	8.2532	3528.6310
4	27	25	0.70	5.9640	569.5834	2	190	77	0.65	8.2553	3529.8855
4	36	34	0.70	6.6902	638.5545	2	190	76	0.65	8.2572	3531.3037
4	43	40	0.65	6.8980	729.3834	2	190	75	0.65	8.2591	3532.6645
4	45	39	0.70	6.9569	761.5631	2	190	74	0.65	8.2608	3534.2299
2	40	39	0.65	7.3231	799.2733	2	190	72	0.65	8.2639	3537.3681
2	46	40	0.75	7.3692	925.1974	2	190	69	0.65	8.2676	3542.6321
2	49	47	0.70	7.5808	940.0515	2	190	62	0.65	8.2693	3558.6563
2	53	52	0.65	7.6398	1003.9920	2	192	60	0.70	8.2715	3601.3553
2	53	51	0.65	7.6451	1005.7809	2	203	80	0.70	8.2768	3764.6320
2	53	48	0.65	7.6526	1012.3609	2	203	78	0.70	8.2805	3767.3458
2	57	56	0.70	7.6727	1073.7925	2	203	77	0.70	8.2820	3768.9142
2	58	57	0.70	7.6799	1091.3118	2	203	76	0.70	8.2836	3770.4043
2	58	56	0.70	7.6871	1092.6309	2	203	75	0.70	8.2851	3771.9770
2	58	54	0.70	7.6996	1095.6516	2	203	74	0.70	8.2864	3773.6576
2	58	53	0.70	7.7045	1097.3687	2	203	72	0.70	8.2886	3777.2528
2	58	52	0.70	7.7083	1099.2644	2	203	69	0.70	8.2910	3783.1159
2	58	51	0.70	7.7111	1101.3748	2	211	75	0.65	8.2929	3923.1169
2	58	50	0.70	7.7124	1103.7255	2	211	74	0.65	8.2941	3924.8553
2	66	64	0.65	7.7283	1234.3443	2	211	72	0.65	8.2963	3928.3404
2	60	47	0.70	7.7283	1151.0835	2	211	69	0.65	8.2986	3934.1862
2	66	62	0.65	7.7452	1236.1641	2	216	82	0.80	8.3009	3998.6755
2	66	60	0.65	7.7608	1238.2040	2	227	84	0.70	8.3063	4204.1572
2	66	58	0.65	7.7750	1240.5655	2	227	80	0.70	8.3127	4209.7116
2	66	56	0.65	7.7872	1243.2832	2	227	78	0.70	8.3155	4212.7463
2	66	55	0.65	7.7924	1244.7943	2	227	77	0.70	8.3166	4214.5001
2	66	54	0.65	7.7969	1246.4572	2	227	76	0.70	8.3177	4216.1664
2	66	53	0.65	7.8007	1248.2644	2	227	75	0.70	8.3187	4217.9250
2	66	52	0.65	7.8035	1250.2541	2	227	74	0.70	8.3195	4219.8043
2	66	51	0.65	7.8054	1252.4819	2	227	72	0.70	8.3208	4223.8246
2	66	50	0.65	7.8059	1254.8955	2	227	69	0.70	8.3218	4230.3808
2	68	52	0.50	7.8177	1289.2939	2	231	70	0.70	8.3262	4302.6695
2	68	49	0.50	7.8252	1295.4080	2	241	80	0.70	8.3304	4469.3414
2	71	50	0.70	7.8457	1351.1122	2	241	78	0.70	8.3327	4472.5632
2	82	66	0.55	7.8678	1534.0174	2	241	77	0.70	8.3335	4474.4252
2	86	67	0.70	7.9066	1604.6209	2	241	75	0.70	8.3352	4478.0613
2	88	62	0.65	7.9460	1648.2188	2	241	74	0.70	8.3358	4480.0565
2	89	52	0.65	7.9760	1685.9488	2	241	72	0.70	8.3367	4484.3248
2	99	69	0.70	7.9845	1844.9667	2	241	69	0.70	8.3370	4491.2853
2	99	60	0.70	8.0208	1856.9478	2	243	75	0.70	8.3374	4515.2237
2	99	58	0.70	8.0252	1860.6909	2	243	74	0.70	8.3380	4517.2354
2	99	56	0.70	8.0274	1865.0080	2	243	72	0.70	8.3388	4521.5391
2	101	61	0.70	8.0284	1892.7587	2	243	69	0.70	8.3390	4528.5574
2	101	60	0.70	8.0311	1894.4619	2	250	80	0.70	8.3407	4636.2463
2	101	58	0.70	8.0351	1898.2807	2	250	78	0.70	8.3427	4639.5884
2	101	56	0.70	8.0369	1902.6850	2	250	77	0.70	8.3434	4641.5199
2	112	72	0.65	8.0372	2085.1841	2	250	76	0.70	8.3442	4643.3550
2	112	69	0.65	8.0506	2088.2872	2	250	75	0.70	8.3448	4645.2918
2	112	61	0.65	8.0775	2099.4495	2	250	74	0.70	8.3453	4647.3616
2	112	60	0.65	8.0796	2101.1947	2	250	72	0.70	8.3459	4651.7892
2	112	58	0.65	8.0826	2105.2021	2	252	71	0.70	8.3479	4691.3645
2	112	56	0.65	8.0837	2109.8139	2	259	80	0.70	8.3502	4803.1512
2	128	51	0.70	8.1177	2430.6213	2	259	77	0.70	8.3526	4808.6146
2	134	72	0.70	8.1317	2493.3574	2	259	75	0.70	8.3538	4812.5223
2	134	69	0.70	8.1407	2497.2277	2	259	74	0.70	8.3542	4814.6666
2	134	62	0.70	8.1546	2509.0903	2	261	86	1.00	8.3603	4820.1051
2	134	61	0.70	8.1555	2511.1862	2	277	80	0.65	8.3627	5140.6391
2	134	60	0.70	8.1560	2513.4459	2	277	76	0.65	8.3653	5148.2691
2	141	69	0.70	8.1627	2627.6815	2	277	75	0.65	8.3657	5150.2530
2	149	77	0.70	8.1637	2766.3441	2	277	74	0.65	8.3659	5152.5351
2	149	76	0.70	8.1669	2767.4378	2	277	72	0.65	8.3662	5157.1104
2	149	75	0.70	8.1699	2768.5922	2	308	92	0.70	8.3827	5691.5973
2	149	74	0.70	8.1728	2769.8275	2	308	85	0.70	8.3895	5702.5459
2	149	69	0.70	8.1852	2776.7698	2	308	84	0.70	8.3903	5704.3190
2	149	64	0.70	8.1931	2785.6431	2	308	80	0.70	8.3926	5711.8554
2	149	61	0.70	8.1947	2792.2891	2	308	78	0.70	8.3933	5715.9729
2	156	69	0.65	8.1997	2908.6874	2	308	77	0.70	8.3934	5718.3525
2	156	61	0.65	8.2087	2924.2348	2	308	75	0.70	8.3935	5722.9995
2	157	61	0.65	8.2108	2942.9799	2	318	80	0.65	8.3948	5901.5279
2	167	78	0.95	8.2187	3094.7539	2	318	78	0.65	8.3954	5905.8141
2	167	76	0.95	8.2224	3097.8993	2	318	75	0.65	8.3958	5912.5648
2	167	74	0.95	8.2256	3101.1971	2	320	72	0.70	8.4000	5954.2901
2	167	72	0.95	8.2281	3104.9101	2	322	76	0.70	8.4027	5980.6413
2	167	69	0.95	8.2303	3111.1743	2	338	84	0.70	8.4111	6259.9345
2	175	76	0.70	8.2321	3250.3485	2	340	84	0.70	8.4124	6296.9755
2	175	75	0.70	8.2342	3251.7043	2	340	80	0.70	8.4136	6305.2950
2	175	74	0.70	8.2362	3253.1531	2	348	77	0.70	8.4181	6460.9957
2	175	72	0.70	8.2398	3256.2524	2	371	84	0.60	8.4185	6882.9977
2	175	69	0.70	8.2443	3261.3068	2	371	80	0.60	8.4196	6891.3404
2	175	61	0.70	8.2466	3279.5342	2	387	93	0.70	8.4347	7149.8639
2	176	63	0.65	8.2473	3293.9333	2	397	75	0.65	8.4361	7381.4095

Table E.8: Pareto–Optimal Solutions for $\sigma = 0.6$ and $q = 0.1$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
8	24	23	0.95	1.0816	482.1846	4	282	100	0.80	13.3747	2119.0839
4	45	44	0.80	1.7129	510.5045	4	282	99	0.80	13.4082	2125.3488
4	50	49	0.80	1.7998	522.3037	4	307	115	0.95	13.7849	2133.8342
4	44	41	0.80	2.1929	530.2067	4	307	114	0.95	13.8244	2138.4273
4	51	46	0.80	2.7648	558.4541	4	307	108	0.95	14.0478	2168.2810
4	49	43	0.65	2.8875	582.7430	4	307	107	0.95	14.0826	2173.6675
4	84	77	0.95	3.3419	655.7478	4	307	106	0.95	14.1167	2179.1983
2	106	99	0.80	4.1350	676.0835	4	307	103	0.95	14.2145	2196.5686
4	84	70	0.95	4.6197	685.1580	4	307	100	0.95	14.3047	2215.2925
2	94	80	0.80	5.3874	691.3841	4	307	99	0.95	14.3331	2221.8408
2	82	66	0.55	5.4999	712.4999	4	307	96	0.95	14.4124	2242.5935
2	94	76	0.80	6.0519	720.6762	4	307	95	0.95	14.4367	2249.8936
2	106	86	0.80	6.4505	745.4130	4	307	94	0.95	14.4601	2257.3807
2	113	92	0.80	6.6102	763.0810	4	307	93	0.95	14.4822	2265.1139
2	106	80	0.80	7.3030	789.7908	4	307	88	0.95	14.5754	2307.1432
2	113	87	0.80	7.3064	799.1850	4	307	87	0.95	14.5902	2316.3117
2	113	86	0.80	7.4536	804.5776	4	307	86	0.95	14.6036	2325.7678
2	106	77	0.80	7.5440	831.6270	4	307	85	0.95	14.6155	2335.5275
2	106	76	0.80	7.6750	838.4724	4	320	96	0.95	14.6417	2337.5567
2	141	110	0.95	7.8908	876.5295	4	320	95	0.95	14.6631	2345.1660
2	113	78	0.80	8.3085	879.6103	4	320	94	0.95	14.6836	2352.9701
2	113	77	0.80	8.4264	886.5457	4	320	93	0.95	14.7028	2361.0308
2	113	76	0.80	8.5399	893.8432	4	320	89	0.95	14.7682	2395.5419
2	133	94	0.80	8.6201	945.6255	4	320	88	0.95	14.7813	2404.8398
2	141	100	0.95	8.8192	968.5614	4	320	87	0.95	14.7931	2414.3966
2	141	99	0.95	8.9484	971.9871	4	320	86	0.95	14.8036	2424.2531
2	131	88	0.65	9.0119	974.9198	4	320	85	0.95	14.8125	2434.4260
2	141	96	0.95	9.3237	983.2123	4	344	103	0.95	14.8758	2461.3017
2	141	95	0.95	9.4444	987.2952	4	344	100	0.95	14.9429	2482.2822
2	141	94	0.95	9.5628	991.5692	4	344	99	0.95	14.9635	2489.6197
2	167	115	0.95	9.5820	1102.7337	4	344	96	0.95	15.0195	2512.8735
2	167	114	0.95	9.7016	1104.9455	4	344	95	0.95	15.0361	2521.0534
4	140	73	0.95	9.9381	1121.9289	4	344	94	0.95	15.0517	2529.4429
2	135	73	0.80	10.1506	1185.4403	4	344	93	0.95	15.0661	2538.1081
2	141	77	0.95	10.2458	1211.5940	4	360	103	0.95	15.1197	2575.7808
2	141	75	0.95	10.3925	1226.4849	4	344	87	0.95	15.1275	2595.4763
2	141	74	0.95	10.4585	1234.7695	4	360	100	0.95	15.1782	2597.7372
2	141	73	0.95	10.5192	1243.6733	4	360	99	0.95	15.1960	2605.4159
2	167	93	0.95	10.5568	1355.4545	4	360	96	0.95	15.2434	2629.7513
4	198	108	0.95	10.5770	1378.0303	4	360	95	0.95	15.2572	2638.3117
4	167	76	0.80	10.8044	1379.8561	4	360	94	0.95	15.2699	2647.0914
2	167	86	0.95	11.1804	1380.1322	4	371	103	0.95	15.2752	2654.4852
4	198	99	0.95	11.2162	1413.8372	4	360	93	0.95	15.2814	2656.1596
4	198	96	0.95	11.3049	1446.3632	4	371	100	0.95	15.3282	2677.1125
4	198	95	0.95	11.3690	1451.0714	4	371	99	0.95	15.3442	2685.0259
4	198	94	0.95	11.4320	1455.9003	4	371	96	0.95	15.3861	2710.1048
4	198	93	0.95	11.4939	1460.8878	4	371	95	0.95	15.3981	2718.9268
4	198	87	0.95	11.8403	1493.9079	4	371	94	0.95	15.4090	2727.9747
4	198	86	0.95	11.8934	1500.0066	4	371	93	0.95	15.4187	2737.3201
4	198	85	0.95	11.9451	1506.3011	4	387	103	0.95	15.4855	2768.9644
4	198	80	0.95	12.1798	1540.9922	4	394	106	0.95	15.5221	2796.7561
4	198	78	0.95	12.2614	1556.5868	4	394	103	0.95	15.5722	2819.0490
4	198	77	0.95	12.2993	1564.7940	4	394	100	0.95	15.6148	2843.0790
4	198	75	0.95	12.3690	1582.1176	4	394	99	0.95	15.6273	2851.4830
4	198	74	0.95	12.4004	1591.2887	4	394	96	0.95	15.6588	2878.1167
4	198	73	0.95	12.4296	1600.7821	4	394	95	0.95	15.6672	2887.4856
4	198	58	0.95	12.5055	1803.7844	4	394	94	0.95	15.6747	2897.0944
4	234	78	0.95	13.2976	1839.6026	4	394	93	0.95	15.6809	2907.0192
4	252	75	0.80	13.3429	2093.5234	4	394	87	0.95	15.6932	2972.7258

Table E.9: Pareto–Optimal Solutions for $\sigma = 0.6$ and $q = 0.5$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
8	21	20	0.80	3.0617	606.5836	2	134	82	0.65	8.6797	1964.2666
8	21	19	0.80	3.0866	633.1166	2	134	80	0.65	8.7377	1972.9318
8	24	23	0.65	3.2343	645.0447	2	131	66	0.65	8.9691	2010.5983
8	24	22	0.65	3.2872	662.9600	2	134	66	0.65	8.9983	2066.1964
4	33	32	0.80	4.5097	665.9423	2	134	62	0.65	9.0370	2100.8596
4	44	42	0.65	5.2189	764.0823	2	134	59	0.65	9.0452	2131.9585
4	44	41	0.65	5.2858	771.8553	2	165	96	0.80	9.1452	2338.7237
4	44	40	0.65	5.3477	780.2228	2	165	95	0.80	9.1706	2342.5424
4	44	39	0.65	5.4045	789.3018	2	165	92	0.80	9.2437	2354.7310
4	44	38	0.65	5.4557	799.1628	2	165	89	0.80	9.3119	2368.0712
4	44	37	0.65	5.5004	809.8669	2	165	87	0.80	9.3543	2377.7307
4	44	36	0.65	5.5385	821.5762	2	165	86	0.80	9.3744	2382.8307
4	44	35	0.65	5.5690	834.3767	2	164	82	0.65	9.3797	2404.0278
4	49	43	0.65	5.6496	842.8861	2	165	82	0.80	9.3999	2418.3321
2	53	52	0.65	5.6880	878.6946	2	165	80	0.80	9.4320	2430.5392
4	57	55	0.80	5.6973	883.5793	2	164	76	0.65	9.4392	2451.0462
4	57	54	0.80	5.7677	888.2728	2	165	76	0.80	9.4848	2457.9670
4	57	53	0.80	5.8359	893.2363	2	164	72	0.65	9.4921	2478.3106
4	57	51	0.80	5.9653	904.0813	2	165	72	0.80	9.5194	2490.2428
4	57	50	0.80	6.0262	910.0245	2	164	67	0.65	9.5322	2519.4115
4	57	49	0.80	6.0842	916.3409	2	164	66	0.65	9.5361	2528.7776
4	57	48	0.80	6.1390	923.0690	2	164	65	0.65	9.5384	2538.6201
4	57	47	0.80	6.1906	930.2820	2	173	72	0.65	9.6310	2614.3154
4	57	46	0.80	6.2386	937.9995	2	173	67	0.65	9.6595	2657.6719
4	57	45	0.80	6.2827	946.2612	2	173	66	0.65	9.6611	2667.5520
4	57	44	0.80	6.3226	955.1494	2	176	67	0.80	9.6824	2708.8125
4	57	42	0.80	6.3882	975.0618	2	188	87	0.80	9.7132	2724.6827
4	57	41	0.80	6.4132	986.2627	2	188	82	0.80	9.7801	2755.4329
4	57	40	0.80	6.4323	998.4370	2	188	80	0.80	9.8013	2769.3417
4	57	39	0.80	6.4451	1011.6990	2	193	86	0.80	9.8028	2803.0298
4	57	38	0.80	6.4507	1026.1594	2	188	72	0.80	9.8445	2837.3676
4	60	40	0.80	6.6062	1050.9864	2	205	70	0.80	10.0293	3116.7413
2	68	58	0.50	6.6472	1082.7653	2	216	82	0.80	10.1337	3165.8165
2	68	54	0.50	6.8568	1105.2309	2	222	82	0.65	10.1457	3272.0027
2	68	53	0.50	6.9030	1111.7013	2	222	80	0.65	10.1613	3286.2108
2	68	52	0.50	6.9465	1118.5732	2	222	78	0.65	10.1739	3301.4415
2	68	51	0.50	6.9868	1125.9050	2	222	72	0.65	10.1896	3354.7863
2	68	50	0.50	7.0232	1133.8587	2	230	74	0.80	10.2810	3447.8134
2	68	49	0.50	7.0565	1142.2236	2	243	96	0.80	10.2936	3464.7542
2	68	46	0.50	7.1316	1171.3265	2	241	89	0.70	10.3003	3494.0982
2	68	45	0.50	7.1447	1183.0110	2	241	85	0.70	10.3335	3520.5559
2	68	44	0.50	7.1545	1195.2484	2	243	87	0.80	10.3750	3521.7973
2	68	43	0.50	7.1581	1208.5690	2	243	80	0.80	10.4008	3579.5214
4	80	56	0.80	7.1863	1233.8766	2	259	94	0.80	10.4684	3704.9638
4	80	53	0.80	7.2937	1253.6667	2	271	92	0.65	10.5039	3922.8083
4	80	52	0.80	7.3250	1261.0516	2	295	106	0.95	10.7171	4122.8338
2	76	51	0.65	7.3614	1270.6889	2	313	112	0.80	10.7346	4370.9054
4	80	50	0.80	7.3798	1277.2288	2	313	111	0.80	10.7449	4375.6552
4	80	49	0.80	7.4028	1286.0939	2	305	87	0.65	10.7644	4451.9255
4	80	48	0.80	7.4227	1295.5366	2	313	96	0.80	10.8603	4462.8316
4	80	45	0.80	7.4607	1328.0868	2	313	92	0.80	10.8751	4492.9683
4	80	44	0.80	7.4651	1340.5615	2	338	115	0.80	10.8802	4705.5053
2	88	66	0.65	7.5014	1349.0509	2	338	112	0.80	10.9076	4720.0192
2	89	64	0.65	7.6278	1375.8701	2	338	111	0.80	10.9162	4725.1484
4	91	53	0.80	7.7305	1426.0459	2	338	104	0.80	10.9682	4764.4129
4	91	52	0.80	7.7514	1434.4462	2	338	96	0.80	11.0058	4819.2878
4	88	46	0.85	7.7525	1442.2230	2	338	95	0.80	11.0085	4827.0804
4	91	50	0.80	7.7852	1452.8477	2	338	92	0.80	11.0137	4851.8316
4	94	53	0.80	7.8318	1473.0584	2	338	89	0.80	11.0139	4878.9640
4	94	46	0.80	7.8997	1546.8901	2	354	105	0.80	11.0528	4983.6045
4	105	62	0.80	7.9644	1578.2705	2	371	115	0.80	11.0786	5164.9186
4	105	58	0.80	8.0666	1604.2988	2	371	112	0.80	11.1003	5180.8495
4	105	54	0.80	8.1407	1636.2943	2	371	111	0.80	11.1069	5186.4795
4	105	53	0.80	8.1539	1645.4376	2	373	112	0.80	11.1109	5208.7786
4	105	51	0.80	8.1730	1665.4331	2	373	111	0.80	11.1174	5214.4389
2	110	69	0.80	8.3080	1670.8476	2	371	104	0.80	11.1456	5229.5775
2	119	80	0.80	8.3164	1741.2082	2	373	106	0.80	11.1461	5244.5778
2	119	78	0.80	8.3695	1753.0259	2	373	104	0.80	11.1553	5257.7693
2	119	77	0.80	8.3995	1758.1192	2	371	96	0.80	11.1678	5289.8100
2	119	74	0.80	8.4826	1774.6982	2	371	95	0.80	11.1686	5298.3634
2	119	72	0.80	8.5319	1786.9638	2	373	96	0.80	11.1767	5318.3265
2	119	68	0.80	8.6130	1815.0873	2	373	95	0.80	11.1774	5326.9260
2	119	67	0.80	8.6292	1822.9924	2	388	112	0.80	11.1868	5418.2469
2	119	66	0.80	8.6434	1831.2737	2	388	111	0.80	11.1925	5424.1348
2	119	65	0.80	8.6557	1840.0329	2	388	104	0.80	11.2252	5469.2077
2	119	62	0.80	8.6792	1869.1775	2	388	96	0.80	11.2405	5532.2002

Table E.10: Pareto–Optimal Solutions for $\sigma = 0.6$ and $q = 0.9$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
8	20	19	0.80	4.4211	712.0432	2	241	78	0.70	8.3404	5270.6887
2	31	26	0.25	6.6935	807.2328	2	241	77	0.70	8.3410	5272.7466
4	39	37	0.80	6.9757	908.1188	2	241	76	0.70	8.3414	5274.8697
4	44	43	0.80	7.2158	983.0341	2	241	75	0.70	8.3418	5276.9815
4	45	44	0.80	7.2471	1000.4334	2	260	95	0.80	8.3475	5655.1146
2	46	45	0.80	7.3657	1065.5176	2	261	99	1.00	8.3477	5668.0079
2	49	48	0.90	7.4056	1128.3589	2	261	95	1.00	8.3528	5673.9279
4	57	55	0.80	7.4836	1224.3075	2	261	93	1.00	8.3550	5677.0893
4	57	54	0.80	7.4841	1226.8717	2	261	89	1.00	8.3584	5684.2392
2	57	54	0.80	7.6501	1272.9937	2	261	88	1.00	8.3591	5686.0937
2	66	56	0.80	7.7607	1468.5539	2	261	87	1.00	8.3596	5688.2069
2	68	58	0.50	7.7948	1505.2696	2	261	86	1.00	8.3600	5690.3274
2	68	56	0.50	7.8069	1507.9087	2	261	82	1.00	8.3606	5699.4040
2	68	55	0.50	7.8122	1509.3825	2	263	77	0.90	8.3642	5753.9560
2	68	54	0.50	7.8169	1510.9754	2	269	89	0.80	8.3653	5859.9126
2	68	50	0.50	7.8275	1519.0360	2	269	87	0.80	8.3673	5863.1766
2	76	51	0.80	7.8279	1710.1501	2	269	86	0.80	8.3682	5864.8578
2	82	66	0.55	7.8783	1804.7541	2	269	77	0.80	8.3714	5884.0144
2	100	65	0.80	8.0055	2202.5175	2	271	78	0.90	8.3721	5926.2778
2	109	73	0.90	8.0235	2389.7177	2	282	95	0.90	8.3752	6131.3354
2	110	64	0.80	8.0577	2424.7774	2	282	93	0.90	8.3773	6134.4438
2	125	77	0.90	8.0857	2734.7684	2	282	89	0.90	8.3806	6141.4300
2	129	64	0.80	8.1311	2843.6039	2	282	88	0.90	8.3812	6143.3341
2	133	64	0.80	8.1438	2931.7776	2	282	87	0.90	8.3818	6145.2561
2	134	64	0.80	8.1469	2953.8211	2	282	86	0.90	8.3822	6147.2888
2	145	75	0.80	8.1654	3174.5546	2	286	80	0.90	8.3860	6248.8011
2	152	60	0.90	8.1714	3370.7711	2	296	87	0.80	8.3928	6451.6739
2	164	86	0.80	8.1912	3575.5989	2	296	86	0.80	8.3934	6453.5238
2	164	77	0.80	8.2134	3587.2802	2	296	77	0.80	8.3939	6474.6032
2	164	76	0.80	8.2152	3588.9221	2	310	88	0.80	8.4038	6754.8298
2	164	64	0.80	8.2216	3615.1243	2	311	89	0.80	8.4040	6774.8432
2	174	78	0.90	8.2338	3805.0640	2	311	88	0.80	8.4046	6776.6196
2	174	77	0.90	8.2351	3806.7998	2	311	87	0.80	8.4051	6778.6168
2	174	76	0.90	8.2363	3808.6910	2	311	86	0.80	8.4055	6780.5605
2	174	75	0.90	8.2373	3810.5736	2	312	89	0.90	8.4069	6794.7736
2	187	89	0.80	8.2398	4073.6195	2	312	87	0.90	8.4075	6799.0067
2	187	87	0.80	8.2447	4075.8886	2	312	86	0.90	8.4076	6801.2557
2	187	86	0.80	8.2470	4077.0573	2	320	80	0.90	8.4116	6991.6656
2	187	78	0.80	8.2618	4088.6954	2	333	93	0.90	8.4211	7243.8645
2	187	77	0.80	8.2632	4090.3743	2	333	88	0.90	8.4226	7254.3626
2	187	76	0.80	8.2643	4092.2465	2	336	93	0.90	8.4232	7309.1245
2	187	75	0.80	8.2654	4094.0830	2	336	89	0.90	8.4245	7317.4485
2	195	78	0.90	8.2754	4264.2958	2	336	88	0.90	8.4246	7319.7172
2	195	77	0.90	8.2762	4266.2411	2	356	89	0.90	8.4374	7753.0109
2	195	76	0.90	8.2768	4268.3606	2	377	105	0.80	8.4385	8183.6986
2	195	75	0.90	8.2773	4270.4705	2	377	99	0.80	8.4434	8192.9204
2	200	75	0.90	8.2855	4379.9697	2	377	95	0.80	8.4457	8199.9162
2	202	78	0.90	8.2873	4417.3731	2	377	93	0.80	8.4465	8203.8182
2	202	77	0.90	8.2880	4419.3883	2	377	89	0.80	8.4473	8212.5913
2	202	76	0.90	8.2884	4421.5838	2	377	88	0.80	8.4475	8214.7447
2	202	75	0.90	8.2887	4423.7694	2	383	101	0.90	8.4497	8315.6270
2	204	75	0.90	8.2919	4467.5691	2	383	99	0.90	8.4508	8319.2897
2	205	70	0.80	8.2962	4499.9777	2	383	95	0.90	8.4522	8327.3102
2	210	78	0.90	8.3000	4592.3186	2	383	93	0.90	8.4527	8331.5318
2	210	77	0.90	8.3005	4594.4135	2	389	101	0.90	8.4533	8445.8979
2	210	76	0.90	8.3007	4596.6960	2	389	99	0.90	8.4542	8449.6180
2	210	75	0.90	8.3009	4598.9682	2	389	96	0.90	8.4554	8455.4837
2	216	82	0.80	8.3060	4715.6547	2	389	95	0.90	8.4555	8457.7641
2	217	73	0.80	8.3135	4755.4892	2	389	93	0.90	8.4559	8462.0519
2	225	78	0.90	8.3214	4920.3413	2	394	89	0.80	8.4561	8582.9203
2	241	89	0.70	8.3272	5253.1880	2	398	95	0.80	8.4573	8656.6754
2	241	88	0.70	8.3288	5254.4709	2	398	93	0.80	8.4578	8660.7948
2	241	87	0.70	8.3303	5255.8997	2	398	89	0.80	8.4581	8670.0566
2	241	86	0.70	8.3318	5257.2659						

Table E.11: Pareto–Optimal Solutions for $\sigma = 0.8$ and $q = 0.1$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
4	40	39	0.90	1.6003	572.3656	4	259	87	0.95	14.6678	2204.8949
4	49	43	0.65	2.9521	648.8098	4	259	86	0.95	14.6854	2214.1994
4	66	61	0.95	3.0709	656.9047	4	276	102	0.95	14.6941	2231.1267
4	67	61	0.90	3.2746	672.0162	4	276	100	0.95	14.7605	2244.2382
2	75	69	0.35	3.3980	721.7626	4	276	99	0.95	14.7919	2251.0551
2	70	62	0.95	3.6486	743.3753	4	276	98	0.95	14.8220	2258.0739
2	75	64	0.35	4.3944	754.4130	4	276	96	0.95	14.8785	2272.6482
2	67	53	0.35	4.8462	763.0047	4	276	94	0.95	14.9294	2288.0678
2	82	66	0.55	5.5338	775.7964	4	276	92	0.95	14.9745	2304.3787
2	75	57	0.35	5.5746	812.7838	4	286	100	0.95	15.0102	2325.5511
2	111	93	0.95	6.2698	815.3383	4	276	89	0.95	15.0305	2330.7009
2	80	57	0.35	6.3019	870.8308	4	286	99	0.95	15.0385	2332.6151
2	80	55	0.35	6.5507	891.1966	4	286	98	0.95	15.0654	2339.8881
2	141	120	0.95	6.7672	904.8020	4	286	96	0.95	15.1155	2354.9905
2	141	117	0.95	7.1466	928.5723	4	286	94	0.95	15.1601	2370.9688
2	141	114	0.95	7.6108	938.5886	4	286	92	0.95	15.1989	2387.8707
2	141	113	0.95	7.7620	942.1270	4	286	89	0.95	15.2454	2415.1466
2	131	100	0.65	7.8672	976.8011	4	286	88	0.95	15.2575	2424.8071
2	131	99	0.65	8.0095	980.7535	4	286	87	0.95	15.2679	2434.7488
2	141	107	0.95	8.4203	991.2166	4	286	86	0.95	15.2763	2445.0233
2	141	106	0.95	8.5589	995.0520	4	305	105	0.95	15.3082	2445.1651
2	131	92	0.65	8.6184	1053.0347	4	309	107	0.95	15.3393	2464.2739
2	131	89	0.65	8.9953	1066.2043	4	305	102	0.95	15.3904	2465.5567
2	141	99	0.95	8.9991	1082.9416	4	309	105	0.95	15.3971	2477.2328
2	141	98	0.95	9.1244	1086.8355	4	305	100	0.95	15.4396	2480.0458
2	141	96	0.95	9.3679	1095.1867	4	305	98	0.95	15.4839	2495.3352
2	141	94	0.95	9.6011	1104.3618	4	305	96	0.95	15.5231	2511.4409
4	117	58	0.95	9.7605	1210.7953	4	309	100	0.95	15.5233	2512.5710
4	140	78	0.95	10.1996	1228.6812	4	309	99	0.95	15.5450	2520.2030
4	135	71	0.95	10.3554	1238.1972	4	309	98	0.95	15.5654	2528.0610
4	140	73	0.95	10.5635	1266.7329	4	309	96	0.95	15.6026	2544.3779
4	154	75	0.95	11.1962	1391.1716	4	309	94	0.95	15.6341	2561.6411
4	153	63	0.75	11.2093	1565.2193	4	309	92	0.95	15.6599	2579.9023
4	208	118	0.95	11.3660	1583.3111	4	309	89	0.95	15.6868	2609.3716
4	208	114	0.95	11.6720	1598.0383	4	309	88	0.95	15.6924	2619.8091
4	208	107	0.95	12.1733	1627.1673	4	309	87	0.95	15.6963	2630.5503
4	194	90	0.95	12.3775	1631.8985	4	309	86	0.95	15.6982	2641.6510
4	208	100	0.95	12.4256	1691.3099	4	338	114	0.95	15.7512	2650.9719
4	208	99	0.95	12.4865	1696.4473	4	344	118	0.95	15.7527	2675.1656
4	208	98	0.95	12.5462	1701.7368	4	342	114	0.95	15.8315	2682.3444
4	208	96	0.95	12.6619	1712.7204	4	338	107	0.95	15.9351	2695.5488
4	208	94	0.95	12.7721	1724.3410	4	338	105	0.95	15.9791	2709.7239
4	208	92	0.95	12.8765	1736.6332	4	342	107	0.95	16.0094	2727.4488
4	208	88	0.95	13.0662	1763.4961	4	338	102	0.95	16.0374	2732.3219
4	208	87	0.95	13.1094	1770.7264	4	342	105	0.95	16.0517	2741.7916
4	219	96	0.95	13.1138	1803.2969	4	338	100	0.95	16.0706	2748.3786
4	219	94	0.95	13.2119	1815.5321	4	342	104	0.95	16.0714	2749.2129
4	219	92	0.95	13.3042	1828.4744	4	338	99	0.95	16.0854	2756.7269
4	219	88	0.95	13.4698	1856.7579	4	344	105	0.95	16.0873	2757.8255
4	219	87	0.95	13.5070	1864.3706	4	342	102	0.95	16.1073	2764.6571
4	253	107	0.95	13.8021	2017.6741	4	344	102	0.95	16.1417	2780.8246
4	258	110	0.95	13.8243	2042.2914	4	342	99	0.95	16.1528	2789.3509
4	259	107	0.95	13.9986	2065.5241	4	344	100	0.95	16.1723	2797.1664
4	259	105	0.95	14.0876	2076.3861	4	344	99	0.95	16.1858	2805.6629
4	259	104	0.95	14.1307	2082.0063	4	344	98	0.95	16.1980	2814.4109
4	259	102	0.95	14.2134	2093.7023	4	344	96	0.95	16.2187	2832.5760
4	259	100	0.95	14.2917	2106.0061	4	344	94	0.95	16.2338	2851.7947
4	259	99	0.95	14.3291	2112.4032	4	344	92	0.95	16.2432	2872.1242
4	259	98	0.95	14.3651	2118.9896	4	344	89	0.95	16.2453	2904.9315
4	259	96	0.95	14.4334	2132.6662	4	397	123	0.90	16.4166	3090.9416
4	259	94	0.95	14.4963	2147.1361	4	397	114	0.90	16.6104	3147.7727
4	259	92	0.95	14.5533	2162.4423	4	390	99	0.95	16.8532	3180.8388
4	259	88	0.95	14.6483	2195.8917						

Table E.12: Pareto–Optimal Solutions for $\sigma = 0.8$ and $q = 0.5$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
8	15	14	0.65	2.3127	661.8219	2	131	72	0.65	8.8693	2076.7187
8	20	19	0.65	2.9519	671.3739	2	131	70	0.65	8.9083	2089.7811
8	24	23	0.65	3.3045	708.5869	2	131	67	0.65	8.9567	2111.6505
8	29	27	0.95	3.8514	746.8732	2	131	66	0.65	8.9698	2119.6131
4	36	34	0.65	4.8151	753.1532	2	131	61	0.65	8.9725	2174.6842
4	45	43	0.65	5.3046	841.0220	2	141	75	0.65	9.0521	2216.2482
4	45	42	0.65	5.3697	848.8844	2	145	76	0.65	9.1234	2273.1328
4	45	41	0.65	5.4300	857.3728	2	145	74	0.65	9.1608	2285.3644
4	45	39	0.65	5.5345	876.4810	2	145	73	0.65	9.1778	2291.9014
4	45	38	0.65	5.5777	887.2810	2	145	68	0.65	9.2041	2339.4613
4	45	36	0.65	5.6432	911.8767	2	145	66	0.65	9.2220	2356.4765
4	49	43	0.65	5.7008	915.7795	2	145	62	0.65	9.2368	2395.8819
2	55	53	0.65	5.8079	953.9355	4	154	62	0.95	9.2598	2502.6493
4	54	47	0.65	5.9137	977.5025	4	165	77	0.95	9.2607	2566.9740
4	54	46	0.65	5.9694	984.6409	4	165	76	0.95	9.2780	2572.6750
4	54	45	0.65	6.0214	992.2821	4	165	74	0.95	9.3105	2584.6136
4	54	41	0.65	6.1873	1028.8474	4	165	73	0.95	9.3255	2590.8911
4	54	40	0.65	6.2165	1039.8590	4	165	72	0.95	9.3395	2597.4463
2	63	53	0.65	6.4936	1092.6915	2	175	97	0.95	9.3813	2609.3509
4	71	62	0.95	6.5317	1153.7952	4	165	66	0.95	9.4009	2643.2502
4	71	61	0.95	6.5885	1158.4285	2	175	89	0.95	9.5291	2651.4667
4	71	59	0.95	6.6967	1168.4388	2	175	88	0.95	9.5444	2657.5715
4	71	57	0.95	6.7970	1179.5347	2	175	83	0.95	9.5584	2705.8605
4	71	56	0.95	6.8439	1185.6087	2	167	63	0.65	9.5736	2747.2029
4	71	55	0.95	6.8884	1192.0081	2	184	95	0.95	9.5939	2753.5331
4	71	54	0.95	6.9303	1198.7890	2	175	76	0.95	9.6035	2765.6219
4	71	53	0.95	6.9694	1206.0031	2	175	75	0.95	9.6047	2775.6963
4	71	52	0.95	7.0055	1213.6809	2	184	91	0.95	9.6080	2791.0212
4	71	51	0.95	7.0385	1221.8729	2	184	89	0.95	9.6369	2803.1265
4	71	50	0.95	7.0681	1230.6556	2	184	85	0.95	9.6851	2830.0197
4	71	49	0.95	7.0938	1239.9972	2	184	83	0.95	9.7037	2845.0190
4	71	48	0.95	7.1156	1250.0440	2	184	78	0.95	9.7315	2888.0089
2	68	44	0.50	7.1320	1255.9469	2	184	76	0.95	9.7338	2907.8539
4	71	47	0.95	7.1330	1260.8384	2	189	83	0.65	9.7510	2930.8273
2	75	57	0.65	7.1339	1266.1304	2	189	76	0.65	9.8303	2977.6490
4	71	46	0.95	7.1457	1272.4479	2	189	74	0.65	9.8448	2993.4122
2	75	56	0.65	7.1665	1274.1739	2	189	73	0.65	9.8505	3001.7868
2	75	55	0.65	7.1968	1282.5376	2	189	72	0.65	9.8551	3010.5032
2	75	54	0.65	7.2237	1291.4465	2	202	89	0.95	9.9108	3077.3453
2	75	52	0.65	7.2661	1311.2252	2	202	85	0.95	9.9443	3106.8694
2	75	50	0.65	7.2908	1334.0742	2	210	91	0.95	9.9988	3185.4046
4	76	43	0.95	7.3560	1405.7774	2	210	89	0.95	10.0175	3199.2204
4	85	56	0.65	7.3634	1463.2022	2	210	85	0.95	10.0452	3229.9137
4	85	55	0.65	7.3963	1469.8165	2	210	83	0.95	10.0536	3247.0326
4	85	54	0.65	7.4273	1476.8041	2	216	82	0.80	10.1391	3343.9824
4	85	52	0.65	7.4827	1491.9769	2	222	85	0.65	10.1407	3429.1679
4	85	51	0.65	7.5069	1500.2244	2	219	78	0.65	10.1613	3433.3420
4	85	49	0.65	7.5472	1518.2579	2	222	76	0.65	10.1980	3497.5560
4	85	47	0.65	7.5755	1538.6625	2	222	75	0.65	10.2001	3506.6277
2	99	76	0.65	7.6284	1549.7942	2	222	74	0.65	10.2011	3516.0714
2	99	75	0.65	7.6710	1553.9355	2	248	94	0.95	10.4056	3739.2999
2	99	74	0.65	7.7106	1558.6908	2	257	99	0.95	10.4557	3840.6443
2	99	73	0.65	7.7512	1563.1527	2	260	102	0.95	10.4578	3866.9956
2	99	72	0.65	7.7907	1567.8185	2	257	89	0.95	10.5101	3915.2364
4	88	46	0.85	7.7930	1580.1359	2	263	99	0.95	10.5126	3930.3091
2	99	63	0.65	8.0775	1621.8124	2	261	94	1.00	10.5263	3935.5177
2	99	62	0.65	8.1017	1629.1663	2	263	88	0.95	10.5612	4015.7829
2	99	61	0.65	8.1238	1636.9025	2	292	101	0.65	10.6181	4402.8474
2	99	57	0.65	8.1877	1672.5753	2	292	99	0.65	10.6365	4413.7618
2	99	56	0.65	8.1967	1682.9037	2	292	95	0.65	10.6688	4437.4355
2	99	55	0.65	8.2024	1693.9069	2	292	89	0.65	10.7038	4478.5551
2	99	54	0.65	8.2045	1705.6286	2	311	120	0.95	10.7262	4524.6355
2	120	83	0.95	8.2485	1843.3711	2	312	120	0.95	10.7344	4539.1842
2	120	77	0.95	8.4002	1880.4273	2	309	99	0.95	10.8754	4617.7396
2	120	76	0.95	8.4227	1887.1168	2	311	101	0.95	10.8809	4632.7155
2	120	72	0.95	8.4975	1917.0384	2	311	99	0.95	10.8887	4647.6279
2	121	72	0.95	8.5276	1933.0137	2	312	99	0.95	10.8953	4662.5721
2	123	72	0.95	8.5863	1964.9643	2	311	95	0.95	10.8978	4680.3351
4	121	62	0.95	8.6234	1966.3316	2	312	95	0.95	10.9041	4695.3844
4	121	61	0.95	8.6387	1974.2276	2	327	98	0.95	10.9915	4894.9571
2	127	75	0.95	8.6553	2004.5871	2	355	124	0.95	11.0092	5145.3129
2	127	73	0.95	8.6856	2020.4042	2	355	120	0.95	11.0440	5164.7769
2	127	72	0.95	8.6982	2028.8656	2	355	99	0.95	11.1442	5305.1701
2	127	70	0.95	8.7179	2047.1081	2	389	124	0.90	11.1922	5648.8344
2	131	76	0.65	8.7771	2053.6580	2	389	120	0.90	11.2212	5669.4276
2	131	75	0.65	8.8018	2059.0675	2	389	106	0.90	11.2878	5758.7984
2	131	74	0.65	8.8254	2064.7085	2	389	99	0.90	11.2943	5817.3648
2	131	73	0.65	8.8479	2070.6144						

Table E.13: Pareto–Optimal Solutions for $\sigma = 0.8$ and $q = 0.9$

D	F	M	p	TH	$Delay$	D	F	M	p	TH	$Delay$
8	4	4	0.10	1.2232	580.6966	2	153	78	0.90	8.1797	3473.7790
4	9	9	0.10	2.7377	582.7466	2	153	77	0.90	8.1815	3475.4053
2	15	12	0.10	3.9136	697.7605	2	153	75	0.90	8.1847	3478.8946
2	16	12	0.10	3.9348	744.2778	2	153	74	0.90	8.1860	3480.7919
2	31	28	0.25	6.8607	806.5829	2	153	72	0.90	8.1879	3484.8211
4	45	40	0.90	7.1671	1099.1132	2	153	68	0.90	8.1884	3494.4148
4	49	48	0.65	7.3153	1159.4253	2	167	82	0.95	8.2082	3785.8383
2	55	53	0.90	7.5638	1286.6573	2	174	88	0.90	8.2137	3935.5270
2	56	54	0.65	7.6617	1292.8592	2	171	55	0.40	8.2167	3942.9031
2	57	54	0.35	7.6664	1317.6165	2	174	81	0.90	8.2283	3945.4780
2	57	53	0.35	7.6741	1318.8991	2	174	78	0.90	8.2327	3950.5744
2	57	52	0.35	7.6812	1320.2712	2	174	77	0.90	8.2339	3952.4238
2	57	51	0.35	7.6874	1321.7345	2	174	75	0.90	8.2356	3956.3899
2	57	50	0.35	7.6930	1323.3494	2	174	74	0.90	8.2362	3958.5477
2	57	49	0.35	7.6978	1325.1288	2	174	72	0.90	8.2367	3963.1299
2	57	48	0.35	7.7015	1327.0671	2	185	72	0.90	8.2579	4213.6726
2	57	47	0.35	7.7042	1329.2275	2	206	65	0.35	8.2595	4727.8212
2	57	46	0.35	7.7055	1331.5755	2	210	74	0.90	8.2989	4777.5576
2	66	59	0.35	7.7410	1519.7440	2	216	82	0.80	8.3063	4895.4456
2	66	58	0.35	7.7488	1520.7795	2	222	82	0.65	8.3105	5034.3272
2	66	53	0.35	7.7814	1527.1463	2	222	81	0.65	8.3119	5035.8132
2	66	49	0.35	7.7965	1534.3596	2	222	78	0.65	8.3156	5040.7448
2	66	48	0.35	7.7981	1536.6040	2	222	77	0.65	8.3165	5042.5951
2	76	51	0.35	7.8763	1762.3127	2	222	75	0.65	8.3182	5046.4049
2	77	57	0.65	7.8813	1770.8244	2	222	74	0.65	8.3189	5048.3225
2	97	73	0.90	7.9503	2208.0011	2	222	72	0.65	8.3196	5052.6777
2	100	69	0.35	7.9533	2291.1564	2	222	70	0.65	8.3198	5057.3573
2	100	68	0.35	7.9588	2292.0469	2	231	88	0.90	8.3217	5224.7544
2	100	65	0.35	7.9743	2295.0578	2	231	82	0.90	8.3270	5235.7903
2	100	64	0.35	7.9791	2296.1463	2	231	81	0.90	8.3274	5237.9622
2	100	60	0.35	7.9959	2301.2085	2	231	78	0.90	8.3280	5244.7280
2	100	58	0.35	8.0027	2304.2115	2	232	77	0.90	8.3292	5269.8984
2	100	57	0.35	8.0055	2305.8556	2	233	82	0.90	8.3296	5281.1218
2	100	56	0.35	8.0079	2307.6701	2	233	81	0.90	8.3300	5283.3125
2	100	54	0.35	8.0114	2311.6079	2	233	77	0.90	8.3304	5292.6135
2	100	53	0.35	8.0124	2313.8581	2	234	77	0.90	8.3316	5315.3286
2	100	52	0.35	8.0127	2316.2652	2	240	88	0.90	8.3341	5428.3163
2	107	72	0.65	8.0199	2435.2983	2	240	82	0.90	8.3385	5439.7821
2	107	70	0.65	8.0287	2437.5539	2	240	81	0.90	8.3387	5442.0386
2	107	69	0.65	8.0327	2438.7950	2	240	78	0.90	8.3389	5449.0681
2	107	68	0.65	8.0366	2440.1145	2	241	81	0.70	8.3390	5465.3048
2	107	65	0.65	8.0469	2444.3455	2	241	78	0.70	8.3414	5470.8672
2	107	64	0.65	8.0497	2445.9472	2	241	77	0.70	8.3419	5472.8820
2	107	60	0.65	8.0575	2453.4791	2	241	75	0.70	8.3425	5477.3519
2	107	58	0.65	8.0584	2458.1665	2	261	87	1.00	8.3587	5906.2749
2	121	78	0.90	8.0635	2747.2370	2	268	83	0.90	8.3692	6072.0590
2	121	77	0.90	8.0669	2748.5231	2	290	100	0.90	8.3786	6538.7217
2	121	75	0.90	8.0731	2751.2812	2	290	94	0.90	8.3848	6548.1023
2	121	74	0.90	8.0760	2752.7817	2	290	88	0.90	8.3887	6559.2155
2	121	72	0.90	8.0810	2755.9682	2	290	82	0.90	8.3893	6573.0700
2	121	69	0.90	8.0865	2761.4835	2	293	88	0.90	8.3914	6627.0695
2	121	68	0.90	8.0877	2763.5555	2	305	77	0.90	8.3976	6928.0992
2	121	65	0.90	8.0887	2770.5941	2	310	100	0.90	8.3979	6989.6680
2	127	75	0.65	8.1006	2886.9058	2	310	94	0.90	8.4029	6999.6956
2	128	75	0.90	8.1023	2910.4462	2	310	88	0.90	8.4057	7011.5752
2	128	74	0.90	8.1047	2912.0336	2	355	100	0.90	8.4334	8004.2972
2	128	72	0.90	8.1090	2915.4044	2	355	94	0.90	8.4362	8015.7804
2	128	68	0.90	8.1140	2923.4306	2	357	94	0.90	8.4375	8060.9397
2	129	56	0.35	8.1147	2976.8955	2	357	88	0.90	8.4380	8074.6205
2	133	77	0.65	8.1166	3021.0125	2	365	94	0.90	8.4425	8241.5770
2	133	75	0.65	8.1235	3023.2950	2	365	88	0.90	8.4427	8255.5644
2	133	74	0.65	8.1269	3024.4439	2	371	100	0.90	8.4439	8365.0543
2	133	72	0.65	8.1330	3027.0531	2	371	94	0.90	8.4461	8377.0550
2	133	70	0.65	8.1385	3029.8567	2	380	94	0.90	8.4513	8580.2720
2	133	69	0.65	8.1410	3031.3993	2	389	100	0.90	8.4547	8770.9060
2	133	68	0.65	8.1432	3033.0395	2	389	95	0.90	8.4562	8781.1819
2	133	65	0.65	8.1486	3038.3001	2	389	94	0.90	8.4563	8783.4890
2	133	60	0.65	8.1511	3049.6529	2	391	94	0.90	8.4573	8828.6483
2	147	79	0.90	8.1595	3336.0577	2	392	94	0.90	8.4579	8851.2279
2	153	81	0.90	8.1731	3469.2977						

Table E.14: Pareto–Optimal Solutions with $D = 2$ for $\sigma = 0.1$ and $q = 0.1$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
35	31	0.90	1.8753	127.1952	202	39	0.90	7.8158	634.8245
38	32	0.90	2.2487	134.2682	206	40	0.90	7.8218	638.8094
42	33	0.90	2.7312	144.7044	202	38	0.90	7.8412	643.9364
42	32	0.70	2.8720	163.2214	206	39	0.90	7.8477	647.3953
58	44	0.90	3.0340	168.9063	200	37	0.90	7.8491	647.6156
55	39	0.90	3.4313	169.8965	207	39	0.90	7.8555	650.5380
44	30	0.80	3.5247	172.6169	202	37	0.90	7.8646	654.0918
58	40	0.90	3.6043	176.7670	207	38	0.90	7.8798	659.8754
58	39	0.90	3.7444	179.1636	215	40	0.90	7.8912	666.7185
58	37	0.90	4.0204	184.7209	206	37	0.90	7.8948	667.0441
58	35	0.90	4.2899	191.4523	207	37	0.90	7.9021	670.2821
58	34	0.90	4.4214	195.4144	215	39	0.90	7.9152	675.6795
58	33	0.85	4.5351	204.0413	209	37	0.90	7.9166	676.7583
58	30	0.90	4.9142	217.6605	216	38	0.90	7.9448	688.5657
58	29	0.90	5.0247	225.7213	215	37	0.90	7.9585	696.1868
74	37	0.90	5.1924	235.6784	215	35	0.90	7.9933	721.3374
63	30	0.85	5.2203	245.2353	215	34	0.90	8.0067	736.1158
79	39	0.90	5.2566	248.2729	234	37	0.90	8.0769	757.7103
81	40	0.90	5.2657	251.1823	261	44	0.90	8.0922	773.9638
63	29	0.85	5.3160	254.2610	262	44	0.90	8.0977	776.9291
81	39	0.90	5.3604	254.5583	261	43	0.90	8.1136	781.7019
67	30	0.90	5.4746	255.2908	261	42	0.90	8.1337	790.2795
81	37	0.90	5.5456	262.2843	271	44	0.90	8.1456	803.6176
79	35	0.90	5.6311	265.0496	261	40	0.90	8.1709	809.3653
79	34	0.90	5.7194	270.4798	261	39	0.90	8.1874	820.2435
81	35	0.90	5.7233	271.7597	265	39	0.90	8.2066	832.8143
85	37	0.90	5.7278	275.2366	271	40	0.90	8.2191	840.3754
81	34	0.90	5.8087	277.3273	283	43	0.90	8.2234	847.5925
85	35	0.90	5.8947	285.1799	271	39	0.90	8.2343	851.6705
85	34	0.90	5.9747	291.0225	282	42	0.90	8.2363	853.8652
83	32	0.90	6.0492	297.9668	283	42	0.90	8.2408	856.8931
84	32	0.90	6.0875	301.5568	282	40	0.90	8.2682	874.4866
85	32	0.90	6.1249	305.1467	283	40	0.90	8.2725	877.5876
84	31	0.90	6.1576	310.1310	282	39	0.90	8.2821	886.2401
85	31	0.90	6.1937	313.8230	283	39	0.90	8.2863	889.3828
85	30	0.90	6.2572	323.8763	304	43	0.90	8.3134	910.4880
85	29	0.90	6.3136	335.8063	311	44	0.90	8.3249	922.2327
101	35	0.90	6.4446	338.8608	311	43	0.90	8.3407	931.4532
119	42	0.95	6.5051	354.5080	316	44	0.90	8.3441	937.0596
110	37	0.90	6.5663	356.1886	311	42	0.90	8.3552	941.6740
119	40	0.95	6.6271	362.8811	316	43	0.90	8.3594	946.4283
119	39	0.95	6.6863	367.6119	316	42	0.90	8.3734	956.8135
107	33	0.90	6.7199	374.6307	311	40	0.90	8.3811	964.4161
119	37	0.95	6.8001	378.5494	311	39	0.90	8.3919	977.3783
119	35	0.95	6.9062	391.9717	316	40	0.90	8.3984	979.9212
119	34	0.95	6.9556	399.8521	316	39	0.90	8.4088	993.0918
129	37	0.90	6.9861	417.7121	333	43	0.90	8.4190	997.3438
124	34	0.90	7.0317	424.5505	343	44	0.90	8.4382	1017.1248
129	35	0.90	7.0785	432.8025	343	43	0.90	8.4513	1027.2940
129	34	0.90	7.1210	441.6695	343	42	0.90	8.4630	1038.5665
151	42	0.90	7.1302	457.2115	353	44	0.90	8.4694	1046.7786
119	29	0.95	7.1383	460.8672	355	44	0.90	8.4754	1052.7093
151	40	0.90	7.2185	468.2535	353	43	0.90	8.4817	1057.2443
151	39	0.90	7.2607	474.5470	355	43	0.90	8.4876	1063.2344
135	32	0.90	7.2873	484.6448	353	42	0.90	8.4927	1068.8454
167	44	0.95	7.2936	487.4915	355	42	0.90	8.4984	1074.9012
151	37	0.90	7.3403	488.9498	359	43	0.90	8.4992	1075.2145
167	42	0.95	7.3755	497.5029	359	42	0.90	8.5097	1087.0128
151	35	0.90	7.4117	506.6137	355	40	0.90	8.5170	1100.8608
167	40	0.95	7.4532	509.2534	359	40	0.90	8.5278	1113.2649
167	39	0.95	7.4903	515.8923	359	39	0.90	8.5347	1128.2277
167	37	0.95	7.5596	531.2415	382	43	0.90	8.5610	1144.1001
167	35	0.95	7.6208	550.0780	396	44	0.90	8.5856	1174.2899
183	39	0.85	7.6238	585.9487	393	43	0.90	8.5880	1177.0454
167	32	0.95	7.6906	588.0055	396	43	0.90	8.5951	1186.0304
202	43	0.90	7.7010	604.9953	393	42	0.90	8.5964	1189.9611
206	44	0.90	7.7057	610.8680	396	42	0.90	8.6033	1199.0447
202	42	0.90	7.7314	611.6339	393	40	0.90	8.6100	1218.6994
206	43	0.90	7.7365	616.9754	396	40	0.90	8.6166	1228.0025
207	43	0.90	7.7451	619.9705	396	39	0.90	8.6211	1244.5074
206	42	0.90	7.7659	623.7455	396	38	0.90	8.6242	1262.3704
202	40	0.90	7.7890	626.4053	396	37	0.90	8.6252	1282.2789
195	37	0.90	7.8089	631.4252					

Table E.15: Pareto–Optimal Solutions with $D = 4$ for $\sigma = 0.1$ and $q = 0.1$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
15	14	0.90	0.7306	116.9667	187	18	1.00	4.3440	1131.6741
28	26	0.95	0.7969	146.3892	200	20	0.95	4.3541	1159.9393
24	18	0.95	1.5421	147.5957	213	23	0.95	4.3552	1162.8256
22	15	0.95	1.7935	155.8577	213	22	0.95	4.3657	1182.8624
28	17	0.95	2.1403	178.7617	213	21	0.95	4.3738	1207.6106
41	25	0.90	2.1817	219.9218	222	23	0.95	4.3742	1211.9590
41	24	0.90	2.2834	223.2549	222	22	0.95	4.3838	1232.8425
41	23	0.90	2.3841	226.9657	232	24	0.95	4.3841	1246.3708
34	16	0.95	2.7010	227.3394	233	24	0.95	4.3860	1251.7431
32	14	0.95	2.7783	244.5266	222	21	0.95	4.3911	1258.6364
41	17	0.90	2.9497	266.5662	232	23	0.95	4.3936	1266.5518
41	16	0.90	3.0307	279.3819	243	25	0.95	4.3944	1287.5142
54	22	0.90	3.0422	304.3452	222	20	0.95	4.3968	1287.5326
54	21	0.90	3.1134	310.7936	232	22	0.95	4.4023	1288.3759
52	19	0.90	3.1916	314.9895	237	23	0.95	4.4027	1293.8481
54	19	0.90	3.2495	327.1045	233	22	0.95	4.4041	1293.9292
54	18	0.90	3.3130	337.8956	232	21	0.95	4.4087	1315.3317
54	17	0.90	3.3724	351.0871	233	21	0.95	4.4103	1321.0013
54	16	0.90	3.4257	367.9664	252	25	0.95	4.4104	1335.1999
66	20	0.90	3.4721	388.8944	237	21	0.95	4.4169	1343.6794
66	19	0.90	3.5233	399.7943	252	24	0.95	4.4198	1353.8165
66	18	0.90	3.5712	412.9835	252	23	0.95	4.4278	1375.7373
66	16	0.90	3.6522	449.7367	252	22	0.95	4.4349	1399.4428
86	23	0.95	3.6629	469.4976	260	23	0.95	4.4399	1419.4115
86	22	0.95	3.7050	477.5876	252	20	0.95	4.4430	1461.5235
75	17	0.95	3.7561	478.8267	260	21	0.95	4.4508	1474.0787
86	20	0.95	3.7833	498.7739	269	21	1.00	4.4709	1503.9981
86	19	0.95	3.8187	512.5226	281	22	0.95	4.4739	1560.4898
85	18	0.95	3.8406	522.7357	286	23	0.95	4.4749	1561.3526
86	18	0.95	3.8512	528.8856	291	24	0.95	4.4753	1563.3358
86	17	0.95	3.8792	549.0546	294	24	0.95	4.4790	1579.4526
86	16	0.95	3.9008	575.0351	306	26	1.00	4.4861	1582.8646
99	20	0.90	3.9061	583.3416	306	25	1.00	4.4935	1602.0943
103	21	0.95	3.9190	583.9619	306	24	1.00	4.5001	1623.8486
105	21	0.95	3.9357	595.3010	306	23	1.00	4.5053	1649.0165
110	22	0.95	3.9467	610.8679	306	22	1.00	4.5094	1677.3951
110	21	0.95	3.9750	623.6487	306	21	1.00	4.5117	1710.8677
110	20	0.95	4.0018	637.9666	306	20	1.00	4.5122	1749.3075
110	19	0.95	4.0255	655.5522	332	25	0.95	4.5146	1759.0729
119	21	0.95	4.0374	674.6745	339	26	0.95	4.5162	1772.3652
110	18	0.95	4.0462	676.4815	335	25	0.95	4.5175	1774.9681
119	20	0.95	4.0610	690.1639	332	24	0.95	4.5196	1783.5996
110	17	0.95	4.0624	702.2791	339	25	0.95	4.5214	1796.1618
118	19	0.90	4.0663	714.7838	335	24	0.95	4.5225	1799.7164
133	23	1.00	4.0801	716.7294	332	23	0.95	4.5232	1812.4793
133	22	1.00	4.1036	729.0639	339	24	0.95	4.5261	1821.2056
133	21	1.00	4.1254	743.6124	351	26	0.95	4.5276	1835.1038
133	20	1.00	4.1454	760.3199	339	23	0.95	4.5294	1850.6942
135	20	0.95	4.1468	782.9590	351	25	0.95	4.5324	1859.7427
141	21	0.95	4.1563	799.4042	351	24	0.95	4.5367	1885.6731
135	19	0.95	4.1627	804.5413	356	25	0.95	4.5367	1886.2348
133	18	1.00	4.1767	804.8805	362	26	0.95	4.5375	1892.6142
133	17	1.00	4.1863	835.1050	365	26	0.95	4.5400	1908.2988
151	21	0.95	4.1989	856.0995	356	24	0.95	4.5408	1912.5345
151	20	0.95	4.2144	875.7541	365	25	0.95	4.5443	1933.9205
152	20	0.95	4.2181	881.5538	362	24	0.95	4.5457	1944.7682
157	21	0.95	4.2219	890.1167	375	26	0.95	4.5483	1960.5810
158	21	0.95	4.2255	895.7863	376	26	0.95	4.5491	1965.8092
157	20	0.95	4.2362	910.5523	375	25	0.95	4.5522	1986.9046
158	20	0.95	4.2396	916.3520	376	25	0.95	4.5530	1992.2030
179	24	1.00	4.2576	949.8984	382	26	0.95	4.5539	1997.1785
179	23	1.00	4.2734	964.6207	375	24	0.95	4.5557	2014.6080
174	21	0.95	4.2783	986.4988	376	24	0.95	4.5564	2019.9802
187	24	1.00	4.2826	992.3519	389	26	0.95	4.5592	2033.7760
179	21	1.00	4.3010	1000.8017	382	24	0.95	4.5608	2052.2140
187	22	1.00	4.3109	1025.0748	398	25	0.90	4.5618	2134.8613
187	21	1.00	4.3227	1045.5303	398	24	0.90	4.5637	2167.2155
187	20	1.00	4.3328	1069.0212	398	23	0.90	4.5647	2203.2373
199	21	0.95	4.3438	1128.2371	400	22	0.90	4.5658	2254.4086

Table E.16: Pareto–Optimal Solutions with $D = 2$ for $\sigma = 0.1$ and $q = 0.9$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
33	32	0.95	7.1012	245.3070	179	76	1.00	8.0556	1033.4483
41	40	0.95	7.3404	270.1097	179	75	1.00	8.0578	1034.7290
47	34	1.00	7.3949	333.1892	179	70	1.00	8.0675	1042.0235
53	49	0.80	7.4412	338.7860	179	68	1.00	8.0705	1045.3901
59	58	0.85	7.4762	361.2183	179	67	1.00	8.0718	1047.1864
63	61	0.85	7.5030	382.7137	179	65	1.00	8.0739	1051.0750
63	60	0.85	7.5116	383.6504	179	63	1.00	8.0753	1055.2551
63	59	0.85	7.5198	384.6414	179	61	1.00	8.0759	1059.8911
63	58	0.85	7.5279	385.7077	187	70	1.00	8.0816	1088.5944
63	56	0.85	7.5429	387.9479	196	70	1.00	8.0960	1140.9866
66	64	1.00	7.5620	388.2742	213	75	0.95	8.0996	1240.3549
63	48	0.85	7.5866	400.2863	221	86	1.00	8.1058	1262.6918
63	47	0.85	7.5894	402.3536	221	85	1.00	8.1079	1263.8408
63	46	0.85	7.5915	404.5593	221	82	1.00	8.1137	1267.4084
63	44	0.85	7.5928	409.5500	221	76	1.00	8.1235	1275.9334
70	61	1.00	7.6322	414.4820	221	75	1.00	8.1248	1277.5146
66	43	1.00	7.6700	420.8535	221	70	1.00	8.1299	1286.5206
70	48	1.00	7.7016	433.1317	221	68	1.00	8.1311	1290.6772
70	47	1.00	7.7033	435.3956	221	65	1.00	8.1317	1297.6959
77	60	1.00	7.7047	456.9995	229	76	1.00	8.1336	1322.1210
77	56	1.00	7.7274	461.9959	229	70	1.00	8.1392	1333.0915
77	54	1.00	7.7367	464.9136	229	68	1.00	8.1401	1337.3985
77	48	1.00	7.7532	476.4448	229	65	1.00	8.1403	1344.6713
77	47	1.00	7.7537	478.9352	238	77	1.00	8.1430	1372.4190
84	58	0.95	7.7540	505.0006	239	70	1.00	8.1499	1391.3051
88	61	0.95	7.7664	525.1011	239	68	1.00	8.1505	1395.8002
88	47	0.95	7.8008	551.6826	255	83	1.00	8.1541	1460.9271
96	63	0.95	7.8058	570.2683	260	85	1.00	8.1565	1486.8716
96	61	0.95	7.8146	572.8376	264	83	1.00	8.1634	1512.4892
96	60	0.95	7.8186	574.2269	264	79	1.00	8.1680	1518.7843
97	61	0.95	7.8201	578.8046	264	75	1.00	8.1713	1526.0808
97	60	0.95	7.8240	580.2084	266	76	1.00	8.1724	1535.7389
96	55	0.95	7.8345	582.3037	264	70	1.00	8.1732	1536.8391
96	54	0.95	7.8366	584.1936	266	70	1.00	8.1749	1548.4819
97	56	0.95	7.8369	586.5320	281	79	1.00	8.1833	1616.5848
97	54	0.95	7.8414	590.2789	292	86	1.00	8.1856	1668.3530
97	49	0.95	7.8446	602.1189	292	85	1.00	8.1867	1669.8711
105	67	1.00	7.8506	614.2706	292	83	1.00	8.1889	1672.9048
105	65	1.00	7.8595	616.5516	292	82	1.00	8.1898	1674.5848
105	64	1.00	7.8637	617.7090	292	81	1.00	8.1906	1676.3327
105	63	1.00	7.8677	619.0036	292	79	1.00	8.1922	1679.8675
105	61	1.00	7.8751	621.7231	292	77	1.00	8.1933	1683.8082
105	60	1.00	7.8784	623.1812	292	76	1.00	8.1938	1685.8487
105	54	1.00	7.8923	633.9731	292	75	1.00	8.1942	1687.9379
105	49	1.00	7.8924	646.6391	292	70	1.00	8.1946	1699.8372
115	61	0.95	7.9022	686.2116	310	89	1.00	8.1969	1766.6876
119	65	0.95	7.9057	704.0403	310	86	1.00	8.2000	1771.1966
119	64	0.95	7.9089	705.4300	310	85	1.00	8.2010	1772.8084
119	63	0.95	7.9118	706.8951	310	83	1.00	8.2028	1776.0290
119	61	0.95	7.9171	710.0799	310	82	1.00	8.2035	1777.8127
119	57	0.95	7.9246	717.4567	310	81	1.00	8.2042	1779.6682
119	56	0.95	7.9257	719.5607	310	79	1.00	8.2054	1783.4210
119	54	0.95	7.9269	724.1573	310	77	1.00	8.2062	1787.6046
129	65	0.95	7.9411	763.2034	310	76	1.00	8.2065	1789.7709
129	63	0.95	7.9461	766.2989	319	86	1.00	8.2066	1822.6185
129	61	0.95	7.9502	769.7513	319	85	1.00	8.2075	1824.2770
129	57	0.95	7.9555	777.7472	319	83	1.00	8.2091	1827.5912
129	56	0.95	7.9561	780.0280	319	82	1.00	8.2098	1829.4266
129	54	0.95	7.9561	785.0109	319	81	1.00	8.2104	1831.3360
138	63	0.95	7.9727	819.7616	319	76	1.00	8.2123	1841.7319
146	75	1.00	7.9782	843.9677	322	81	1.00	8.2124	1848.5586
146	70	1.00	7.9933	849.9186	322	79	1.00	8.2134	1852.4566
146	68	1.00	7.9985	852.6645	322	76	1.00	8.2142	1859.0523
146	67	1.00	8.0009	854.1297	322	75	1.00	8.2143	1861.3561
146	65	1.00	8.0051	857.3014	332	86	1.00	8.2155	1896.8945
146	64	1.00	8.0070	858.9107	332	82	1.00	8.2183	1903.9800
146	63	1.00	8.0087	860.7109	332	81	1.00	8.2188	1905.9673
146	61	1.00	8.0115	864.4922	332	76	1.00	8.2202	1916.7868
146	57	1.00	8.0139	873.4266	344	89	1.00	8.2208	1960.4533
155	63	0.95	8.0145	920.7468	344	86	1.00	8.2232	1965.4569
155	61	0.95	8.0164	924.8950	344	85	1.00	8.2238	1967.2455
155	56	0.95	8.0166	937.2429	344	83	1.00	8.2251	1970.8193
167	75	0.95	8.0182	972.4848	344	82	1.00	8.2255	1972.7986
167	70	0.95	8.0293	979.4227	344	81	1.00	8.2259	1974.8576
167	68	0.95	8.0328	982.6771	344	79	1.00	8.2266	1979.0220
167	65	0.95	8.0370	988.0242	344	76	1.00	8.2269	1986.0683
167	63	0.95	8.0389	992.0304	362	89	1.00	8.2317	2063.0352
167	61	0.95	8.0400	996.4998	360	82	1.00	8.2345	2064.5566
167	60	0.95	8.0402	998.9166	362	82	1.00	8.2355	2076.0264
173	70	0.95	8.0414	1014.6116	362	81	1.00	8.2358	2078.1932
173	68	0.95	8.0446	1017.9829	362	76	1.00	8.2361	2089.9905
173	65	0.95	8.0483	1023.5221	395	82	0.95	8.2372	2281.5367
173	63	0.95	8.0499	1027.6722	396	70	1.00	8.2474	2305.2587
179	77	1.00	8.0533	1032.1975					

Table E.17: Pareto–Optimal Solutions with $D = 4$ for $\sigma = 0.1$ and $q = 0.9$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
20	19	1.00	4.1950	142.9385	159	39	1.00	4.6884	892.2140
28	26	0.95	4.2882	176.6527	168	46	1.00	4.6915	917.6848
30	28	0.90	4.2901	187.9799	168	45	1.00	4.6922	921.0810
30	26	0.90	4.3060	192.6414	168	42	1.00	4.6940	931.1127
30	25	0.90	4.3124	195.4275	168	40	1.00	4.6943	939.2653
30	24	0.90	4.3176	198.4739	168	39	1.00	4.6947	942.7167
30	23	0.90	4.3214	201.9226	179	49	1.00	4.6970	969.2203
30	22	0.90	4.3234	206.1104	179	46	1.00	4.6996	977.7713
36	35	1.00	4.3263	206.4702	179	45	1.00	4.7001	981.3899
36	34	1.00	4.3359	207.8107	179	43	1.00	4.7011	988.1975
36	33	1.00	4.3451	209.1466	179	42	1.00	4.7014	992.0784
36	32	1.00	4.3540	210.7434	179	39	1.00	4.7016	1004.4422
36	30	1.00	4.3705	214.2174	187	49	1.00	4.7026	1012.5374
36	29	1.00	4.3780	216.1423	187	46	1.00	4.7049	1021.4706
36	28	1.00	4.3849	218.3544	187	45	1.00	4.7053	1025.2509
36	27	1.00	4.3911	220.6340	187	43	1.00	4.7061	1032.3628
36	26	1.00	4.3964	223.4830	187	42	1.00	4.7062	1036.4171
36	24	1.00	4.4039	229.9165	191	46	1.00	4.7074	1043.3203
36	22	1.00	4.4058	238.4012	198	49	1.00	4.7096	1072.0984
43	34	1.00	4.4095	248.2184	198	46	1.00	4.7115	1081.5571
43	33	1.00	4.4165	249.8140	198	45	1.00	4.7117	1085.5597
43	32	1.00	4.4231	251.7213	198	42	1.00	4.7122	1097.3828
43	29	1.00	4.4404	258.1700	219	58	1.00	4.7138	1161.6927
43	28	1.00	4.4450	260.8122	217	52	1.00	4.7182	1165.9556
43	26	1.00	4.4521	266.9380	217	49	1.00	4.7201	1174.9766
43	24	1.00	4.4552	274.6225	217	46	1.00	4.7213	1185.3429
52	36	1.00	4.4640	296.5489	219	46	1.00	4.7223	1196.2677
52	35	1.00	4.4696	298.2347	219	43	1.00	4.7223	1209.0238
52	34	1.00	4.4749	300.1711	224	49	1.00	4.7235	1212.8790
52	33	1.00	4.4800	302.1007	224	46	1.00	4.7245	1223.5798
52	32	1.00	4.4846	304.4072	238	52	1.00	4.7284	1278.7900
52	30	1.00	4.4927	309.4251	238	49	1.00	4.7297	1288.6840
52	29	1.00	4.4960	312.2056	238	46	1.00	4.7304	1300.0535
52	28	1.00	4.4986	315.4008	264	64	1.00	4.7306	1386.8496
52	26	1.00	4.5016	322.8088	264	62	1.00	4.7325	1390.6721
61	31	1.00	4.5311	359.6015	264	61	1.00	4.7331	1393.8399
67	36	1.00	4.5383	382.0919	264	58	1.00	4.7357	1400.3967
64	29	1.00	4.5458	384.2531	264	55	1.00	4.7375	1408.8282
64	28	1.00	4.5466	388.1856	263	53	1.00	4.7379	1410.3125
64	27	1.00	4.5469	392.2382	264	52	1.00	4.7388	1418.4898
67	30	1.00	4.5542	398.6824	264	49	1.00	4.7395	1429.4646
67	29	1.00	4.5554	402.2649	264	46	1.00	4.7395	1442.0762
79	42	1.00	4.5584	437.8447	272	43	1.00	4.7407	1501.6186
79	40	1.00	4.5654	441.6783	294	64	1.00	4.7426	1544.4461
79	36	1.00	4.5774	450.5263	294	62	1.00	4.7442	1548.7030
79	35	1.00	4.5798	453.0874	294	61	1.00	4.7445	1552.2308
79	29	1.00	4.5867	474.3124	294	59	1.00	4.7457	1557.5254
81	30	1.00	4.5911	481.9891	294	58	1.00	4.7465	1559.5326
90	35	1.00	4.6058	516.1755	294	55	1.00	4.7478	1568.9224
100	43	1.00	4.6097	552.0656	296	55	1.00	4.7484	1579.5953
101	34	1.00	4.6267	583.0245	294	52	1.00	4.7485	1579.6818
116	48	1.00	4.6275	630.1263	296	52	1.00	4.7491	1590.4279
116	46	1.00	4.6323	633.6395	296	49	1.00	4.7492	1602.7330
116	45	1.00	4.6343	635.9845	309	44	1.00	4.7507	1699.2107
116	42	1.00	4.6400	642.9112	317	46	1.00	4.7536	1731.5839
116	40	1.00	4.6430	648.5403	340	69	1.00	4.7538	1774.2927
116	39	1.00	4.6447	650.9234	340	64	1.00	4.7569	1786.0941
116	38	1.00	4.6456	654.3937	340	62	1.00	4.7580	1791.0171
118	40	1.00	4.6458	659.7220	340	58	1.00	4.7594	1803.5411
116	36	1.00	4.6471	661.5322	340	55	1.00	4.7600	1814.4000
116	35	1.00	4.6475	665.2929	340	52	1.00	4.7601	1826.8429
116	34	1.00	4.6476	669.6124	345	49	1.00	4.7606	1868.0503
118	36	1.00	4.6496	672.9380	359	64	1.00	4.7617	1885.9053
118	35	1.00	4.6500	676.7634	359	62	1.00	4.7627	1891.1034
124	42	1.00	4.6513	687.2499	359	58	1.00	4.7638	1904.3273
124	40	1.00	4.6537	693.2672	359	55	1.00	4.7642	1915.7929
124	38	1.00	4.6557	699.5243	363	52	1.00	4.7648	1950.4234
124	36	1.00	4.6567	707.1552	376	64	1.00	4.7656	1975.2100
126	38	1.00	4.6581	710.8070	376	62	1.00	4.7665	1980.6542
134	42	1.00	4.6635	742.6732	376	58	1.00	4.7674	1994.5043
134	40	1.00	4.6653	749.1759	376	55	1.00	4.7675	2006.5129
146	43	1.00	4.6750	806.0158	392	64	1.00	4.7690	2059.2615
148	42	1.00	4.6777	820.2660	392	62	1.00	4.7697	2064.9374
148	40	1.00	4.6789	827.4480	392	58	1.00	4.7704	2079.3768
148	39	1.00	4.6796	830.4885	392	55	1.00	4.7704	2091.8965
159	40	1.00	4.6878	888.9475					

Table E.18: Pareto–Optimal Solutions with $D = 2$ for $\sigma = 0.3$ and $q = 0.1$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
59	58	1.00	2.1063	370.3544	267	62	1.00	12.1493	2230.7494
93	90	1.00	2.8918	395.9297	267	61	1.00	12.1538	2244.2711
87	81	1.00	3.5163	395.9749	281	69	0.85	12.1607	2288.7751
82	75	0.95	3.7064	399.7555	283	70	0.95	12.1792	2292.0601
93	82	1.00	4.4961	420.0662	281	68	0.85	12.1883	2295.3987
88	74	0.95	5.0647	434.3924	283	69	0.95	12.2077	2298.3471
96	81	0.95	5.1804	443.8206	283	68	0.95	12.2342	2305.2475
93	75	1.00	5.7674	450.2873	281	66	0.85	12.2368	2310.5158
102	83	1.00	5.8360	459.5633	281	65	0.85	12.2573	2319.1435
102	82	1.00	6.0050	463.2229	283	66	0.95	12.2795	2321.1616
96	76	0.95	6.0222	468.5994	283	65	0.95	12.2980	2330.3376
102	81	1.00	6.1459	470.9645	283	64	0.95	12.3132	2340.4481
96	75	0.95	6.1879	473.1560	283	62	0.95	12.3324	2363.9356
93	71	1.00	6.4118	473.6980	283	61	0.95	12.3356	2377.5747
93	70	1.00	6.5682	479.4886	293	66	0.80	12.3550	2414.1383
93	69	1.00	6.6526	494.9640	294	65	0.95	12.4179	2420.9161
93	68	1.00	6.8012	501.1260	293	61	0.80	12.4188	2467.1767
102	76	1.00	6.8517	504.7822	303	68	0.95	12.4570	2468.1625
96	70	0.95	6.8945	509.2894	303	66	0.95	12.4949	2485.2013
102	75	1.00	7.0022	509.4007	303	65	0.95	12.5095	2495.0258
94	67	1.00	7.0813	513.1302	303	64	0.95	12.5209	2505.8507
93	66	1.00	7.0839	514.6344	303	62	0.95	12.5323	2530.9982
120	90	1.00	7.2202	520.3065	311	68	0.95	12.5381	2533.3285
93	64	1.00	7.3446	529.9683	309	65	0.95	12.5676	2544.4322
102	71	1.00	7.5711	530.4633	311	66	0.95	12.5732	2550.8172
102	70	1.00	7.7038	536.4886	311	65	0.95	12.5865	2560.9010
102	69	1.00	7.8323	542.8637	311	64	0.95	12.5965	2572.0118
102	68	1.00	7.9561	549.6221	311	63	0.95	12.6028	2584.2762
120	82	1.00	8.2047	566.1850	311	62	0.95	12.6051	2597.8232
102	62	1.00	8.2763	640.4361	319	68	0.90	12.6056	2601.5668
102	61	1.00	8.3685	649.6459	319	67	0.90	12.6239	2609.9115
154	104	1.00	8.8246	650.2395	319	66	0.90	12.6396	2619.0591
116	69	0.95	8.8496	676.8070	319	65	0.90	12.6525	2629.0941
119	71	0.95	8.9015	682.6612	319	64	0.90	12.6623	2640.1140
120	71	1.00	9.0176	684.2333	319	63	0.90	12.6687	2652.2337
120	70	1.00	9.1197	690.0403	327	69	0.95	12.6695	2655.6873
120	69	1.00	9.2181	696.2207	327	68	0.95	12.6884	2663.6606
120	68	1.00	9.3126	702.8104	327	66	0.95	12.7185	2682.0489
119	66	0.95	9.3851	714.7678	327	65	0.95	12.7292	2692.6516
120	66	1.00	9.4882	717.3777	327	64	0.95	12.7366	2704.3340
120	65	1.00	9.5687	725.4454	334	69	1.00	12.7388	2710.1871
120	64	1.00	9.6437	734.1042	334	68	1.00	12.7560	2718.5418
124	67	0.95	9.6816	737.2330	336	68	0.95	12.7666	2736.9723
119	62	0.95	9.6880	749.3706	334	66	1.00	12.7823	2737.9122
120	62	1.00	9.7758	753.4542	334	65	1.00	12.7907	2749.1351
120	61	1.00	9.8317	764.2893	336	66	0.95	12.7941	2755.8668
146	75	1.00	10.5408	808.1136	334	64	1.00	12.7957	2761.5474
179	93	1.00	10.9044	912.5794	336	65	0.95	12.8035	2766.7612
179	90	1.00	11.1564	921.4987	336	64	0.95	12.8095	2778.7652
283	147	0.95	11.3139	1343.3053	336	63	0.95	12.8119	2792.0154
293	150	0.80	11.3432	1418.4865	355	71	0.95	12.8672	2868.0009
319	162	0.90	11.5494	1510.6355	355	70	0.95	12.8865	2875.1990
386	193	1.00	11.8268	1788.8234	355	69	0.95	12.9038	2883.0856
248	64	0.80	11.8396	2058.9891	355	68	0.95	12.9188	2891.7416
248	62	0.80	11.8827	2077.6122	355	66	0.95	12.9412	2911.7045
248	61	0.80	11.8992	2088.2587	355	65	0.95	12.9479	2923.2150
248	60	0.80	11.9119	2099.9237	355	64	0.95	12.9514	2935.8977
253	61	0.70	11.9250	2141.7683	362	66	0.95	12.9914	2969.1184
254	61	0.70	11.9380	2150.2338	375	71	1.00	13.0260	3026.6079
253	60	0.70	11.9387	2153.0325	375	70	1.00	13.0425	3034.3547
267	71	1.00	11.9437	2154.9449	375	69	1.00	13.0569	3042.8747
267	70	1.00	11.9772	2160.4605	375	68	1.00	13.0689	3052.2550
267	69	1.00	12.0087	2166.5268	375	66	1.00	13.0846	3074.0032
267	68	1.00	12.0379	2173.2056	375	65	1.00	13.0877	3086.6038
267	66	1.00	12.0884	2188.6903	382	65	1.00	13.1320	3144.2204
267	65	1.00	12.1092	2197.6619	389	64	0.80	13.1379	3229.6241
267	64	1.00	12.1265	2207.5843	393	65	1.00	13.1985	3234.7608

Table E.19: Pareto–Optimal Solutions with $D = 4$ for $\sigma = 0.3$ and $q = 0.1$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
34	33	0.90	1.3673	311.6469	146	46	1.00	8.8400	1003.8562
54	51	0.90	1.8711	361.7802	146	45	1.00	8.8671	1016.9040
56	49	0.95	2.6455	376.9664	146	44	1.00	8.8902	1030.9933
36	29	0.85	2.7377	387.2301	170	58	0.95	8.9002	1060.4493
63	54	0.85	2.8639	417.6203	146	42	1.00	8.9227	1062.5956
70	59	0.90	3.0928	439.5373	173	58	0.95	8.9686	1079.1631
70	58	0.90	3.2484	442.5826	173	57	0.95	9.0097	1086.8729
70	57	0.90	3.4027	445.7957	170	54	0.95	9.0603	1093.1592
70	54	0.90	3.8566	456.4441	170	52	0.95	9.1292	1112.3626
44	30	0.80	3.9097	458.9648	173	52	0.95	9.1893	1131.9926
70	52	0.90	4.1506	464.5714	170	49	0.95	9.2153	1145.7427
63	44	0.85	4.4406	465.9087	173	50	0.95	9.2465	1153.9439
70	49	0.90	4.5761	478.6461	173	49	0.95	9.2713	1165.9617
60	39	0.85	4.7796	481.9291	173	48	0.95	9.2933	1178.7386
74	50	0.90	4.8682	500.7292	173	46	0.95	9.3275	1207.0901
70	45	0.90	5.1062	502.1384	173	45	0.95	9.3395	1222.7307
66	41	0.90	5.2017	502.8593	186	52	0.90	9.3675	1235.9681
66	40	0.90	5.3200	511.8136	186	50	0.90	9.4123	1260.0875
70	43	0.90	5.3510	516.5293	186	49	0.90	9.4309	1273.2641
74	46	0.90	5.3677	524.0001	186	48	0.90	9.4465	1287.3358
70	42	0.90	5.4672	524.5982	186	46	0.90	9.4687	1318.2854
74	45	0.90	5.4851	530.8320	186	45	0.90	9.4745	1335.3950
69	40	0.85	5.5557	543.9651	216	63	0.85	9.4770	1345.8400
80	48	0.95	5.7001	544.5331	193	48	0.85	9.4819	1358.3725
70	39	0.90	5.7853	553.1735	224	64	0.85	9.5679	1387.6379
74	42	0.90	5.8141	554.5753	230	65	0.90	9.6841	1395.2227
74	40	0.90	6.0102	573.8516	230	64	0.90	9.7141	1402.8300
74	39	0.90	6.0998	584.7834	230	63	0.90	9.7430	1410.7551
88	50	0.95	6.1094	586.3202	230	62	0.90	9.7707	1419.0987
88	49	0.95	6.2102	592.4600	230	61	0.90	9.7975	1427.7210
88	48	0.95	6.3084	598.9864	230	58	0.90	9.8695	1456.3539
88	46	0.95	6.4959	613.4244	230	57	0.90	9.8907	1466.8442
88	45	0.95	6.5847	621.4098	230	54	0.90	9.9441	1501.7264
85	42	0.90	6.5975	637.5124	230	52	0.90	9.9700	1528.3477
88	43	0.95	6.7480	639.8316	252	64	0.90	9.9974	1537.0137
88	42	0.95	6.8244	649.8497	255	61	0.85	10.0345	1608.2307
88	40	0.95	6.9618	672.5298	292	82	1.00	10.0621	1611.7093
100	48	0.95	7.0354	681.3518	298	82	1.00	10.1307	1644.8266
88	37	0.95	7.1200	715.1902	292	75	1.00	10.2632	1649.5250
91	39	0.90	7.1259	719.5900	292	73	1.00	10.3159	1662.0702
109	52	0.90	7.1427	724.3039	298	75	1.00	10.3253	1683.4193
109	50	0.90	7.2903	738.4384	298	73	1.00	10.3762	1696.2224
109	49	0.90	7.3604	746.1601	298	70	1.00	10.4475	1717.3317
103	44	0.90	7.4107	749.7177	292	65	1.00	10.4982	1723.0556
109	48	0.90	7.4277	754.4065	298	67	1.00	10.5125	1740.8813
119	54	0.95	7.5058	765.2114	292	63	1.00	10.5352	1741.6114
109	46	0.90	7.5536	772.5436	292	62	1.00	10.5518	1751.6434
121	54	0.95	7.5915	778.0721	292	61	1.00	10.5675	1761.9951
119	52	0.95	7.6430	778.6539	298	63	1.00	10.5863	1777.3979
119	51	0.95	7.7083	785.9880	298	62	1.00	10.6020	1787.6361
121	52	0.95	7.7248	791.7405	292	58	1.00	10.6068	1796.3109
119	50	0.95	7.7714	793.7533	298	61	1.00	10.6168	1798.2004
119	49	0.95	7.8319	802.0199	298	58	1.00	10.6532	1833.2214
121	50	0.95	7.8495	807.0937	298	57	1.00	10.6622	1846.2234
119	48	0.95	7.8898	810.8086	298	54	1.00	10.6797	1889.2279
121	49	0.95	7.9081	815.4992	298	52	1.00	10.6815	1922.1129
121	48	0.95	7.9641	824.4357	343	65	0.90	10.7724	2080.7017
119	46	0.95	7.9960	830.3106	334	58	0.95	10.8368	2083.4709
119	45	0.95	8.0441	841.0691	334	57	0.95	10.8404	2098.3557
121	46	0.95	8.0665	844.2654	340	57	0.95	10.8751	2136.0507
119	42	0.95	8.1631	878.7741	383	75	0.95	10.9444	2190.8333
121	42	0.95	8.2259	893.5434	383	73	0.95	10.9747	2207.7073
146	58	1.00	8.2979	898.1554	383	65	0.95	11.0657	2290.3113
146	57	1.00	8.3540	904.5256	383	63	0.95	11.0796	2315.4323
146	54	1.00	8.5130	925.5949	383	62	0.95	11.0850	2328.7791
146	52	1.00	8.6097	941.7063	383	61	0.95	11.0890	2342.8882
146	50	1.00	8.6975	959.8511	383	58	0.95	11.0936	2389.1298
146	48	1.00	8.7748	980.4055	398	61	0.90	11.0941	2470.5781

Table E.20: Pareto–Optimal Solutions with $D = 2$ for $\sigma = 0.3$ and $q = 0.9$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
12	12	0.15	3.5861	546.3375	167	74	0.85	8.2252	3101.4186
18	17	0.15	5.1677	546.8462	167	70	0.85	8.2309	3108.5015
22	20	0.15	5.7844	588.8639	167	68	0.85	8.2326	3112.5065
40	37	0.20	7.2235	818.4041	174	65	0.85	8.2467	3250.1812
44	43	0.25	7.4066	866.6762	179	75	1.00	8.2495	3322.5040
59	55	0.85	7.6948	1115.2087	179	74	1.00	8.2505	3324.3987
72	70	0.85	7.7558	1340.1914	179	70	1.00	8.2524	3333.1606
72	68	0.85	7.7715	1341.9181	181	75	0.85	8.2533	3359.7380
72	66	0.85	7.7863	1343.8637	181	74	0.85	8.2547	3361.4177
72	65	0.85	7.7934	1344.9492	181	70	0.85	8.2587	3369.0944
72	61	0.85	7.8178	1349.9078	181	68	0.85	8.2596	3373.4352
72	60	0.85	7.8228	1351.3605	193	83	0.85	8.2621	3570.4078
72	59	0.85	7.8272	1352.9766	193	75	0.85	8.2754	3582.4831
72	58	0.85	7.8311	1354.7173	193	74	0.85	8.2766	3584.2741
72	57	0.85	7.8342	1356.6164	193	70	0.85	8.2794	3592.4598
72	55	0.85	7.8376	1360.9327	193	68	0.85	8.2797	3597.0884
72	54	0.85	7.8377	1363.4563	211	75	0.60	8.2873	3926.4487
75	61	0.85	7.8467	1406.1540	211	74	0.60	8.2886	3928.0996
76	47	0.75	7.8526	1461.2977	211	70	0.60	8.2930	3935.1506
82	68	0.55	7.8537	1532.3326	211	68	0.60	8.2943	3939.1654
82	67	0.55	7.8608	1533.1597	211	65	0.60	8.2949	3946.0566
82	66	0.55	7.8678	1534.0174	211	66	0.60	8.2949	3943.6401
82	65	0.55	7.8745	1534.8956	215	82	0.85	8.3016	3978.8397
82	61	0.55	7.8992	1539.0160	215	75	0.85	8.3097	3990.8490
82	60	0.55	7.9047	1540.2262	215	74	0.85	8.3103	3992.8443
82	58	0.55	7.9147	1542.8547	215	70	0.85	8.3112	4001.9630
82	55	0.55	7.9264	1547.5694	222	70	0.65	8.3121	4137.2120
82	54	0.55	7.9293	1549.4039	222	68	0.65	8.3127	4141.6381
86	56	0.85	7.9440	1622.8717	224	83	0.85	8.3134	4143.8930
94	58	0.85	7.9916	1768.6587	224	82	0.85	8.3149	4145.3958
97	60	0.85	8.0061	1820.5829	224	75	0.85	8.3217	4157.9078
103	70	0.85	8.0071	1917.2182	224	74	0.85	8.3222	4159.9866
103	68	0.85	8.0155	1919.6885	224	70	0.85	8.3225	4169.4870
103	66	0.85	8.0228	1922.4717	225	75	0.85	8.3230	4176.4699
103	65	0.85	8.0261	1924.0246	225	74	0.85	8.3235	4178.5579
103	64	0.85	8.0290	1925.6067	225	70	0.85	8.3237	4188.1008
103	61	0.85	8.0356	1931.1181	236	82	0.80	8.3287	4368.9233
103	58	0.85	8.0375	1937.9984	236	74	0.80	8.3354	4383.6492
105	61	0.60	8.0433	1969.1950	236	70	0.80	8.3356	4393.1992
105	59	0.60	8.0488	1972.5033	238	83	1.00	8.3359	4400.4692
105	58	0.60	8.0509	1974.3311	238	74	1.00	8.3387	4420.1502
105	57	0.60	8.0527	1976.2596	243	82	0.85	8.3396	4497.0142
105	55	0.60	8.0545	1980.6412	243	75	0.85	8.3443	4510.5875
105	54	0.60	8.0546	1983.1211	243	74	0.85	8.3445	4512.8426
111	65	0.85	8.0650	2073.4634	248	74	0.80	8.3485	4606.5466
119	70	0.55	8.0665	2221.5635	254	85	0.85	8.3492	4695.4470
119	66	0.55	8.0819	2226.1961	260	75	0.65	8.3505	4834.1725
122	53	0.85	8.0917	2315.1104	260	74	0.65	8.3509	4836.3146
131	70	0.85	8.1319	2438.4039	260	70	0.65	8.3516	4845.3834
131	59	0.85	8.1419	2461.6671	260	66	0.85	8.3545	4852.8441
135	65	0.85	8.1543	2521.6923	265	83	0.85	8.3629	4902.3734
138	68	0.70	8.1560	2573.2517	265	82	0.85	8.3637	4904.1512
138	66	0.70	8.1603	2576.4763	265	75	0.85	8.3663	4918.9535
138	65	0.70	8.1622	2578.2161	272	86	0.85	8.3673	5026.5765
138	61	0.70	8.1668	2586.1470	275	83	0.80	8.3704	5089.0919
138	59	0.70	8.1669	2591.0254	275	82	0.80	8.3712	5090.9063
146	75	1.00	8.1680	2709.9736	275	75	0.80	8.3737	5105.7237
146	74	1.00	8.1701	2711.5190	275	74	0.80	8.3737	5108.0658
146	72	1.00	8.1737	2714.8753	289	82	0.75	8.3805	5352.6773
146	70	1.00	8.1764	2718.6673	298	96	0.85	8.3811	5491.4676
146	68	1.00	8.1781	2722.8803	298	75	1.00	8.3936	5531.3195
146	66	1.00	8.1786	2727.5646	301	83	0.85	8.3953	5568.3562
152	68	0.85	8.1981	2832.9401	301	82	0.85	8.3957	5570.3756
153	58	0.75	8.2004	2876.1001	319	93	0.90	8.4051	5880.6255
158	74	0.85	8.2035	2934.2762	319	91	0.90	8.4068	5883.8741
158	70	0.85	8.2104	2940.9774	319	89	0.90	8.4081	5887.5163
158	68	0.85	8.2127	2944.7666	319	83	0.90	8.4106	5899.6129
160	68	0.70	8.2141	2983.4802	319	82	0.90	8.4107	5901.8547
160	66	0.70	8.2167	2987.2189	344	93	0.95	8.4256	6339.5227
160	65	0.70	8.2176	2989.2361	344	91	0.95	8.4267	6343.2073
160	61	0.70	8.2187	2998.4313	358	91	0.85	8.4311	6605.7926
161	61	0.75	8.2209	3017.1244	394	83	0.85	8.4514	7288.8118
163	68	0.85	8.2240	3037.9555					

Table E.21: Pareto–Optimal Solutions with $D = 4$ for $\sigma = 0.3$ and $q = 0.9$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
29	28	0.90	6.2569	566.8349	153	75	1.00	7.9481	2233.2192
41	39	0.90	7.0068	676.2686	155	77	1.00	7.9523	2257.8913
45	42	0.90	7.1211	725.8764	155	75	1.00	7.9534	2262.4116
46	44	0.90	7.1576	733.0485	155	74	1.00	7.9536	2264.8660
48	42	0.90	7.1669	774.2682	158	82	1.00	7.9550	2291.5198
51	44	0.90	7.2319	812.7277	158	79	1.00	7.9586	2297.4021
54	51	0.90	7.3082	834.2926	158	77	1.00	7.9602	2301.5924
57	51	0.90	7.3482	880.6422	158	76	1.00	7.9607	2303.9495
60	58	0.90	7.3808	908.6391	158	75	1.00	7.9611	2306.2002
63	61	0.90	7.4092	947.9864	158	74	1.00	7.9611	2308.7021
63	60	0.90	7.4128	949.9002	172	83	0.95	7.9675	2503.9047
63	58	0.90	7.4184	954.0710	172	82	0.95	7.9685	2505.8595
63	56	0.90	7.4218	958.7782	172	79	0.95	7.9706	2512.2362
64	53	0.90	7.4320	982.2631	172	77	0.95	7.9711	2517.0146
66	64	0.90	7.4334	987.6190	176	85	0.95	7.9747	2558.1539
66	63	0.90	7.4378	989.3630	176	83	0.95	7.9767	2562.1350
66	62	0.90	7.4416	991.1537	176	82	0.95	7.9776	2564.1353
66	61	0.90	7.4453	993.1286	176	81	0.95	7.9783	2566.2918
66	60	0.90	7.4483	995.1335	179	98	1.00	7.9792	2570.3982
66	58	0.90	7.4526	999.5030	176	79	0.95	7.9793	2570.6603
66	57	0.90	7.4539	1001.8760	179	96	1.00	7.9838	2572.9970
68	56	0.90	7.4751	1034.8717	179	93	1.00	7.9902	2577.1130
71	66	0.90	7.4815	1058.9385	179	91	1.00	7.9940	2580.1204
73	63	0.90	7.5130	1094.2954	179	85	1.00	8.0032	2590.0586
75	55	0.90	7.5361	1144.4699	179	83	1.00	8.0052	2594.0675
76	56	1.00	7.5938	1144.8285	179	82	1.00	8.0061	2596.0889
80	71	1.00	7.5949	1172.9996	179	81	1.00	8.0068	2598.1710
80	69	1.00	7.6038	1176.0771	179	79	1.00	8.0077	2602.7530
80	64	1.00	7.6205	1184.9753	179	77	1.00	8.0081	2607.5002
80	63	1.00	7.6226	1187.0144	182	81	1.00	8.0131	2641.7158
80	61	1.00	7.6258	1191.5313	199	84	1.00	8.0438	2881.7215
80	60	1.00	7.6266	1193.9218	209	77	1.00	8.0597	3044.5114
90	58	1.00	7.6927	1349.0562	229	106	1.00	8.0676	3276.8281
96	71	0.95	7.6962	1414.3063	229	105	1.00	8.0693	3278.3783
96	69	0.95	7.7011	1417.9353	229	98	1.00	8.0801	3288.3865
96	65	0.95	7.7075	1426.4422	229	96	1.00	8.0826	3291.7112
96	64	0.95	7.7082	1428.7395	229	93	1.00	8.0858	3296.9770
96	63	0.95	7.7085	1431.3329	229	83	1.00	8.0902	3318.6674
105	76	0.90	7.7114	1546.0900	232	91	1.00	8.0918	3344.0667
105	75	0.90	7.7144	1547.8077	236	96	1.00	8.0930	3392.3312
105	74	0.90	7.7171	1549.4569	238	98	1.00	8.0937	3417.6244
105	71	0.90	7.7238	1554.9436	238	93	1.00	8.0987	3426.5525
105	66	0.90	7.7298	1566.0358	250	105	1.00	8.1018	3579.0156
105	64	0.90	7.7298	1571.2121	250	98	1.00	8.1104	3589.9416
108	67	0.90	7.7438	1608.2930	250	96	1.00	8.1123	3593.5712
111	83	1.00	7.7668	1608.6117	250	81	1.00	8.1146	3628.7305
111	82	1.00	7.7706	1609.8652	277	114	1.00	8.1248	3952.6794
111	79	1.00	7.7812	1613.9977	277	113	1.00	8.1263	3953.9326
111	77	1.00	7.7875	1616.9415	277	106	1.00	8.1355	3963.6742
111	75	1.00	7.7930	1620.1786	277	105	1.00	8.1364	3965.5493
111	74	1.00	7.7955	1621.9363	277	98	1.00	8.1427	3977.6553
111	73	1.00	7.7977	1623.6915	277	96	1.00	8.1438	3981.6769
111	71	1.00	7.8015	1627.5370	277	93	1.00	8.1450	3988.0464
111	69	1.00	7.8042	1631.8070	292	111	1.00	8.1466	4170.7147
111	64	1.00	7.8055	1644.1532	292	106	1.00	8.1521	4178.3136
118	82	1.00	7.8074	1711.3882	292	105	1.00	8.1529	4180.2902
118	77	1.00	7.8219	1718.9108	292	98	1.00	8.1580	4193.0518
118	76	1.00	7.8243	1720.6712	292	96	1.00	8.1588	4197.2912
118	74	1.00	7.8285	1724.2206	292	93	1.00	8.1596	4204.0056
118	64	1.00	7.8339	1747.8385	292	91	1.00	8.1596	4208.9115
143	93	0.90	7.8360	2077.8511	350	143	1.00	8.1613	4959.8613
146	106	1.00	7.8449	2089.1568	350	133	1.00	8.1761	4969.3807
146	105	1.00	7.8489	2090.1451	343	81	1.00	8.1883	4978.6182
143	83	0.90	7.8614	2092.2729	350	119	1.00	8.1930	4986.6821
143	82	0.90	7.8633	2093.9008	350	118	1.00	8.1939	4988.2202
146	100	1.00	7.8677	2094.4486	350	114	1.00	8.1975	4994.3603
146	98	1.00	7.8748	2096.5259	350	113	1.00	8.1984	4995.9437
146	96	1.00	7.8816	2098.6456	350	106	1.00	8.2030	5008.2526
146	93	1.00	7.8912	2102.0028	350	105	1.00	8.2033	5010.6219
146	89	1.00	7.9029	2106.8791	350	101	1.00	8.2046	5018.9711
146	83	1.00	7.9172	2115.8316	350	98	1.00	8.2050	5025.9183
146	82	1.00	7.9192	2117.4804	359	105	1.00	8.2096	5139.4664
146	79	1.00	7.9242	2122.9159	359	98	1.00	8.2109	5155.1562
146	77	1.00	7.9267	2126.7879	364	98	1.00	8.2140	5226.9550
146	76	1.00	7.9276	2128.9660	379	98	1.00	8.2230	5442.3515
146	75	1.00	7.9284	2131.0457	387	109	1.00	8.2263	5531.9175
146	74	1.00	7.9290	2133.3577	389	106	1.00	8.2287	5566.3150
146	73	1.00	7.9294	2135.6663					

Table E.22: Pareto–Optimal Solutions with $D = 2$ for $\sigma = 0.6$ and $q = 0.1$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
43	42	0.30	1.6945	556.4200	270	64	0.55	12.1439	3138.0416
44	42	0.40	1.9839	566.3662	274	66	0.75	12.1719	3160.8198
46	44	0.30	2.0173	573.1780	270	61	0.55	12.1835	3170.9931
48	46	0.30	2.0570	577.6958	288	65	0.45	12.2638	3358.7122
45	42	0.30	2.2419	582.3000	290	65	0.90	12.3041	3376.2918
54	51	0.55	2.4682	584.3526	288	58	0.45	12.3195	3444.0699
68	66	0.90	2.5116	594.1586	308	75	0.90	12.3359	3488.5127
44	39	0.30	2.6756	606.6104	309	75	0.80	12.3521	3498.2727
71	66	0.80	3.2793	612.8076	309	73	0.80	12.4112	3508.6465
85	78	0.90	3.9628	633.6218	309	71	0.80	12.4623	3521.4579
96	88	0.95	4.3198	649.8116	309	70	0.80	12.4843	3528.9593
82	72	0.55	4.4209	668.7963	309	69	0.80	12.5035	3537.3030
82	71	0.55	4.6129	675.4179	309	68	0.80	12.5196	3546.5906
82	70	0.55	4.7990	682.5016	309	67	0.80	12.5323	3556.9294
73	60	0.55	4.8711	682.7800	309	66	0.80	12.5413	3568.4441
82	69	0.55	4.9832	689.3468	309	65	0.80	12.5461	3581.2693
82	68	0.55	5.1565	697.4123	309	64	0.80	12.5463	3595.5581
82	67	0.55	5.3309	704.7924	328	75	0.80	12.5521	3713.3769
82	66	0.55	5.4999	712.4999	328	72	0.80	12.6291	3730.8286
119	105	0.95	5.5705	720.9631	328	71	0.80	12.6505	3737.9877
107	91	0.80	5.8016	722.6116	328	70	0.80	12.6695	3745.9503
95	76	0.90	6.2236	728.1491	328	69	0.80	12.6856	3754.8070
107	88	0.80	6.2866	741.7578	328	68	0.80	12.6987	3764.6657
107	86	0.80	6.6137	752.4452	328	67	0.80	12.7084	3775.6403
104	82	0.90	6.7303	758.7731	328	66	0.80	12.7142	3787.8630
119	97	0.95	6.8484	767.0926	328	65	0.80	12.7158	3801.4767
96	72	0.95	6.8802	778.9473	337	71	0.80	12.7323	3840.5544
107	82	0.80	7.1620	784.4557	337	70	0.80	12.7499	3848.7356
119	94	0.95	7.2076	795.8840	337	69	0.80	12.7648	3857.8353
119	91	0.95	7.6550	810.6303	337	68	0.80	12.7765	3867.9645
129	99	0.90	7.8561	843.7765	337	67	0.80	12.7848	3879.2402
129	97	0.90	8.1429	852.2048	337	66	0.80	12.7893	3891.7982
105	70	0.60	8.1853	886.2183	337	65	0.80	12.7895	3905.7856
119	83	0.95	8.4487	890.0121	337	64	0.55	12.7987	3916.7408
119	82	0.95	8.5680	896.4239	337	61	0.55	12.8018	3957.8692
129	86	0.90	9.2194	947.3741	352	72	0.95	12.8164	4012.4344
146	99	1.00	9.4959	1003.7471	352	71	0.95	12.8287	4022.1635
158	106	0.90	9.7807	1069.4030	352	70	0.95	12.8376	4033.0821
129	71	0.90	9.8320	1152.2513	352	69	0.95	12.8426	4045.3356
129	70	0.90	9.8967	1161.7738	352	68	0.95	12.8433	4059.0850
129	69	0.90	9.9553	1172.0130	358	72	0.70	12.8874	4072.8371
129	68	0.90	10.0072	1183.0291	358	71	0.70	12.9070	4079.8697
129	67	0.90	10.0520	1194.8884	358	70	0.70	12.9245	4087.6294
129	66	0.90	10.0894	1207.6636	358	69	0.70	12.9397	4096.1936
129	65	0.90	10.1189	1221.4370	358	68	0.70	12.9523	4105.6540
146	82	1.00	10.1727	1224.8247	358	67	0.70	12.9620	4116.1075
146	77	1.00	10.5696	1256.6696	358	66	0.70	12.9685	4127.6678
146	75	1.00	10.6973	1272.9974	358	65	0.70	12.9716	4140.4548
158	83	0.90	10.8842	1319.9918	368	71	0.85	12.9769	4196.1210
158	82	0.90	10.9654	1325.0752	368	70	0.85	12.9887	4205.7301
158	80	0.90	11.1184	1336.3175	368	69	0.85	12.9973	4216.4556
178	89	0.90	11.5383	1459.6192	368	68	0.85	13.0025	4228.4293
205	105	0.90	11.5662	1637.4053	368	67	0.85	13.0038	4241.8019
209	105	0.95	11.7683	1667.2308	376	70	0.90	13.0332	4301.7931
260	130	0.65	11.8491	2067.0196	377	70	0.90	13.0401	4313.2340
298	150	1.00	11.9145	2338.8295	381	72	0.90	13.0485	4338.5490
260	68	0.65	11.9418	2982.8238	381	71	0.90	13.0596	4348.1980
260	67	0.65	11.9699	2990.0376	381	70	0.90	13.0676	4358.9978
260	66	0.65	11.9953	2997.9817	382	70	0.90	13.0744	4370.4387
260	65	0.65	12.0177	3006.7333	382	68	0.55	13.0847	4394.1060
260	64	0.65	12.0367	3016.3808	383	66	0.70	13.1369	4415.9128
260	61	0.65	12.0704	3051.7311	394	71	0.90	13.1464	4496.5618
267	65	0.70	12.1080	3087.9928	394	70	0.90	13.1530	4507.7300
267	64	0.70	12.1232	3098.5462	394	69	0.90	13.1561	4520.2314
270	65	0.55	12.1242	3128.8859	392	65	0.55	13.1679	4542.6787
267	61	0.70	12.1424	3137.5299					

Table E.23: Pareto–Optimal Solutions with $D = 4$ for $\sigma = 0.6$ and $q = 0.1$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
34	33	0.85	1.4216	514.2647	146	59	1.00	10.9705	1304.2975
36	34	0.60	1.7352	526.5051	146	58	1.00	10.9879	1317.9969
46	44	0.60	1.9166	545.2758	167	76	0.95	11.1268	1326.9692
48	46	0.60	1.9421	551.6055	167	75	0.95	11.1814	1334.4123
45	42	0.60	2.1434	551.8316	179	85	1.00	11.2423	1346.1325
37	33	0.60	2.2281	556.5068	167	73	0.95	11.2838	1350.1546
54	51	0.60	2.2360	580.5064	179	83	1.00	11.3657	1358.0098
38	32	0.60	2.6662	588.8827	179	81	1.00	11.4825	1370.6911
54	49	0.60	2.6785	595.2342	179	79	1.00	11.5923	1384.2479
48	41	0.60	3.0516	599.4393	179	76	1.00	11.7421	1406.4834
49	40	0.90	3.5575	600.1779	179	75	1.00	11.7878	1414.4449
54	43	0.60	3.8499	650.7778	179	73	1.00	11.8723	1431.3049
60	49	0.60	3.8717	661.3730	179	71	1.00	11.9469	1449.5298
51	39	0.60	3.9826	662.4269	179	70	1.00	11.9802	1459.1964
60	48	0.60	4.0471	670.2267	179	69	1.00	12.0107	1469.2728
55	41	0.60	4.3428	686.8730	179	68	1.00	12.0382	1479.7743
58	43	0.60	4.5214	699.0014	179	67	1.00	12.0625	1490.7464
85	70	0.90	4.7394	701.7309	179	66	1.00	12.0837	1502.1956
85	69	0.90	4.9072	706.6537	179	63	1.00	12.1261	1539.8051
85	67	0.90	5.2339	717.1723	179	59	1.00	12.1265	1599.1044
85	66	0.90	5.3929	722.7366	219	85	0.95	12.6663	1666.0603
81	61	0.85	5.5325	727.5087	238	97	1.00	12.8389	1712.8771
88	68	0.95	5.5662	728.2892	219	77	0.95	12.9344	1730.7570
85	63	0.90	5.8486	741.0920	238	92	1.00	13.0618	1741.7238
85	61	0.90	6.1348	754.6183	238	91	1.00	13.1030	1748.0006
76	47	0.75	6.5449	827.1522	238	88	1.00	13.2189	1767.9611
74	44	0.80	6.6342	839.7195	238	85	1.00	13.3223	1789.8298
107	76	0.80	6.6519	875.2443	238	83	1.00	13.3837	1805.6220
107	75	0.80	6.7779	880.2251	238	81	1.00	13.4383	1822.4832
107	73	0.80	7.0235	890.7210	238	80	1.00	13.4630	1831.3485
107	71	0.80	7.2523	903.2512	238	79	1.00	13.4859	1840.5084
107	70	0.80	7.3671	909.1713	238	76	1.00	13.5423	1870.0730
107	69	0.80	7.4793	915.3168	238	75	1.00	13.5569	1880.6585
107	68	0.80	7.5888	921.7102	238	73	1.00	13.5790	1903.0758
107	67	0.80	7.6955	928.3600	238	71	1.00	13.5911	1927.3078
107	66	0.80	7.7993	935.2836	238	70	1.00	13.5931	1940.1606
107	65	0.80	7.9002	942.4848	273	92	0.90	13.7096	2046.7010
107	64	0.80	7.9978	950.0058	298	114	1.00	13.7304	2051.5702
107	63	0.80	8.0921	957.8463	273	91	0.90	13.7382	2054.0226
107	62	0.80	8.1830	966.0370	273	88	0.90	13.8168	2077.2806
107	59	0.80	8.4333	992.9325	273	85	0.90	13.8834	2102.7212
107	58	0.80	8.5087	1002.7651	298	102	1.00	14.1825	2113.1506
107	57	0.80	8.5797	1013.0775	298	97	1.00	14.3364	2144.6949
107	54	0.80	8.7645	1047.3238	298	92	1.00	14.4649	2180.8138
146	91	1.00	8.8521	1058.9123	298	91	1.00	14.4871	2188.6731
146	89	1.00	9.0488	1067.2062	298	88	1.00	14.5463	2213.6656
146	88	1.00	9.1452	1071.5482	298	85	1.00	14.5929	2241.0474
146	85	1.00	9.4255	1085.3893	298	83	1.00	14.6163	2260.8208
146	83	1.00	9.6046	1095.3777	298	82	1.00	14.6255	2271.2025
146	81	1.00	9.7769	1106.0260	298	81	1.00	14.6330	2281.9327
146	79	1.00	9.9420	1117.4055	298	79	1.00	14.6426	2304.5021
146	76	1.00	10.1745	1136.0482	350	115	1.00	14.7705	2404.3906
146	75	1.00	10.2477	1142.7240	350	114	1.00	14.8015	2409.5623
146	73	1.00	10.3869	1156.8429	350	111	1.00	14.8908	2425.7935
146	71	1.00	10.4478	1182.2981	350	97	1.00	15.2190	2518.9369
146	70	1.00	10.5096	1190.1826	350	92	1.00	15.2918	2561.3585
146	69	1.00	10.5685	1198.4013	350	91	1.00	15.3029	2570.5892
146	67	1.00	10.6775	1215.9161	350	88	1.00	15.3285	2599.9428
146	66	1.00	10.7272	1225.2545	350	85	1.00	15.3417	2632.1026
146	64	1.00	10.8163	1245.2332	350	83	1.00	15.3428	2655.3264
179	92	1.00	10.8657	1293.3781	378	91	0.90	15.3737	2844.0313
179	91	1.00	10.9372	1298.2555					

Table E.24: Pareto–Optimal Solutions with $D = 2$ for $\sigma = 0.6$ and $q = 0.9$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
20	17	0.15	5.1943	677.9470	200	80	0.85	8.2829	4369.3791
26	24	0.15	6.3690	706.4151	200	75	0.85	8.2871	4379.0715
28	26	0.15	6.5685	736.0641	204	75	0.70	8.2932	4466.8225
29	28	0.15	6.7033	743.0282	204	68	0.70	8.2961	4482.0352
40	37	0.20	7.2943	941.9469	214	82	0.85	8.3034	4671.6492
44	43	0.25	7.4570	1006.0878	214	75	0.85	8.3082	4685.6065
53	51	0.85	7.5599	1196.5166	224	88	0.85	8.3107	4880.1638
58	55	0.25	7.6093	1302.5578	224	85	0.85	8.3149	4884.8374
68	67	0.90	7.7224	1496.9358	224	84	0.85	8.3161	4886.4504
69	58	0.85	7.7804	1532.5876	224	82	0.85	8.3182	4889.9506
74	58	0.85	7.8287	1643.6447	224	81	0.85	8.3191	4891.8280
82	68	0.55	7.8649	1802.9216	224	80	0.85	8.3199	4893.7046
82	66	0.55	7.8783	1804.7541	224	75	0.85	8.3217	4904.5601
82	65	0.55	7.8847	1805.7568	230	74	0.90	8.3271	5039.7590
82	64	0.55	7.8909	1806.7762	234	85	0.90	8.3291	5102.7577
82	62	0.55	7.9024	1809.0732	234	84	0.90	8.3300	5104.5858
82	58	0.55	7.9213	1814.5766	234	83	0.90	8.3308	5106.4274
82	53	0.55	7.9333	1824.2790	234	82	0.90	8.3315	5108.3918
88	65	0.70	7.9341	1937.0619	234	80	0.90	8.3324	5112.6555
88	64	0.70	7.9386	1938.4304	234	75	0.90	8.3324	5124.5646
88	62	0.70	7.9465	1941.4039	239	79	0.90	8.3388	5224.0960
88	59	0.70	7.9550	1946.7797	242	85	0.85	8.3393	5277.3690
88	56	0.70	7.9578	1953.5747	242	82	0.85	8.3417	5282.8930
96	65	0.95	7.9675	2118.6615	242	80	0.85	8.3428	5286.9487
96	62	0.95	7.9686	2125.8633	242	75	0.85	8.3432	5298.6766
105	77	0.60	7.9788	2298.9690	247	85	0.90	8.3456	5386.2442
105	75	0.60	7.9898	2300.5876	247	84	0.90	8.3463	5388.1739
105	74	0.60	7.9951	2301.4619	247	83	0.90	8.3469	5390.1179
105	68	0.60	8.0241	2307.6132	247	82	0.90	8.3474	5392.1914
105	67	0.60	8.0283	2308.8115	247	80	0.90	8.3479	5396.6919
105	65	0.60	8.0361	2311.4616	248	75	0.80	8.3503	5429.5861
105	64	0.60	8.0396	2312.8672	251	84	0.85	8.3510	5475.4422
105	62	0.60	8.0457	2315.9792	251	83	0.85	8.3517	5477.3298
105	59	0.60	8.0522	2321.4593	251	82	0.85	8.3522	5479.3643
105	58	0.60	8.0535	2323.5155	251	81	0.85	8.3526	5481.4679
105	56	0.60	8.0543	2328.2545	251	80	0.85	8.3530	5483.5708
115	65	0.85	8.0742	2534.2636	252	75	0.85	8.3537	5517.6301
117	65	0.90	8.0774	2580.0179	254	75	0.70	8.3555	5561.6319
119	58	0.55	8.1088	2633.3502	260	75	0.65	8.3594	5694.4639
126	68	0.85	8.1137	2770.2611	263	84	0.85	8.3641	5737.2163
126	67	0.85	8.1150	2772.2763	263	81	0.90	8.3650	5743.8216
126	65	0.85	8.1164	2776.6715	266	85	0.85	8.3666	5800.7445
144	80	0.60	8.1394	3149.7662	266	84	0.85	8.3673	5802.6599
146	82	1.00	8.1470	3188.1705	266	83	0.85	8.3678	5804.6602
146	79	1.00	8.1540	3192.6242	266	82	0.85	8.3681	5806.8163
146	77	1.00	8.1579	3195.8998	266	81	0.85	8.3683	5809.0457
146	75	1.00	8.1610	3199.5161	266	80	0.85	8.3685	5811.2742
147	79	0.85	8.1612	3212.8265	276	88	0.90	8.3754	6012.6248
146	68	1.00	8.1635	3215.6246	276	85	0.90	8.3769	6018.6373
147	75	0.85	8.1707	3218.6158	276	82	0.90	8.3775	6025.2827
148	75	0.90	8.1721	3241.1758	282	77	0.85	8.3823	6168.7003
148	68	0.90	8.1789	3255.4467	287	75	0.70	8.3848	6284.2062
153	77	0.75	8.1839	3346.8969	289	82	0.75	8.3876	6310.2447
152	65	0.85	8.1919	3349.6370	291	75	0.85	8.3881	6371.5491
153	68	0.75	8.1999	3361.9060	293	81	0.90	8.3922	6399.0104
153	67	0.75	8.2007	3363.9963	299	85	0.85	8.3971	6520.3857
153	65	0.75	8.2012	3368.6732	299	82	0.85	8.3975	6527.2108
158	68	0.65	8.2123	3471.6448	308	88	1.00	8.4036	6710.0263
158	67	0.65	8.2134	3473.5767	310	88	0.85	8.4050	6753.7982
158	65	0.65	8.2149	3477.7798	310	85	0.85	8.4058	6760.2661
158	62	0.65	8.2151	3484.9573	310	84	0.85	8.4060	6762.4983
173	85	0.90	8.2184	3772.5516	310	83	0.85	8.4060	6764.8296
173	84	0.90	8.2206	3773.9032	312	85	0.85	8.4074	6803.8807
173	83	0.90	8.2228	3775.2647	318	85	0.85	8.4118	6934.7246
173	82	0.90	8.2248	3776.7170	318	84	0.85	8.4119	6937.0144
173	80	0.90	8.2284	3779.8692	328	85	0.85	8.4188	7152.7977
173	79	0.90	8.2301	3781.4586	340	85	0.70	8.4219	7419.0439
173	77	0.90	8.2329	3784.9216	340	84	0.70	8.4222	7421.0647
173	75	0.90	8.2351	3788.6738	340	82	0.70	8.4222	7425.8456
173	68	0.90	8.2358	3805.3533	349	82	0.85	8.4314	7618.7176
179	82	1.00	8.2363	3908.7866	359	96	0.90	8.4375	7803.3899
179	75	1.00	8.2425	3922.6965	359	85	0.90	8.4384	7828.5898
180	75	0.75	8.2520	3940.7869	367	96	0.85	8.4412	7978.6160
183	75	0.85	8.2571	4006.8505	371	96	0.90	8.4450	8064.2274
183	68	0.85	8.2578	4023.4765	371	85	0.90	8.4450	8090.2697
186	75	0.90	8.2612	4073.3718	375	82	0.85	8.4455	8186.3012
196	75	1.00	8.2738	4295.2431	392	85	0.70	8.4505	8553.7212
200	85	0.85	8.2755	4361.4620	397	96	0.85	8.4583	8630.8189
200	82	0.85	8.2803	4366.0273					

Table E.25: Pareto–Optimal Solutions with $D = 4$ for $\sigma = 0.6$ and $q = 0.9$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
5	4	0.05	1.3427	671.7824	146	95	1.00	8.0256	3012.2217
28	26	0.55	6.0748	761.5336	146	93	1.00	8.0331	3014.0237
31	30	0.95	6.2403	814.7712	146	89	1.00	8.0472	3018.0662
33	32	0.80	6.5914	816.1951	146	87	1.00	8.0536	3020.3783
37	34	0.95	6.7351	899.1289	146	85	1.00	8.0597	3022.6752
49	43	0.80	7.2956	1094.7426	146	82	1.00	8.0678	3026.6603
51	49	0.95	7.4204	1106.5099	146	81	1.00	8.0702	3028.2256
56	46	0.80	7.4299	1234.2463	146	80	1.00	8.0725	3029.6612
63	47	0.85	7.5299	1380.5150	146	75	1.00	8.0809	3038.5602
71	70	0.80	7.6003	1494.7662	146	74	1.00	8.0821	3040.3956
71	69	0.80	7.6062	1496.1068	146	70	1.00	8.0837	3049.7202
71	68	0.80	7.6119	1497.5212	166	81	0.85	8.0901	3460.8509
71	67	0.80	7.6171	1498.9445	166	76	0.85	8.0950	3470.9141
71	66	0.80	7.6220	1500.5253	176	96	0.95	8.1080	3635.2117
71	65	0.80	7.6265	1502.1638	176	92	0.95	8.1188	3639.6273
71	64	0.80	7.6305	1503.8537	176	89	0.95	8.1260	3643.3084
71	63	0.80	7.6341	1505.6735	176	85	0.95	8.1341	3649.1331
71	62	0.80	7.6373	1507.6354	176	80	0.95	8.1415	3657.6005
71	60	0.80	7.6418	1511.8223	176	76	0.95	8.1444	3666.0849
71	58	0.80	7.6438	1516.5809	176	75	0.95	8.1447	3668.3073
78	75	0.95	7.6892	1625.7271	179	85	1.00	8.1515	3705.8827
78	70	0.95	7.7183	1631.7317	179	81	1.00	8.1575	3712.6875
78	69	0.95	7.7234	1633.1942	179	80	1.00	8.1588	3714.4476
78	68	0.95	7.7279	1634.6080	179	76	1.00	8.1614	3722.9553
78	67	0.95	7.7322	1636.2248	179	75	1.00	8.1616	3725.3580
78	66	0.95	7.7360	1637.8374	194	89	0.95	8.1662	4015.9195
78	65	0.95	7.7395	1639.6100	194	85	0.95	8.1724	4022.3399
78	64	0.95	7.7426	1641.5353	194	80	0.95	8.1775	4031.6733
78	63	0.95	7.7450	1643.5272	194	75	0.95	8.1783	4043.4751
78	61	0.95	7.7484	1647.9015	198	75	0.95	8.1850	4126.8458
78	60	0.95	7.7491	1650.2750	208	93	0.95	8.1857	4300.2009
84	63	0.80	7.7492	1781.3602	208	89	0.95	8.1926	4305.7281
89	71	0.80	7.7648	1872.2325	208	85	0.95	8.1976	4312.6119
89	70	0.80	7.7684	1873.7210	208	81	0.95	8.2006	4320.6645
89	69	0.80	7.7719	1875.4015	208	80	0.95	8.2012	4322.6188
89	68	0.80	7.7749	1877.1744	238	114	1.00	8.2027	4890.3772
89	67	0.80	7.7776	1878.9586	238	107	1.00	8.2188	4896.6556
89	66	0.80	7.7800	1880.9401	238	106	1.00	8.2212	4897.3054
89	65	0.80	7.7819	1882.9940	238	105	1.00	8.2231	4898.5349
89	60	0.80	7.7845	1895.1012	238	104	1.00	8.2252	4899.6469
91	67	0.80	7.7915	1921.1824	238	96	1.00	8.2399	4909.1041
96	83	0.95	7.8019	1992.2163	238	93	1.00	8.2444	4913.2715
96	76	0.95	7.8382	1999.6826	238	89	1.00	8.2491	4919.8613
96	75	0.95	7.8426	2000.8949	238	85	1.00	8.2523	4927.3747
96	74	0.95	7.8468	2002.2370	238	81	1.00	8.2534	4936.4225
96	72	0.95	7.8544	2005.0973	238	80	1.00	8.2534	4938.7628
96	71	0.95	7.8579	2006.6400	255	89	0.95	8.2601	5278.6571
96	70	0.95	7.8610	2008.2852	255	85	0.95	8.2620	5287.0963
96	69	0.95	7.8639	2010.0852	286	80	0.85	8.2632	5965.9746
96	68	0.95	7.8663	2011.8253	292	114	1.00	8.2790	5999.9586
96	67	0.95	7.8685	2013.8151	292	107	1.00	8.2903	6007.6614
96	66	0.95	7.8702	2015.7999	292	106	1.00	8.2921	6008.4587
96	65	0.95	7.8715	2017.9816	292	105	1.00	8.2933	6009.9672
96	64	0.95	7.8723	2020.3512	292	96	1.00	8.3039	6022.9345
96	62	0.95	7.8724	2025.3530	292	94	1.00	8.3056	6026.3610
107	71	0.80	7.8780	2250.8862	292	93	1.00	8.3064	6028.0473
107	70	0.80	7.8800	2252.6759	292	89	1.00	8.3084	6036.1324
107	69	0.80	7.8817	2254.6962	292	85	1.00	8.3088	6045.3505
107	68	0.80	7.8831	2256.8277	311	93	0.95	8.3126	6429.6273
107	67	0.80	7.8842	2258.9727	319	85	0.95	8.3191	6614.0538
107	66	0.80	7.8848	2261.3550	333	107	0.95	8.3195	6860.0816
107	65	0.80	7.8851	2263.8243	333	106	0.95	8.3206	6861.4376
112	69	0.95	7.9509	2345.0994	333	105	0.95	8.3218	6862.6680
117	74	0.90	7.9519	2444.6034	333	96	0.95	8.3286	6877.9858
117	73	0.90	7.9542	2446.2270	333	93	0.95	8.3295	6884.4562
117	71	0.90	7.9579	2449.9849	333	89	0.95	8.3301	6893.3051
117	70	0.90	7.9594	2452.0870	344	93	0.95	8.3371	7111.8707
117	69	0.90	7.9605	2454.2291	350	107	1.00	8.3426	7200.9640
117	68	0.90	7.9613	2456.3459	350	106	1.00	8.3438	7201.9197
117	67	0.90	7.9617	2458.8227	350	105	1.00	8.3446	7203.7278
117	66	0.90	7.9617	2461.2168	350	96	1.00	8.3507	7219.2708
118	68	0.95	7.9782	2472.8686	350	93	1.00	8.3517	7225.3992
119	70	0.95	7.9806	2489.4368	380	107	0.95	8.3537	7828.3213
127	68	0.95	8.0128	2661.4772	380	106	0.95	8.3544	7829.8687
142	81	0.85	8.0199	2960.4869	380	105	0.95	8.3553	7831.2727
146	96	1.00	8.0217	3011.4672	380	96	0.95	8.3592	7848.7525

Table E.26: Pareto–Optimal Solutions with $D = 2$ for $\sigma = 0.8$ and $q = 0.1$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
44	43	0.25	1.7017	600.8897	260	65	0.65	12.0159	3223.8894
57	56	0.55	2.0544	614.0245	260	64	0.65	12.0324	3234.2510
60	59	0.55	2.1288	618.2430	260	63	0.65	12.0448	3245.7228
40	36	0.30	2.3483	627.2699	260	61	0.65	12.0558	3272.4985
59	54	0.30	2.9577	673.7504	267	63	0.65	12.1329	3333.1076
91	89	0.95	2.9873	684.2659	276	68	0.65	12.1588	3395.2317
97	95	0.95	3.0863	696.7224	279	68	0.65	12.1961	3432.1364
96	93	0.95	3.2905	699.6015	280	68	0.65	12.2083	3444.4380
82	75	0.55	3.8645	709.4152	279	65	0.65	12.2535	3459.4813
91	84	0.95	4.0488	715.2210	279	64	0.65	12.2656	3470.6001
96	86	0.95	4.7126	741.3697	279	63	0.65	12.2736	3482.9102
82	70	0.55	4.8434	743.5710	280	63	0.65	12.2848	3495.3938
96	85	0.95	4.8959	748.5847	288	65	0.45	12.2982	3588.2408
96	84	0.95	5.0777	755.5158	288	64	0.45	12.3150	3597.7771
82	67	0.55	5.3678	767.5174	288	63	0.45	12.3290	3608.1371
96	82	0.95	5.4126	772.3125	288	58	0.45	12.3474	3675.4129
102	87	0.95	5.6738	784.4173	304	71	0.55	12.3889	3724.0508
82	65	0.55	5.6788	786.6372	304	70	0.55	12.4155	3730.2511
96	80	0.95	5.7418	787.9804	304	68	0.55	12.4626	3744.4504
110	94	0.95	5.9247	802.4320	304	65	0.55	12.5152	3771.2317
110	93	0.95	6.0951	807.9929	304	64	0.55	12.5269	3781.9350
96	75	0.95	6.4178	840.1570	304	63	0.55	12.5352	3793.6869
82	58	0.55	6.5882	863.6424	311	65	0.55	12.5823	3858.0693
110	86	0.95	7.0674	867.5393	311	64	0.55	12.5928	3869.0190
107	82	0.80	7.1944	869.4495	311	63	0.55	12.5999	3881.0415
110	84	0.95	7.3467	881.4777	324	71	0.65	12.6134	3962.1920
110	82	0.95	7.4028	923.0120	324	70	0.65	12.6354	3969.2646
110	80	0.95	7.6578	937.4525	324	68	0.65	12.6721	3985.7068
117	86	0.95	7.8110	955.4556	326	68	0.55	12.6746	4015.4303
105	70	0.60	8.2309	968.0403	324	65	0.65	12.7050	4017.4622
113	78	0.65	8.2870	976.4386	324	64	0.65	12.7088	4030.3743
117	82	0.95	8.3057	981.7492	326	65	0.55	12.7164	4044.1498
120	84	0.65	8.3735	1000.4928	326	64	0.55	12.7245	4055.6276
120	82	0.65	8.6229	1011.5636	326	63	0.55	12.7292	4068.2300
120	80	0.65	8.8588	1023.6668	331	68	0.65	12.7345	4071.8178
123	82	0.65	8.9589	1036.8527	333	68	0.65	12.7518	4096.4209
144	100	0.65	9.0117	1126.4429	331	65	0.65	12.7642	4104.2592
120	71	0.65	9.1372	1167.9533	331	64	0.65	12.7669	4117.4503
120	70	0.65	9.2286	1175.2921	333	65	0.65	12.7806	4129.0584
158	111	0.90	9.2531	1176.7974	333	64	0.65	12.7831	4142.3292
119	68	0.55	9.2896	1184.8956	341	70	0.65	12.7874	4177.5285
120	68	0.65	9.3970	1191.4604	345	70	0.55	12.8012	4233.3442
123	70	0.65	9.4774	1204.6744	345	68	0.55	12.8359	4249.4585
119	65	0.55	9.5287	1210.3456	345	65	0.55	12.8696	4279.8518
119	64	0.55	9.5980	1219.9356	345	64	0.55	12.8749	4291.9986
120	65	0.65	9.6090	1220.0482	345	63	0.55	12.8768	4305.3354
123	68	0.65	9.6356	1221.2469	352	70	0.65	12.8780	4312.2875
119	63	0.55	9.6615	1230.1494	351	68	0.65	12.8991	4317.8491
120	64	0.65	9.6676	1230.9135	352	68	0.65	12.9068	4330.1506
162	109	0.65	9.6841	1238.6866	352	67	0.65	12.9170	4340.5131
120	63	0.65	9.7195	1242.5390	351	65	0.65	12.9203	4352.2507
123	65	0.65	9.8319	1250.5494	352	65	0.65	12.9277	4364.6503
123	64	0.65	9.8852	1261.6864	363	71	0.65	12.9463	4439.1225
130	71	0.65	9.9377	1265.2827	363	70	0.65	12.9630	4447.0465
130	70	0.65	10.0132	1273.2331	363	68	0.65	12.9891	4465.4678
130	68	0.65	10.1494	1290.7488	363	65	0.65	13.0057	4501.0456
130	65	0.65	10.3121	1321.7189	368	68	0.65	13.0249	4526.9757
142	75	0.95	10.3989	1361.8485	368	65	0.65	13.0397	4563.0435
144	75	0.65	10.5999	1371.4073	382	71	0.90	13.0418	4687.1731
146	75	1.00	10.6147	1406.0819	382	70	0.90	13.0453	4699.6648
155	82	0.95	10.7620	1431.5148	389	75	0.65	13.0563	4730.3406
152	76	0.65	11.0231	1440.8067	389	71	0.65	13.1312	4757.0762
160	82	0.70	11.0631	1478.3410	389	70	0.65	13.1450	4765.5677
197	99	0.65	11.6124	1759.8632	389	68	0.65	13.1652	4785.3085
219	111	0.95	11.7534	1917.1502	389	65	0.65	13.1728	4823.4346
230	115	0.95	11.9104	2009.2016	399	63	0.65	13.2151	4980.9361
363	182	0.65	11.9459	3153.5321					

Table E.27: Pareto–Optimal Solutions with $D = 4$ for $\sigma = 0.8$ and $q = 0.1$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
28	26	0.55	1.4831	592.3037	205	85	0.95	13.0779	1760.1407
54	51	0.75	2.3948	622.4623	205	80	0.95	13.2532	1802.0625
53	49	0.75	2.6137	628.0454	205	79	0.95	13.2816	1811.3291
41	35	0.75	2.7582	647.8631	205	75	0.95	13.3697	1851.8843
48	41	0.75	3.1639	653.6783	226	96	0.95	13.3785	1860.9366
60	49	0.75	4.1245	710.9943	205	74	0.95	13.3849	1862.9950
63	49	0.75	4.6692	746.5570	207	75	0.95	13.4331	1869.9514
77	63	0.95	4.8990	751.7379	226	94	0.95	13.4695	1873.5628
77	60	0.95	5.4033	774.3313	226	93	0.95	13.5128	1880.1487
88	69	0.95	5.7243	816.9058	238	104	1.00	13.5450	1894.1442
77	54	0.95	6.2567	830.8322	226	90	0.95	13.6337	1901.0776
95	72	0.95	6.3053	862.9321	226	89	0.95	13.6708	1908.4725
95	70	0.95	6.5928	875.7549	226	86	0.95	13.7715	1932.0813
88	58	0.95	7.2075	905.4007	226	85	0.95	13.8014	1940.4478
95	63	0.95	7.4792	929.1747	226	83	0.95	13.8552	1958.0223
95	58	0.95	7.9819	977.4213	226	82	0.95	13.8790	1967.2461
122	82	0.95	8.2622	1044.7565	226	80	0.95	13.9199	1986.6640
122	79	0.95	8.5780	1064.7536	226	79	0.95	13.9369	1996.8799
122	75	0.95	8.9943	1089.8792	237	89	0.95	14.0192	2001.3627
122	73	0.95	9.1853	1103.8672	237	85	0.95	14.1292	2034.8944
139	85	0.95	9.4995	1176.8557	257	96	0.95	14.3772	2116.1978
140	85	0.95	9.5826	1185.3223	257	94	0.95	14.4415	2130.5559
140	84	0.95	9.6774	1190.9345	257	93	0.95	14.4716	2138.0452
122	63	0.95	9.9296	1193.2559	257	90	0.95	14.5524	2161.8449
139	80	0.95	9.9586	1206.5658	257	89	0.95	14.5761	2170.2541
139	79	0.95	10.0435	1213.1209	257	88	0.95	14.5981	2178.9350
146	86	1.00	10.0564	1218.3083	257	86	0.95	14.6367	2197.1013
140	79	0.95	10.1186	1221.8484	257	85	0.95	14.6533	2206.6155
139	75	0.95	10.3567	1241.7476	257	82	0.95	14.6907	2237.0896
146	82	1.00	10.4022	1242.1694	257	79	0.95	14.7085	2270.7882
140	75	0.95	10.4266	1250.6810	267	85	0.95	14.8859	2292.4760
139	73	0.95	10.4963	1257.6848	267	84	0.95	14.8968	2302.6859
146	79	1.00	10.6374	1262.2411	296	107	0.95	15.0343	2360.5990
144	75	0.95	10.6962	1286.4148	296	100	0.95	15.2431	2406.8641
146	75	1.00	10.9141	1292.5080	296	96	0.95	15.3366	2437.3328
146	74	1.00	10.9760	1300.7959	296	94	0.95	15.3753	2453.8698
146	73	1.00	11.0348	1309.4027	296	93	0.95	15.3925	2462.4957
146	71	1.00	11.0402	1341.2956	296	90	0.95	15.4348	2489.9070
146	70	1.00	11.0915	1350.6393	296	89	0.95	15.4457	2499.5922
144	64	0.95	11.1442	1406.5761	316	108	0.95	15.4644	2513.7287
172	93	0.95	11.1766	1408.0888	316	107	0.95	15.4931	2520.0989
146	64	1.00	11.3197	1415.8757	316	106	0.95	15.5208	2526.6376
172	90	0.95	11.3975	1424.8373	323	111	0.95	15.5268	2550.7311
172	89	0.95	11.4678	1430.7415	310	96	0.95	15.6221	2552.6121
179	96	1.00	11.4786	1434.4314	310	94	0.95	15.6532	2569.9312
172	88	0.95	11.5363	1436.8343	310	93	0.95	15.6665	2578.9651
179	94	1.00	11.6260	1444.9506	316	96	0.95	15.7367	2602.0175
172	86	0.95	11.6676	1449.5863	316	94	0.95	15.7647	2619.6718
179	93	1.00	11.6972	1450.4454	316	93	0.95	15.7765	2628.8805
179	90	1.00	11.9013	1467.8828	316	90	0.95	15.8029	2658.1439
179	86	1.00	11.9893	1515.8515	316	89	0.95	15.8084	2668.4836
179	85	1.00	12.0483	1522.4868	316	88	0.95	15.8122	2679.1575
179	84	1.00	12.1052	1529.3569	326	96	0.95	15.9184	2684.3598
179	83	1.00	12.1601	1536.4503	326	94	0.95	15.9415	2702.5729
179	82	1.00	12.2128	1543.7798	343	107	0.95	16.0277	2735.4238
179	80	1.00	12.3114	1559.2258	349	111	0.95	16.0448	2756.0532
179	79	1.00	12.3571	1567.3602	349	107	0.95	16.1352	2783.2738
179	75	1.00	12.5141	1602.9893	349	106	0.95	16.1555	2790.4953
179	74	1.00	12.5463	1612.7660	345	100	0.95	16.1889	2805.2977
179	73	1.00	12.5755	1622.9150	353	106	0.95	16.2243	2822.4781
179	72	1.00	12.6015	1633.4845	349	100	0.95	16.2544	2837.8229
189	80	0.95	12.6458	1661.4137	349	96	0.95	16.2966	2873.7472
189	79	0.95	12.6845	1669.9571	349	94	0.95	16.3097	2893.2452
191	80	0.95	12.7273	1678.9948	349	93	0.95	16.3140	2903.4155
193	82	0.95	12.7283	1679.9933	350	93	0.95	16.3287	2911.7347
191	79	0.95	12.7646	1687.6286	369	107	0.95	16.4685	2942.7737
193	80	0.95	12.8071	1696.5759	384	116	0.95	16.5311	2998.7116
189	75	0.95	12.8141	1707.3469	384	113	0.95	16.5937	3018.4873
189	74	0.95	12.8397	1717.5905	384	107	0.95	16.6957	3062.3987
205	90	0.95	12.8536	1724.4288	384	106	0.95	16.7094	3070.3444
191	75	0.95	12.8886	1725.4141	399	116	0.95	16.7632	3115.8488
205	89	0.95	12.9020	1731.1365	384	100	0.95	16.7692	3122.4183
191	74	0.95	12.9128	1735.7661	384	96	0.95	16.7854	3161.9453
205	88	0.95	12.9487	1738.0610	399	107	0.95	16.9058	3182.0236
193	75	0.95	12.9616	1743.4813	399	106	0.95	16.9170	3190.2797
205	86	0.95	13.0367	1752.5517	399	96	0.95	16.9686	3285.4588

Table E.28: Pareto–Optimal Solutions with $D = 2$ for $\sigma = 0.8$ and $q = 0.9$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
20	18	0.15	5.3964	664.5433	176	75	0.95	8.2370	4003.2917
28	27	0.15	6.6434	747.6567	176	74	0.95	8.2372	4005.6578
29	28	0.15	6.7201	765.2599	179	76	1.00	8.2394	4071.0091
36	34	0.15	7.0618	905.2943	179	75	1.00	8.2396	4073.2941
40	37	0.20	7.3080	973.4298	180	82	0.80	8.2408	4079.5380
44	43	0.25	7.4664	1041.4535	180	81	0.80	8.2428	4080.9293
59	54	0.15	7.4677	1406.0278	180	76	0.80	8.2504	4089.1666
59	52	0.15	7.4768	1409.5849	180	75	0.80	8.2516	4090.9850
59	51	0.15	7.4806	1411.5584	180	74	0.80	8.2524	4093.0505
59	50	0.15	7.4836	1413.6476	180	70	0.80	8.2542	4101.8341
59	48	0.15	7.4872	1418.2790	186	76	0.90	8.2591	4227.0327
59	47	0.15	7.4878	1420.8607	191	82	0.80	8.2635	4328.8431
61	54	0.15	7.4923	1453.6898	193	82	0.85	8.2673	4374.1135
61	52	0.15	7.5005	1457.3674	193	81	0.85	8.2688	4375.7045
61	51	0.15	7.5038	1459.4079	193	76	0.85	8.2738	4385.1256
61	48	0.15	7.5090	1466.3563	193	74	0.85	8.2746	4389.4975
63	52	0.85	7.6817	1473.3629	206	83	0.65	8.2839	4670.0944
76	47	0.75	7.7678	1797.7916	207	85	0.85	8.2882	4686.4175
82	78	0.55	7.7912	1863.6665	207	82	0.85	8.2921	4691.4067
82	76	0.55	7.8073	1864.8762	207	81	0.85	8.2933	4693.1132
82	75	0.55	7.8152	1865.5293	207	76	0.85	8.2967	4703.2176
82	74	0.55	7.8230	1866.2081	207	74	0.85	8.2970	4707.9066
82	70	0.55	7.8527	1869.1900	217	82	0.85	8.3079	4918.0447
82	67	0.55	7.8734	1871.7928	219	82	0.80	8.3108	4963.4379
82	65	0.55	7.8862	1873.7696	219	81	0.80	8.3119	4965.1306
82	62	0.55	7.9034	1877.1343	219	76	0.80	8.3152	4975.1526
82	59	0.55	7.9177	1881.2530	219	74	0.80	8.3154	4979.8781
82	58	0.55	7.9217	1882.8535	224	82	0.85	8.3181	5076.6913
82	56	0.55	7.9279	1886.3770	224	81	0.85	8.3189	5078.5379
82	54	0.55	7.9315	1890.5226	224	76	0.85	8.3207	5089.4721
82	52	0.55	7.9317	1895.5330	228	82	0.80	8.3236	5167.4148
96	74	0.95	7.9360	2184.9031	228	81	0.80	8.3245	5169.1771
96	70	0.95	7.9513	2190.6620	228	76	0.80	8.3270	5179.6110
96	68	0.95	7.9569	2194.0694	229	76	0.80	8.3283	5202.3286
96	67	0.95	7.9591	2195.9599	231	76	0.80	8.3307	5247.7637
96	65	0.95	7.9620	2200.1892	236	75	0.80	8.3367	5363.7359
97	62	0.80	7.9915	2223.8903	243	85	0.80	8.3404	5501.7137
105	75	0.60	7.9915	2387.6223	243	82	0.80	8.3428	5507.3763
105	74	0.60	7.9968	2388.4497	243	81	0.80	8.3435	5509.2545
105	70	0.60	8.0164	2392.4961	243	76	0.80	8.3447	5520.3748
105	68	0.60	8.0252	2394.7789	248	85	0.80	8.3465	5614.9177
105	67	0.60	8.0292	2396.0495	248	82	0.80	8.3486	5620.6968
105	65	0.60	8.0368	2398.7069	248	81	0.80	8.3493	5622.6137
105	58	0.60	8.0528	2411.2470	248	76	0.80	8.3502	5633.9628
113	58	0.80	8.0618	2602.9643	255	82	0.80	8.3565	5779.3455
119	74	0.55	8.0642	2708.2777	255	76	0.80	8.3574	5792.9860
119	70	0.55	8.0806	2712.6051	258	85	0.80	8.3580	5841.3257
119	68	0.55	8.0878	2715.0759	258	82	0.80	8.3597	5847.3378
119	67	0.55	8.0912	2716.3824	258	76	0.80	8.3604	5861.1387
119	65	0.55	8.0972	2719.2511	260	76	0.65	8.3606	5907.8372
119	62	0.55	8.1044	2724.1340	262	75	0.80	8.3639	5954.6560
119	58	0.55	8.1091	2732.4349	267	85	0.70	8.3644	6047.3935
129	75	0.80	8.1097	2931.8711	268	76	0.80	8.3699	6088.3146
129	74	0.80	8.1126	2933.3513	276	77	0.80	8.3773	6267.4063
129	70	0.80	8.1221	2939.6463	283	85	0.80	8.3831	6407.3456
129	67	0.80	8.1267	2945.2916	283	82	0.80	8.3839	6413.9402
129	65	0.80	8.1281	2949.7176	283	81	0.80	8.3841	6416.1277
131	67	0.80	8.1335	2990.9551	285	85	0.80	8.3849	6452.6272
131	65	0.80	8.1347	2995.4510	285	82	0.80	8.3857	6459.2684
134	70	0.85	8.1368	3054.6574	285	81	0.80	8.3859	6461.4714
134	67	0.85	8.1394	3061.0190	289	82	0.75	8.3884	6550.4506
144	81	0.85	8.1455	3264.7725	302	91	0.80	8.3962	6826.0297
144	75	0.85	8.1605	3273.4292	302	82	0.80	8.3997	6844.5581
144	68	0.85	8.1690	3287.0572	316	93	0.80	8.4064	7138.8415
156	85	0.85	8.1722	3531.7909	328	75	0.80	8.4135	7454.6838
156	82	0.85	8.1803	3535.5509	333	85	0.80	8.4221	7539.3854
156	81	0.85	8.1829	3536.8369	338	85	0.80	8.4254	7652.5894
156	76	0.85	8.1933	3544.4538	344	93	0.80	8.4272	7771.3970
156	74	0.85	8.1964	3547.9876	344	85	0.80	8.4292	7788.4342
156	70	0.85	8.2000	3556.1701	354	93	0.80	8.4338	7997.3097
158	75	0.65	8.2003	3591.5855	358	93	0.80	8.4363	8087.6748
158	74	0.65	8.2026	3592.9503	358	91	0.80	8.4370	8091.7835
158	71	0.65	8.2085	3597.6463	358	85	0.80	8.4375	8105.4054
158	70	0.65	8.2101	3599.3805	368	93	0.85	8.4438	8312.2292
158	67	0.65	8.2136	3605.1435	368	91	0.85	8.4443	8316.4121
158	65	0.65	8.2149	3609.4091	378	104	0.90	8.4459	8515.8260
176	85	0.95	8.2246	3985.1157	378	93	0.90	8.4504	8537.2505
176	82	0.95	8.2299	3989.8655	378	91	0.90	8.4504	8541.9813
176	81	0.95	8.2314	3991.4920	397	104	0.85	8.4556	8945.7620
176	76	0.95	8.2365	4001.1297	397	93	0.85	8.4598	8967.2690

Table E.29: Pareto–Optimal Solutions with $D = 4$ for $\sigma = 0.8$ and $q = 0.9$

F	M	p	TH	$Delay$	F	M	p	TH	$Delay$
11	11	0.15	3.2561	597.1464	186	82	1.00	8.1882	4156.9028
15	13	0.15	3.8510	696.4835	186	81	1.00	8.1895	4158.5300
16	14	0.15	4.0429	706.2196	186	76	1.00	8.1927	4169.0692
22	20	0.15	4.8789	797.8329	208	98	1.00	8.2020	4625.0178
28	26	0.55	6.0574	810.6548	208	95	1.00	8.2086	4628.4422
40	39	0.75	7.0602	982.6642	208	89	1.00	8.2198	4636.4889
54	44	1.00	7.3929	1287.8192	208	85	1.00	8.2253	4642.9824
63	57	0.85	7.5976	1445.4107	208	82	1.00	8.2281	4648.5795
67	65	1.00	7.6505	1514.6468	208	81	1.00	8.2290	4650.3991
69	68	1.00	7.6585	1555.3764	208	76	1.00	8.2296	4662.1850
71	62	0.80	7.6627	1620.3219	209	76	1.00	8.2311	4684.5993
72	69	0.95	7.6791	1623.4435	216	85	1.00	8.2383	4821.5586
74	60	0.80	7.6959	1693.2097	216	82	1.00	8.2406	4827.3710
82	79	1.00	7.7306	1835.1225	216	81	1.00	8.2413	4829.2606
82	75	1.00	7.7562	1839.0999	223	82	0.95	8.2419	4989.4742
83	64	0.80	7.7656	1889.7732	223	81	0.95	8.2422	4991.7341
85	63	1.00	7.8276	1925.8063	229	89	0.95	8.2447	5110.2484
92	76	1.00	7.8361	2062.1203	229	85	0.95	8.2486	5117.3761
92	75	1.00	7.8410	2063.3804	229	82	0.95	8.2501	5123.7202
92	73	1.00	7.8501	2065.9404	229	81	0.95	8.2503	5126.0409
92	72	1.00	7.8542	2067.3002	247	104	1.00	8.2551	5485.2201
92	70	1.00	7.8617	2070.4072	247	98	1.00	8.2660	5492.2086
92	69	1.00	7.8649	2072.1080	247	96	1.00	8.2691	5494.9644
92	68	1.00	7.8677	2073.8352	247	95	1.00	8.2706	5496.2751
92	65	1.00	7.8738	2079.8135	247	91	1.00	8.2758	5502.3463
92	62	1.00	7.8752	2087.0339	247	89	1.00	8.2778	5505.8305
96	71	0.95	7.8767	2161.2758	247	85	1.00	8.2806	5513.5416
96	70	0.95	7.8798	2162.8727	247	82	1.00	8.2814	5520.1882
96	68	0.95	7.8852	2166.5408	247	81	1.00	8.2816	5522.3489
96	62	0.95	7.8906	2180.0784	251	89	1.00	8.2827	5594.9938
108	85	1.00	7.9009	2410.7793	251	87	1.00	8.2842	5598.7834
108	82	1.00	7.9156	2413.6855	251	85	1.00	8.2853	5602.8297
108	81	1.00	7.9202	2414.6303	251	82	1.00	8.2860	5609.5840
108	76	1.00	7.9406	2420.7499	251	81	1.00	8.2861	5611.7797
108	75	1.00	7.9441	2422.2292	255	85	1.00	8.2898	5692.1178
108	73	1.00	7.9503	2425.2344	278	104	1.00	8.2956	6173.6485
108	72	1.00	7.9531	2426.8306	278	98	1.00	8.3041	6181.5142
108	70	1.00	7.9576	2430.4780	278	95	1.00	8.3075	6186.0910
108	68	1.00	7.9608	2434.5022	278	91	1.00	8.3111	6192.9242
108	67	1.00	7.9618	2436.6122	278	89	1.00	8.3122	6196.8457
116	66	1.00	7.9981	2619.6635	278	87	1.00	8.3131	6201.0430
119	70	0.95	7.9997	2681.0609	278	85	1.00	8.3135	6205.5245
119	68	0.95	8.0013	2685.6079	286	76	1.00	8.3148	6410.5043
119	67	0.95	8.0016	2687.8877	292	98	1.00	8.3186	6492.8135
132	85	0.95	8.0149	2949.7539	292	96	1.00	8.3206	6496.0713
132	81	0.95	8.0281	2954.7484	292	95	1.00	8.3216	6497.6207
132	76	0.95	8.0410	2962.0994	292	93	1.00	8.3233	6500.9784
132	75	0.95	8.0429	2963.8298	292	91	1.00	8.3245	6504.7980
132	70	0.95	8.0489	2973.9499	292	90	1.00	8.3251	6506.7438
132	68	0.95	8.0490	2978.9936	292	89	1.00	8.3254	6508.9170
146	93	1.00	8.0495	3250.4892	292	87	1.00	8.3259	6513.3257
146	90	1.00	8.0603	3253.3719	292	85	1.00	8.3260	6518.0330
146	85	1.00	8.0765	3259.0165	315	104	1.00	8.3334	6995.3211
146	82	1.00	8.0848	3262.9452	315	98	1.00	8.3397	7004.2337
146	81	1.00	8.0873	3264.2224	315	95	1.00	8.3420	7009.4196
146	76	1.00	8.0971	3272.4952	315	87	1.00	8.3446	7026.3617
146	75	1.00	8.0984	3274.4950	321	85	1.00	8.3486	7165.3718
146	72	1.00	8.1010	3280.7155	327	87	1.00	8.3533	7294.0326
146	70	1.00	8.1011	3285.6462	335	95	1.00	8.3575	7454.4622
162	75	0.95	8.1327	3637.4275	335	89	1.00	8.3590	7467.4220
164	82	1.00	8.1376	3665.2262	340	98	1.00	8.3594	7560.1253
164	70	1.00	8.1459	3690.7259	340	95	1.00	8.3611	7565.7228
167	73	1.00	8.1539	3750.1310	340	91	1.00	8.3623	7574.0799
171	70	1.00	8.1608	3848.2569	340	89	1.00	8.3624	7578.8760
173	81	1.00	8.1615	3867.8800	352	95	1.00	8.3693	7832.7483
173	76	1.00	8.1665	3877.6827	375	108	1.00	8.3756	8321.7776
173	75	1.00	8.1668	3880.0523	378	82	1.00	8.3800	8447.8993
175	79	1.00	8.1684	3916.4199	383	89	1.00	8.3876	8537.3809
178	76	1.00	8.1770	3989.7544	391	98	1.00	8.3917	8694.1440
179	76	1.00	8.1791	4012.1688	391	95	1.00	8.3924	8700.5812
179	75	1.00	8.1792	4014.6206	398	104	1.00	8.3928	8838.5328
180	75	1.00	8.1812	4037.0487	398	95	1.00	8.3961	8856.3461
186	85	1.00	8.1838	4151.8977					

Bibliography

[ABM96] M. Azizoglu, R. A. Barry, and A. Mokhtar. Impact of Tuning Delay on the Performance of Bandwidth–Limited Optical Broadcast Networks with Uniform Traffic. *IEEE Journal on Selected Areas in Communications*, 14(5):935–944, June 1996.

[ACG+88] E. Arthurs, J. M. Cooper, M. S. Goodman, H. Kobrinski, M. Tur, and M. P. Vecchi. Multiwavelength Optical Crossconnect For Parallel–Processing Computers. *Electronics Letters*, 24(2):119–120, 1988.

[AGKV88] E. Arthurs, M. S. Goodman, H. Kobrinski, and M. P. Vecchi. HYPASS: An Optoelectronic Hybrid Packet Switching System. *IEEE Journal on Selected Areas in Communications*, 6(9):1500–1510, Dec. 1988.

[AIX] http://pma.nlanr.net/PMA/Sites/AIX.html.

[Ali] Optical Service Network^{TM} Solutions: Bringing Service Intelligence To Next Generation Metropolitan Area Networks. http://www.lightreading.com/.

[AM95] M. Azizoglu and A. Mokhtar. Multiaccess in All–Optical Networks with Wavelength and Code Concurrency. *Fiber and Integrated Optics*, 14:37–51, 1995.

[AR01] D. Awduche and Y. Rekhter. Multiprotocol Lambda Switching: Combining MPLS Traffic Engineering Control with Optical Crossconnects. *IEEE Communications Magazine*, pages 111–116, March 2001.

[AS91] A. S. Acampora and S. I. A. Shah. Multihop lightwave networks: A comparison of store–and–forward and hot–potato routing. In *Proc., IEEE INFOCOM*, pages 10–19, Bal Harbour, FL, April 1991.

[BBC+98] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. *An Architecture for Differentiated Services*, Dec. 1998. RFC 2475.

[BBLN02] A. Bianco, M. Bonsignori, E. Leonardi, and F. Neri. Variable–Size Packets in Slotted WDM Ring Networks. In *Optical Network Design and Modelling*, Feb. 2002.

[BCC98] A. Borella, G. Cancellieri, and F. Chiaraluce. *Wavelength Division Multiple Access Optical Networks*. Artech House, 1998.

[BCS94] R. Braden, D. Clark, and S. Shenker. *Integrated Services in the Internet Architecture: An Overview.*, June 1994. RFC 1633.

[BD91] K. Bogineni and P. W. Dowd. Collisionless Media Access Protocols For High Speed Communication In Optically Interconnected Parallel Computers. In *Proc., SPIE High–Speed Fiber Networks and Channels*, volume 1577, pages 276–287, Sept. 1991.

[BD92] K. Bogineni and P. W. Dowd. A Collisionless Multiple Access Protocol for a Wavelength Division Multiplexed Star–Coupled Configuration: Architecture and Performance Analysis. *IEEE/OSA Journal of Lightwave Technology*, 10(11):1688–1699, Nov. 1992.

[Ben99] K. Bengi. Performance of Single–Hop WDM LANs Supporting Real–Time Services. *Photonic Network Communications*, 1(4):287–301, 1999.

[BGL+01] A. Bianco, G. Galante, E. Leonardi, F. Neri, and A. Nucci. Scheduling Algorithms for Multicast Traffic in TDM/WDM Networks with Arbitrary Tuning Latencies. In *Proc., IEEE GLOBECOM*, pages 1551–1556, San Antonio, TX, Nov. 2001.

[BH93] R. A. Barry and P. A. Humblet. Latin Routers, Design and Implementation. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):891–899, 1993.

[BJM99] M. S. Borella, J. P. Jue, and B. Mukherjee. Simple Scheduling Algorithms for Use with a Waveguide Grating Multiplexer Based Local Optical Network. *Photonic Network Communications*, 1(1), 1999.

[BLMN00a] A. Bianco, E. Leonardi, M. Mellia, and F. Neri. Network Controller Design for SONATA — A Large–Scale All–Optical Passive Network. *IEEE Journal on Selected Areas in Communications*, 18(10):2017–2028, Oct. 2000.

[BLMN00b] A. Bianco, E. Leonardi, M. Mellia, and F. Neri. Network Controller Design for SONATA, A Large Scale All–Optical WDM Network. In *Proc., Optical Network Workshop*, University of Dallas, TX, Jan./Feb. 2000.

[BM93a] S. Banerjee and B. Mukherjee. FairNet: A WDM–based Multiple Channel Lightwave Network with Adaptive and Fair Scheduling Policy. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):1104–1112, May/June 1993.

[BM93b] M. Borella and B. Mukherjee. Multicasting in a WDM Local Lightwave Network. In *Proc., ITC sponsored seminar on 'Teletraffic Analysis Methods for Current and Future Telecom Networks'*, pages 71–77, Bangalore, India, Nov. 1993.

[BM95a] M. S. Borella and B. Mukherjee. A Reservation–Based Multicasting Protocol for WDM Local Lightwave Networks. In *Proc., IEEE International Conference on Communications (ICC)*, pages 1277–1281, 1995.

[BM95b] M. S. Borella and B. Mukherjee. Limits of Multicasting in a Packet–Switched WDM Single–Hop Local Lightwave Network. *Journal of High Speed Networks*, 4:155–167, 1995.

[BM96] M. S. Borella and B. Mukherjee. Efficient Scheduling of Nonuniform Packet Traffic in a WDM/TDM Local Lightwave Network with Arbitrary Transceiver Tuning Latencies. *IEEE Journal on Selected Areas in Communications*, 14(5):923–996, June 1996.

[BNS99] S. Binetti, G. Notaro, and R. Sabella. Transmission Features and Limitations of Large–Scale Switchless All–Optical Networks. *IEEE Photonics Technology Letters*, 11(7):913–915, 1999.

[BO00] G. Barish and K. Obraczka. World Wide Web Caching: Trends and Techniques. *IEEE Communications Magazine*, pages 178–185, May 2000.

[BR99a] I. Baldine and G. N. Rouskas. Dynamic Reconfiguration Policies for WDM Networks. In *Proc., IEEE INFOCOM*, pages 313–320, New York, NY, 1999.

[BR99b] R. Breyer and S. Riley. *Switched, Fast, And Gigabit Ethernet*. New Riders, third edition, 1999.

[BR01] I. Baldine and G. N. Rouskas. Traffic Adaptive WDM Networks: A Study of Reconfiguration Issues. *IEEE/OSA Journal of Lightwave Technology*, 19(4):433–455, April 2001.

[BSD93] K. Bogineni, K. M. Sivalingam, and P. W. Dowd. Low–Complexity Multiple Access Protocols for Wavelength–Division Multiplexed Photonic Networks. *IEEE Journal on Selected Areas in Communications*, 11(4):590–604, May 1993.

[BSdMN01] A. Bertaina, P. Sillard, L.-A. de Montmorillon, and P. Nouchi. Line fibers for new transmission windows. In *Proc., IEEE Lasers and Electro–Optics Society (LEOS)*, pages 473–474, Nov. 2001.

[BSS97] M. Bandai, S. Shiokawa, and I. Sasase. Performance Analysis of Multicasting Protocol in WDM–Based Single–Hop Lightwave Networks. In *Proc., IEEE GLOBECOM*, Phoenix, AZ, Nov. 1997.

[CAM+96] T.-K. Chiang, S. K. Agrawal, D. T. Mayweather, D. Sadot, C. F. Barry, M. Hickey, and L. G. Kazovsky. Implementation of STARNET: A WDM Computer Communications Network. *IEEE Journal on Selected Areas in Communications*, 14(5):824–839, June 1996.

[CDR90] M. S. Chen, N. R. Dono, and R. Ramaswami. A Media–Access Protocol for Packet–Switched Multiwavelength Optical Metropolitan Area Networks. *IEEE Journal on Selected Areas in Communications*, 8(6):1048–1057, Aug. 1990.

[CEFS99] I. Chlamtac, V. Elek, A. Fumagalli, and C. Szabó. Scalable WDM Access Network Architecture Based on Photonic Slot Routing. *IEEE/ACM Transactions on Networking*, 7(1):1–9, Feb. 1999.

[CF91] I. Chlamtac and A. Fumagalli. QUADRO–Stars: High Performance Optical WDM Star Networks. In *Proc., IEEE GLOBECOM*, pages 1224–1229, 1991.

[CF94a] I. Chlamtac and A. Fumagalli. Quadro–Star: A High Performance Optical WDM Star Network. *IEEE Transactions on Communications*, 42(8):2582–2591, 1994.

[CF94b] I. Chlamtac and A. Fumagalli. Quadro: A solution to packet switching in optical transmission networks. *Computer Networks and ISDN Systems*, 26:945–963, 1994.

[CG88a] I. Chlamtac and A. Ganz. A Multibus Train Communication (AMTRAC) Architecture for High–Speed Fiber Optic Networks. *IEEE Journal on Selected Areas in Communications*, 6(6):903–912, July 1988.

[CG88b] I. Chlamtac and A. Ganz. Channel Allocation Protocols in Frequency–Time Controlled High Speed Networks. *IEEE Transactions on Communications*, 36(4):430–440, 1988.

[CG88c] I. Chlamtac and A. Ganz. Channel Allocation Protocols in Frequency–Time Controlled High Speed Networks. *IEEE Transactions on Communications*, 36(4):430–440, April 1988.

[CG90] I. Chlamtac and A. Ganz. Toward Alternative High–Speed Network Concepts: The SWIFT Architecture. *IEEE Transactions on Communications*, 38(4):431–439, 1990.

[CGMJ+93] D. Chiaroni, P. Gavignet-Morin, J. B. Jacob, J. M. Gabriagues, J. Jacquet, D. De Bouard, G. Da Loura, and C. Chauzat. Feasibility demonstration of a 2.5 Gbit/s 16×16 ATM photonic switching matrix. In *Proc., OFC, paper WD2*, 1993.

[CHA+01] M. C. Chia, D. K. Hunter, I. Andonovic, P. Ball, I. Wright, S. P. Ferguson, K. M. Guild, and M. J. O'Mahony. Packet Loss and Delay Performance of Feedback and Feed–Forward Arrayed–Waveguide Gratings–Based Optical Packet Switches With WDM Inputs–Outputs. *IEEE/OSA Journal of Lightwave Technology*, 19(9):1241–1254, 2001.

[CHNS00] N. P. Caponio, A. M. Hill, F. Neri, and R. Sabella. Single–Layer Optical Platform Based on WDM/TDM Multiple Access for Large–Scale "Switchless" Networks. *European Transactions on Telecommunications*, 11(1):73–82, Jan./Feb. 2000.

[CJZ+00] D. Chiaroni, A. Jourdan, T. Zami, N. Le Sauze, E. Dotaro, J. Y. Emery, and L. Tancevski. Towards 10Tbit/s optical packet routers for the backbone. In *Proc., ECOC*, pages 73–74, Sept. 2000.

[CO90] I. Cidon and Y. Ofek. Metaring – A Full–duplex Ring with Fairness and Spatial Reuse. In *Proc., IEEE INFOCOM*, pages 969–978, San Francisco, CA, June 1990.

[CO93] I. Cidon and Y. Ofek. MetaRing – A Full–Duplex Ring with Fairness and Spatial Reuse. *IEEE Transactions on Communications*, 41(1):110–120, Jan. 1993.

[CY91] M. Chen and T.-S. Yum. A Conflict–Free Protocol for Optical WDMA Networks. In *Proc., IEEE GLOBECOM*, pages 1276–1281, Dec. 1991.

[CY94] M. Chen and T. P. Yum. Buffer Sharing in Conflict–Free WDMA Networks. *IEICE Transactions on Communications*, E77–B(9):1144–1151, Sept. 1994.

[CZ95] R. Chipalkatti and Z. Zhang. A Hybrid Dynamic Reservation Protocol for an Optical Star Network. In *Proc., IEEE GLOBECOM*, pages 998–1006, Singapore, 1995.

[CZA92] R. Chipalkatti, Z. Zhang, and A. S. Acampora. High Speed Communication Protocols for Optical Star Coupler Using WDM. In *Proc., IEEE INFOCOM*, pages 2124–2133, Florence, Italy, 1992.

[CZA93] R. Chipalkatti, Z. Zhang, and A. S. Acampora. Protocols for Optical Star–Coupler Network Using WDM: Performance and Complexity Study. *IEEE Journal on Selected Areas in Communications*, 11(4):579–589, May 1993.

[DB92] P. W. Dowd and K. Bogineni. Simulation Analysis of a Collisionless Multiple Access Protocol for a Wavelength Division Multiplexed Star–Coupled Configuraion. In *Proc., 25th Annual Simulation Symposium*, pages 129–138, Orlando, FL, April 1992.

[dBDH+02] J. H. den Besten, M. P. Dessens, C. G. P. Herben, X. J. M. Leijtens, F. H. Groen, M. R. Leys, and M. K. Smit. Low–Loss, Compact, and Polarization Independent PHASAR Demultiplexer Fabricated by Using a Double–Etch Process. *IEEE Photonics Technology Letters*, 14(1):62–64, 2002.

[DDR98] B. Davie, P. Doolan, and Y. Rekhter. *Switching in IP Networks*. Morgan Kaufmann, 1998.

[Deb01] K. Deb. *Multiobjective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, 2001.

[DEK91] C. Dragone, C. A. Edwards, and R. C. Kistler. Integrated Optics $N \times N$ Multiplexer on Silicon. *IEEE Photonics Technology Letters*, 3(10):896–899, 1991.

[DFS+01] B. N. Desai, N. J. Frigo, A. Smiljanić, K. C. Reichmann, P. P. Iannone, and R. S. Roman. An optical implementation of a packet–based (Ethernet) MAC in a WDM passive optical network overlay. In *Proc., OFC, paper WN5*, 2001.

[DGL+90] N. R. Dono, P. E. Green, K. Liu, R. Ramaswami, and F. F.-K. Tong. A Wavelength Division Multiple Access Network for Computer Communication. *IEEE Journal on Selected Areas in Communications*, 8(6):983–994, Aug. 1990.

[DHKK89] C. Dragone, C. H. Henry, I. P. Kaminow, and R. C. Kistler. Efficient Multichannel Integrated Optics Star Coupler on Silicon. *IEEE Photonics Technology Letters*, 1(8):241–243, 1989.

[Dow91] P. W. Dowd. Random Access Protocols for High–Speed Interprocessor Communication Based on an Optical Passive Star Topology. *IEEE/OSA Journal of Lightwave Technology*, 9(6):799–808, June 1991.

[DPC+96] P. Dowd, J. Perreault, J. Chu, D. C. Hoffmeister, R. Minnich, D. Burns, F. Hady, Y.-J. Chen, M. Dagenais, and D. Stone. LIGHTNING Network and Systems Architecture. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1371–1387, June 1996.

[DR00] B. Davie and Y. Rekhter. *MPLS — Technology and Applications*. Morgan Kaufmann, 2000.

[Dra89] C. Dragone. Efficient $N \times N$ Star Couplers Using Fourier Optics. *IEEE/OSA Journal of Lightwave Technology*, 7(3):479–489, March 1989.

[DS94] P. W. Dowd and K. M. Sivalingam. A Multi–Level WDM Access Protocol For An Optically Interconnected Parallel Computer. In *Proc., IEEE INFOCOM*, pages 400–408, Toronto, Ontario, Canada, 1994.

[DS99] A. Dasylva and R. Srikant. Optimal WDM Schedules for Optical Star Networks. *IEEE/ACM Transactions on Networking*, 7(3):446–456, June 1999.

[DY01] R. Doverspike and J. Yates. Challenges for MPLS in Optical Network Restoration. *IEEE Communications Magazine*, pages 89–96, Feb. 2001.

[EBS02] T. S. El-Bawab and J.-D. Shin. Optical Packet Switching in Core Networks: Between Vision and Reality. *IEEE Communications Magazine*, pages 60–65, Sept. 2002.

[EFW01] V. Elek, A. Fumagalli, and G. Wedzinga. Photonic Slot Routing: A Cost–Effective Approach to Designing All–Optical Access and Metro Networks. *IEEE Communications Magazine*, pages 164–172, Nov. 2001.

[EHY87] K. Y. Eng, M. G. Hluchyi, and Y.-S. Yeh. A Knockout Switch for Variable–Length Packets. *IEEE Journal on Selected Areas in Communications*, SAC–5(9):1426–1435, Dec. 1987.

[Eng88] K. Y. Eng. A Photonic Knockout Switch for High–Speed Packet Networks. *IEEE Journal on Selected Areas in Communications*, 6(7):1107–1116, Aug. 1988.

[FHA98] A. Franzen, D. K. Hunter, and I. Andonovic. A Low Loss Optical Packet Synchronisation Architecture. In *Proc., SPIE*, volume 3531, pages 390–395, Nov. 1998.

[FIM+94] N. J. Frigo, P. P. Iannone, P. D. Magill, T. E. Darcie, M. M. Downs, B. N. Desai, U. Koren, T. L. Koch, C. Dragone, H. M. Presby, and G. E. Bodeep. A Wavelength–Division Multiplexed Passive Optical Network with Cost–Shared Components. *IEEE Photonics Technology Letters*, 6(11):1365–1367, 1994.

[FMR03] C. Fan, M. Maier, and M. Reisslein. The AWG||PSC Network: A Performance Enhanced Single–Hop WDM Network with Heterogeneous Protection. In *Proc., IEEE INFOCOM*, San Francisco, CA, 2003. accepted for publication.

[Fri97] N. J. Frigo. A Survey of Fiber Optics in Local Access Architectures. *Optical Fiber Telecommunications*, IIIA:461–522, 1997.

[FSA+00] Y. Fukashiro, K. Shrikhande, M. Avenarius, M. S. Rogge, I. M. White, D. Wonglumson, and L. G. Kazovsky. Fast and fine tuning of a GCSR laser using a digitally controlled driver. In *Proc., OFC, paper WM43*, volume 2, pages 338–340, 2000.

[FT83] A. Fukuda and S. Tasaka. The Equilibrium Point Analysis – A Unified Analytic Tool For Packet Broadcast Networks. In *Proc., IEEE GLOBECOM*, pages 1133–1140, Nov. 1983.

[FV00] A. Fumagalli and L. Valcarenghi. IP Restoration vs. WDM Protection: Is There an Optimal Choice? *IEEE Network*, pages 34–41, Nov./Dec. 2000.

[FWLB01] M. D. Feuer, S. L. Woodward, C. F. Lam, and M. L. Boroditsky. Upgradable metro networks using frequency–cyclic optical add/drop. In *Proc., OFC, paper WBB5*, 2001.

[GA96] D. Guo and A. S. Acampora. Scalable Multihop WDM Passive Ring with Optimal Wavelength Assignment and Adaptive Wavelength Routing. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1264–1277, 1996.

[GBC+00] J. Gripp, P. Bernasconi, C. Chan, K. L. Sherman, and M. Zirngibl. Demonstration of a 1 Tb/s Optical Packet Switch Fabric (80 * 12.5 Gb/s), Scalable To 128 Tb/s (6400 * 20 Gb/s). In *Proc., ECOC, Post Deadline Paper 2.7*, Sept. 2000.

[GCF+01] R. Gaudino, A. Carena, V. Ferrero, A. Pozzi, V. De Feo, P. Gigante, F. Neri, and P. Poggiolini. RINGO: A WDM Ring Optical Packet Network Demonstrator. In *Proc., ECOC*, volume 4, pages 620–621, Oct. 2001.

[GE95] E. L. Goldstein and L. Eskildsen. Scaling Limitations in Transparent Optical Networks Due to Low–Level Crosstalk. *IEEE Photonics Technology Letters*, 7(1):93–94, Jan. 1995.

[Ger96] O. Gerstel. On The Future of Wavelength Routing Networks. *IEEE Network*, pages 14–20, Nov./Dec. 1996.

[GG92a] A. Ganz and Y. Gao. A Time–Wavelength Assignment Algorithm For A WDM Star Network. In *Proc., IEEE INFOCOM*, pages 2144–2150, Florence, Italy, 1992.

[GG92b] A. Ganz and Y. Gao. Traffic Scheduling in Multiple WDM Star Systems. In *Proc., IEEE International Conference on Communications (ICC)*, pages 1468–1472, 1992.

[GG94] A. Ganz and Y. Gao. Time–Wavelength Assignment Algorithms for High Performance WDM Star Based Systems. *IEEE Transactions on Communications*, 42(2/3/4):1827–1836, 1994.

[GGHR98] L. Giehmann, A. Gladisch, N. Hanik, and J. Rudolph. The application of code division multiple access for transport overhead information in transparent optical networks. In *Proc., OFC, paper WM42*, 1998.

[GGR99] L. Giehmann, A. Gladisch, and J. Rudolph. Field Trial of OAM–Signal Transport Capabilities with a 10Mchip/s LED–Direct Sequence Spread Spectrum System suited for OAM–Signal–Transport in Transparent Optical WDM–Networks. In *Proc., OFC, paper TuR3*, 1999.

[GK91] A. Ganz and Z. Koren. WDM Passive Star – Protocols and Performance Analysis. In *Proc., IEEE INFOCOM*, pages 991–1000, Bal Harbor, FL, 1991.

[GKL86] M. S. Goodman, H. Kobrinski, and K. W. Loh. Application of Wavelength Division Multiplexing to Communication Network Architectures. In *Proc., IEEE International Conference on Communications (ICC)*, pages 931–933, 1986.

[GKV+90] M. S. Goodman, H. Kobrinski, M. P. Vecchi, R. M. Bulley, and J. L. Gimlett. The LAMBDANET Multiwavelength Network: Architecture, Applications, and Demonstrations. *IEEE Journal on Selected Areas in Communications*, 8(6):995–1004, Aug. 1990.

[GKW94] B. Glance, I. P. Kaminow, and R. W. Wilson. Applications of the Integrated Waveguide Grating Router. *IEEE/OSA Journal of Lightwave Technology*, 12(6):957–962, 1994.

[Gla92] B. S. Glance. Protection–Agains–Collision Optical Packet Network. *IEEE/OSA Journal of Lightwave Technology*, 10(9):1323–1328, Sept. 1992.

[GLG99] L. Gillner, C. P. Larsen, and M. Gustavsson. Scalability of Optical Multiwavelength Switching Networks: Crosstalk Analysis. *IEEE/OSA Journal of Lightwave Technology*, 17(1):58–67, Jan. 1999.

[Goo89] M. S. Goodman. Multiwavelength Networks and New Approaches to Packet Switching. *IEEE Communications Magazine*, 27(10):27–35, Oct. 1989.

[Gre91] P. E. Green. The Future of Fiber–Optic Computer Networks. *IEEE Computer*, pages 78–87, Sept. 1991.

[Gre92] P. E. Green. An All–Optical Computer Network: Lessons Learned. *IEEE Network*, pages 56–60, March 1992.

[Gre93] P. E. Green. *Fiber Optic Networks*. Prentice Hall, 1993.

[Gre01] P. Green. Progress in Optical Networking. *IEEE Communications Magazine*, pages 54–61, Jan. 2001.

[GWW+99] S. M. Gemelos, I. M. White, D. Wonglumson, K. Shrikande, T. Ono, and L. G. Kazovsky. WDM Metropolitan Area Network Based on CSMA/CA Packet Switching. *IEEE Photonics Technology Letters*, 11(11):1512–1514, 1999.

[GYZ94] D. Guo, Y. Yemini, and Z. Zhang. Scalable High–Speed Protocols for WDM Optical Star Networks. In *Proc., IEEE INFOCOM*, pages 1544–1551, Toronto, Ontario, Canada, 1994.

[GZS01] C. Gan, M. Zhang, and X. Sun. Expanding WDM Star Networks by using EDFAs. *Journal of Optical Communications*, 22(3):95–99, 2001.

[HA00] D. K. Hunter and I. Andonovic. Approaches to Optical Internet Packet Switching. *IEEE Communications Magazine*, 38(9):116–122, Sept. 2000.

[HBP+98] A. M. Hill, M. Brierley, R. M. Percival, R. Wyatt, D. Pitcher, K. M. I. Pati, I. Hall, and J.-P. Laude. Multiple–Star Wavelength–Router Network and Its Protection Strategy. *IEEE Journal on Selected Areas in Communications*, 16(7):1134–1145, Sept. 1998.

[HCA+96] A. M. Hill, S. Carter, J. Armitage, M. Shabeer, R. A. Harmon, and P. Rose. A Scalable and Switchless Optical Network Structure, Employing a Single 32×32 Free–Space Grating Multiplexer. *IEEE Photonics Technology Letters*, 8(4):569–571, April 1996.

[HCL01] D. Huang, T. Chiu, and Y. Lai. Arrayed Waveguide Grating DWDM Interleaver. In *Proc., OFC, paper WDD80*, 2001.

[HCT02] J. He, S.-H. G. Chan, and D. H. Tsang. Multicasting in WDM Networks. *IEEE Communications Surveys and Tutorials*, Dec. 2002.

[Her02] M. Herzog. Design und Untersuchung von optischen "Metropolitan Area Networks" (MANs) unter Berücksichtigung von neuen MAC–Protokollen (only in German available). Master's thesis, Technical University Berlin, Institute of Photonics, Sept. 2002.

[HHK+01] Y. Hida, Y. Hibino, T. Kitoh, Y. Inoue, M. Itoh, T. Shibata, A. Sugita, and A. Himeno. 400–channel 25–GHz spacing arrayed–waveguide grating covering a full range of C– and L–bands. In *Proc., OFC, paper WB2*, 2001.

[HHSW96] U. Hilbk, T. Hermes, J. Saniter, and F.-J. Westphal. High capacity WDM overlay on a passive optical network. *Electronics Letters*, 32(23):2162–2163, 1996.

[Hil00] A. M. Hill. Unconstrained Signalling in the "Switchless" WDMA/TDMA optical transport network. In *Proc., Terabit Optical Networking: Architecture, Control, and Management Issues, Part of SPIE Photonics East*, volume 4213, pages 209–219, Nov. 2000.

[HK02] A. M. Hamad and A. E. Kamal. A Survey of Multicasting Protocols for Broadcast–And–Select Single–Hop Networks. *IEEE Network*, 16(4):36–48, July/Aug. 2002.

[HKR+96] E. Hall, J. Kravitz, R. Ramaswami, M. Halvorson, S. Tenbrink, and R. Thomsen. The Rainbow–II Gigabit Optical Network. *IEEE Journal on Selected Areas in Communications*, 14(5):814–823, June 1996.

[HKS87] I. M. I. Habbab, M. Kavehrad, and C.-E. W. Sundberg. Protocols for Very High–Speed Optical Fiber Local Area Networks Using a Passive Star Topology. *IEEE/OSA Journal of Lightwave Technology*, LT–5(12):1782–1794, 1987.

[HM02] P.-H. Ho and H. T. Mouftah. A Framework of Scalable Optical Metropolitan Networks for Improving Survivability and Class of Service. *IEEE Network*, pages 29–35, July/Aug. 2002.

[HMH97] B. Hamidzadeh, M. Maode, and M. Hamdi. Message Sequencing Techniques for On–Line Scheduling in WDM Networks. In *Proc., IEEE GLOBECOM*, pages 868–872, Phoenix, AZ, 1997.

[HMH99] B. Hamidzadeh, M. Maode, and M. Hamdi. Efficient Sequencing Techniques for Variable–Length Messages in WDM Networks. *IEEE/OSA Journal of Lightwave Technology*, 17(8):1309–1319, Aug. 1999.

[HNC+99] D. K. Hunter, M. H. M. Nizam, M. C. Chia, I. Andonovic, K. M. Guild, A. Tzanakaki, M. J. O'Mahony, J. D. Bainbridge, M. F. C. Stephens, R. V. Penty, and I. H. White. WASPNET: A Wavelength Switched Packet Network. *IEEE Communications Magazine*, 37(3):120–129, March 1999.

[HRS93] P. A. Humblet, R. Ramaswami, and K. N. Sivarajan. An Efficient Communication Protocol for High–Speed Packet–Switched Multichannel Networks. *IEEE Journal on Selected Areas in Communications*, 11(4):568–578, May 1993.

[HS01] J. J. He and D. Simeonidou. A flow–routing approach for optical IP networks. In *Proc., OFC, paper MN2*, 2001.

[HSW83] D. E. Hubber, W. Steinlin, and P. J. Wild. SILK: An implementation of a buffer insertion ring. *IEEE Journal on Selected Areas in Communications*, SAC–1(5):1214–1223, Nov. 1983.

[HW94] N.-F. Huang and C.-S. Wu. An Efficient Transmission Scheduling Algorithm for a Wavelength–Reusable Local Lightwave Network. *IEEE/OSA Journal of Lightwave Technology*, 12(7):1278–1290, July 1994.

[HYM+00] A. Hirano, K. Yonenaga, Y. Miyamoto, H. Toba, H. Takenouchi, and H. Tsuda. 640Gbit/s (16 channel × 42.7Gbit/s) WDM L–band DSF transmission experiment using 25nm bandwidth AWG dispersion slope compensator. *Electronics Letters*, 36(19):1638–1639, 2000.

[IFD95] P. P. Iannone, N. J. Frigo, and T. E. Darcie. WDM passive–optical–network architecture with bidirectional optical spectral slicing. In *Proc., OFC, Technical Digest, paper TuK2*, 1995.

[IIH+01] Y. Inoue, M. Itoh, Y. Hashizume, Y. Hibino, A. Sugita, and A. Himeno. Novel birefringence compensating AWG design. In *Proc., OFC, paper WB4*, 2001.

[J. 01] J. Späth. Dynamic wavelengths under decentralized management. In *Proc., TransiNet Workshop*, Berlin, Oct. 2001.

[JBM96] J. P. Jue, M. S. Borella, and B. Mukherjee. Performance Analysis of the Rainbow WDM Optical Network Prototype. *IEEE Journal on Selected Areas in Communications*, 14(5):945–951, June 1996.

[JCD+01] A. Jourdan, D. Chiaroni, E. Dotaro, G. J. Eilenberger, F. Masetti, and M. Renaud. The Perspective of Optical Packet Switching in IP–Dominant Backbone and Metropolitan Networks. *IEEE Communications Magazine*, pages 136–141, March 2001.

[JFW+00] M. Jaeger, H.-M. Foisel, F.-J. Westphal, J. Chawki, K. Ovsthus, and J.-C. Bischoff. Evaluation of Network Architectures for the Integration of IP over Optical Networks. In *Design of Reliable Communication Networks (DRCN)*, pages 261–266, April 2000.

[JM92a] F. Jia and B. Mukherjee. Bimodal Throughput, Nonmonotonic Delay, Optimal Bandwidth Dimensioning, and Analysis of Receiver Collisions In A Single–Hop WDM Local Lightwave Network. In *Proc., IEEE GLOBECOM*, pages 1896–1900, 1992.

[JM92b] F. Jia and B. Mukherjee. Performance Analysis of a Generalized Receiver Collision Avoidance (RCA) Protocol for Single–Hop WDM Local Lightwave Networks. In *Proc., SPIE High–Speed Fiber Networks and Channels II*, volume 1784, pages 229–240, Sept. 1992.

[JM93a] F. Jia and B. Mukherjee. A High–Capacity, Packet–Switched, Single–Hop Local Lightwave Network. In *Proc., IEEE GLOBECOM*, pages 1110–1114, Dec. 1993.

[JM93b] F. Jia and B. Mukherjee. The Receiver Collision Avoidancd (RCA) Protocol For A Single–Hop WDM Lightwave Network. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):1053–1065, May/June 1993.

[JM93c] F. Jia and B. Mukherjee. The Receiver Collision Avoidance (RCA) Protocol For A Single–Hop WDM Lightwave Network. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):1053–1065, May/June 1993.

[JM97] J. P. Jue and B. Mukherjee. The Advantages of Partitioning Multicast Transmissions in a Single–Hop Optical WDM Network. In *Proc., IEEE International Conference on Communications (ICC)*, 1997.

[JM00] V. K. Jain and G. De Marchis. Performance Evaluation of Optical Code Division Multiple Access Networks. *Journal of Optical Communications*, 21(3):110–115, 2000.

[JMI95] F. Jia, B. Mukherjee, and J. Iness. Scheduling Variable–Length Messages in a Single–Hop Multichannel Local Lightwave Network. *IEEE/ACM Transactions on Networking*, 3(4):477–488, Aug. 1995.

[JMIO94] F. Jia, B. Mukherjee, J. Iness, and S. Ojha. Variable–Length Message Scheduling Algorithms for A WDM–Based Local Lightwave Network. In *Proc., IEEE INFOCOM*, pages 1362–1369, Toronto, Ontario, Canada, 1994.

[JRS92] F. J. Janniello, R. Ramaswami, and D. G. Steinberg. A Prototype Circuit–Switched Multi–Wavelength Optical Metropolitan–Area Network. In *Proc., IEEE International Conference on Communications (ICC)*, pages 818–823, Chicago, IL, 1992.

[JRS93] F. J. Janniello, R. Ramaswami, and D. G. Steinberg. A Prototype Circuit–Switched Multi–Wavelength Optical Metropolitan–Area Network. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):777–782, May/June 1993.

[JU90] H. B. Jeon and C. K. Un. Contention–Based Reservation Protocol in Fibre Optic Local Area Network with Passive Star Topology. *Electronics Letters*, 26(12):780–781, 1990.

[JU92] H. B. Jeon and C. K. Un. Contention–Based Reservation Protocols In Multiwavelength Optical Networks With A Passive Star Topology. In *Proc., IEEE International Conference on Communications (ICC)*, pages 1473–1477, 1992.

[JU95] H. B. Jeon and C.-K. Un. Contention–Based Reservation Protocols in Multiwavelength Optical Networks with a Passive Star Topology. *IEEE Transactions on Communications*, 43(11):2794–2802, 1995.

[KBG+87] H. Kobrinski, R. M. Bulley, M. S. Goodman, M. P. Vecchi, and C. A. Brackett. Demonstration of High Capacity in the LAMBDANET Architecture: A Multiwavelength Optical Network. *Electronics Letters*, 23(16):824–826, 1987.

[KG93a] M. J. Karol and B. Glance. Performance of the PAC Optical Packet Network. *IEEE/OSA Journal of Lightwave Technology*, 11(8):1394–1399, Aug. 1993.

[KG93b] M. Kovačević and M. Gerla. Analysis of a T/WDMA Scheme with Subframe Tuning. In *Proc., IEEE International Conference on Communications (ICC)*, pages 1239–1244, 1993.

[KG94] M. J. Karol and B. Glance. A collision–avoidance WDM optical star network. *Computer Networks and ISDN Systems*, 26:931–943, 1994.

[KGB93] M. Kovačević, M. Gerla, and J. Bannister. Time and Wavelength Division Multiaccess with Acoustooptic Tunable Filters. *Fiber and Integrated Optics*, 12:113–132, 1993.

[Kir90] P. A. Kirkby. SYMFONET: Ultra–High–Capacity Distributed Packet Switching Network For Telecoms And Multiporcessor Computer Applications. *Electronics Letters*, 26(1):19–21, 1990.

[KIS01] T. Kitamura, M. Iizuka, and M. Sakuta. A New Partition Scheduling Algorithm by Prioritizing the Transmission of Multicast Packets with Less Destination Address Overlap in WDM Single–Hop Networks. In *Proc., IEEE GLOBECOM*, San Antonio, TX, Nov. 2001.

[KOS+00] K. Kato, A. Okada, Y. Sakai, K. Noguchi, T. Sakamoto, A. Takahara, S. Kamei, A. Kaneko, S. Suzuki, and M. Matsuoka. 10–Tbps Full–Mesh WDM Network Based On Cyclic–Frequency Arrayed–Waveguide Grating Router. In *Proc., ECOC*, volume 1, pages 105–107, Sept. 2000.

[KP93] L. G. Kazovsky and P. T. Poggiolini. STARNET: A Multi–gigabit–per–second Optical LAN Utilizing a Passive WDM Star. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):1009–1027, May/June 1993.

[KP02] G. Kramer and G. Pesavento. Ethernet Passive Optical Network (EPON): Building a Next–Generation Optical Access Network. *IEEE Communications Magazine*, pages 66–73, Feb. 2002.

[KS01] T. Kamalakis and T. Sphicopoulos. An Efficient Technique for the Design of an Arrayed–Waveguide Grating With Flat Spectral Response. *IEEE/OSA Journal of Lightwave Technology*, 19(11):1716–1725, 2001.

[KSLU95] H. S. Kim, B. C. Shin, J. H. Lee, and C. K. Un. Performance evaluation of reservation protocol with priority control for single–hop WDM networks. *Electronics Letters*, 31(17):1472–1473, 1995.

[KSW+01] L. G. Kazovsky, K. Shrikhande, I. M. White, M. Rogge, and D. Wonglumson. Optical Metropolitan Area Networks. In *Proc., OFC, paper WU1*, 2001.

[KT75] L. Kleinrock and F. A. Tobagi. Packet Switching in Radio Channels: Part I – Carrier Sense Multiple–Access Modes and Their Throughput–Delay Characteristics. *IEEE Transactions on Communications*, COM–23(12):1400–1416, 1975.

[KVG+90] H. Kobrinski, M. P. Vecchi, M. S. Goodman, E. L. Goldstein, T. E. Chapuran, J. M. Cooper, M. Tur, C. Zah, and S. G. Menocal. Fast Wavelength–Switching of Laser Transmitters and Amplifiers. *IEEE Journal on Selected Areas in Communications*, 8(6):1190–1202, Aug. 1990.

[KWSR01] L. G. Kazovsky, I. M. White, K. Shrikande, and M. S. Rogge. High Capacity Metropolitan Area Networks for the Next Generation Internet. In *Proc., 35th Asilomar Conference on Signals, Systems, and Computers*, volume 1, pages 3–7, 2001.

[KYZ+01] N. Keil, H. H. Yao, C. Zawadzki, J. Bauer, M. Bauer, C. Dreyer, and J. Schneider. Athermal polarization–independent all–polymer arrayed waveguide grating (AWG) multi/demultiplexer. In *Proc., OFC, paper PD7*, 2001.

[LA90] J.-F. P. Labourdette and A. S. Acampora. Partially Reconfigurable Multihop Lightwave Networks. In *Proc., IEEE GLOBECOM*, pages 34–40, 1990.

[LA91] J.-F. Labourdette and A. S. Acampora. Logically Rearrangeable Multihop Lightwave Networks. *IEEE Transactions on Communications*, 39(8):1223–1230, 1991.

[LA95] D. A. Levine and I. F. Akyildiz. PROTON: A Media Access Control Protocol for Optical Networks with Star Topology. *IEEE/ACM Transactions on Networking*, 3(2):158–168, April 1995.

[Lee91] H. W. Lee. Protocols for Multichannel Optical Fibre LAN Using Passive Star Topology. *Electronics Letters*, 27(17):1506–1507, 1991.

[LES00] M. Listanti, V. Eramo, and R. Sabella. Architectural and Technological Issues for Future Optical Internet Networks. *IEEE Communications Magazine*, 38(9):82–92, Sept. 2000.

[LGK95] B. Li, A. Ganz, and C. M. Krishna. A Novel Transmission Coordination Scheme for Single Hop Lightwave Networks. In *Proc., IEEE GLOBECOM*, pages 1784–1788, Singapore, 1995.

[LHA94] J.-F. Labourdette, G. W. Hart, and A. S. Acampora. Branch–Exchange Sequences for Reconfiguration of Lightwave Networks. *IEEE Transactions on Communications*, 42(10):2822–2832, 1994.

[LHH02] T.-L. Liu, C.-F. Hsu, and N.-F. Huang. Multicast QoS Traffic Scheduling with Arbitrary Tuning Latencies in Single–Hop WDM Networks. In *Proc., IEEE International Conference on Communications (ICC)*, pages 2886–2890, New York, NY, April 2002.

[LK92a] J. C. Lu and L. Kleinrock. A wavelength division multiple access protocol for high–speed local area networks with a passive star topology. *Performance Evaluation*, 16:223–239, Nov. 1992.

[LK92b] J. C. Lu and L. Kleinrock. On The Performance Of Wavelength Division Multiple Access Networks. In *Proc., IEEE International Conference on Communications (ICC)*, pages 1151–1157, June 1992.

[LK93] J. H. Laarhuis and A. M. J. Koonen. An Efficient Medium Access Control Strategy for High–Speed WDM Multiaccess Networks. *IEEE/OSA Journal of Lightwave Technology*, 11(5/6):1078–1087, May/June 1993.

[LLC00] H.-C. Lin, P.-S. Liu, and H. Chu. A Reservation–Based Multicast Scheduling Algorithm With Reservation Window For Single–Hop WDM Network. In *Proc., IEEE International Conference of Networks*, Singapore, Sept. 2000.

[LQ98] B. Li and Y. Qin. Traffic Scheduling in a Photonic Packet Switching System with QoS Guarantee. *IEEE/OSA Journal of Lightwave Technology*, 16(12):2281–2295, Dec. 1998.

[LRB00] O. A. Lavrova, G. Rossi, and D. J. Blumenthal. Rapid Tunable Transmitter with Large Number of ITU Channels Accessible in less than 5 ns. In *Proc., ECOC*, pages 169–170, Sept. 2000.

[LRI00] C. F. Lam, K. C. Reichmann, and P. P. Iannone. Cascadable Modular Transmitter and Receiver for Delivering Multiple Broadcast Services on WDM Passive Optical Newtorks. In *Proc., ECOC*, pages 109–110, Sept. 2000.

[LTWW94] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self–Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, Feb. 1994.

[LU96] J. H. Lee and G. K. Un. Dynamic Scheduling Protocol for Variable–Sized Messages in a WDM–Based Local Network. *IEEE/OSA Journal of Lightwave Technology*, 14(7):1595–1600, July 1996.

[LW00] H.-C. Lin and C.-H. Wang. Minimizing the Number of Multicast Transmissions in Single–Hop WDM Networks. In *Proc., IEEE International Conference on Communications (ICC)*, New Orleans, LA, June 2000.

[LW01a] H.-C. Lin and C.-H. Wang. A Hybrid Multicast Scheduling Algorithm for Single–Hop WDM Networks. In *Proc., IEEE INFOCOM*, pages 169–178, Anchorage, Alaska, April 2001.

[LW01b] H.-C. Lin and C.-H. Wang. A Hybrid Multicast Scheduling Algorithm for Single–Hop WDM Networks. *IEEE/OSA Journal of Lightwave Technology*, 19(11):1654–1664, Nov. 2001.

[LYWK02] G. Li, J. Yates, D. Wang, and C. Kalmanek. Control Plane Design for Reliable Optical Networks. *IEEE Communications Magazine*, pages 90–96, Feb. 2002.

[MA95] A. Mokhtar and M. Azizoglu. Hybrid Multiaccess for All–Optical LANs with Nonzero Tuning Delays. In *Proc., IEEE International Conference on Communications (ICC)*, pages 1272–1276, 1995.

[Mai01a] M. Maier. Next–Generation Metro WDM Networks. In *Proc., IEEE Computer Communications Workshop*, Oct. 2001.

[Mai01b] M. Maier. On the Design of High–Efficiency WDM Networks. Invited Talk, Feb. 2001. IEEE Communications and Signal Processing, Phoenix Chapter, Arizona State University, Goldwater Research Center.

[Mas00] B. Mason. Widely Tunable Semiconductor Lasers. In *Proc., ECOC*, volume 2, pages 157–158, Sept. 2000.

[MBL+96] M. A. Marsan, A. Bianco, E. Leonardi, M. Meo, and F. Neri. MAC Protocols and Fairness Control in WDM Multirings with Tunable Transmitters and Fixed Receivers. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1230–1244, June 1996.

[McG98] K. A. McGreer. Arrayed Waveguide Gratings for Wavelength Routing. *IEEE Communications Magazine*, pages 62–68, Dec. 1998.

[MCG+02] S. De Maesschalck, D. Colle, A. Groebbens, C. Develder, I. Lievens, P. Lagasse, M. Pickavet, P. Demeester, F. Saluta, and M. Quagliotti. Intelligent Optical Networking for Multilayer Survivability. *IEEE Communications Magazine*, pages 42–49, Jan. 2002.

[MEF] http://www.metroethernetforum.org/.

[Meh90] N. Mehravari. Performance and Protocol Improvements for Very High Speed Optical Fiber Local Area Networks Using a Passive Star Topology. *IEEE/OSA Journal of Lightwave Technology*, 8(4):520–530, 1990.

[MGLA96] A. Muir and J. J. Garcia-Luna-Aceves. Distributed Queue Packet Scheduling Algorithms for WDM–Based Networks. In *Proc., IEEE INFOCOM*, pages 938–945, 1996.

[MMPS00] G. Maier, M. Martinelli, A. Pattavina, and E. Salvadori. Design and Cost Performance of the Multistage WDM–PON Access Network. *IEEE/OSA Journal of Lightwave Technology*, 18(2):125–143, 2000.

[Mod98a] E. Modiano. A novel architecture and medium access control (MAC) protocol for WDM networks. In *Proc., OFC*, pages 90–91, 1998.

[Mod98b] E. Modiano. Unscheduled Multicasts in WDM Broadcast–and–Select Networks. In *Proc., IEEE INFOCOM*, pages 86–93, 1998.

[Mod99] E. Modiano. Random Algorithms for Scheduling Multicast Traffic in WDM Broadcast–and–Select Networks. *IEEE/ACM Transactions on Networking*, 7(3):425–434, June 1999.

[Mok00] S. Mokbel. Canada's Optical Research and Education Network: CA*net3. In *Design of Reliable Communication Networks (DRCN)*, pages 10–32, April 2000.

[MRW00] M. Maier, M. Reisslein, and A. Wolisz. High–performance switchless WDM network using multiple free spectral ranges of an arrayed–waveguide grating. In *Proc., Terabit Optical Networking: Architecture, Control, and Management Issues, Part of SPIE Photonics East*, volume 4213, pages 101–112, Nov. 2000.

[MRW02] M. Maier, M. Reisslein, and A. Wolisz. Towards Efficient Packet Switching Metro WDM Networks. *Optical Networks Magazine*, 3(6), Nov./Dec. 2002.

[MRW03] M. Maier, M. Reisslein, and A. Wolisz. A Hybrid MAC Protocol for a Metro WDM Network Using Multiple Free Spectral Ranges of an Arrayed-Waveguide Grating. *Computer Networks*, 41(4):407–433, March 2003.

[MS00] D. K. Mynbaev and L. L. Scheiner. *Fiber–Optic Communications Technology*. Pearson Education, 2000.

[MSMO97] M. Mathis, J. Semke, J. Mahdavi, and T. Ott. The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm. *Computer Communication Review*, 27(3), July 1997.

[MSRW02] M. Maier, M. Scheutzow, M. Reisslein, and A. Wolisz. Wavelength Reuse for Efficient Transport of Variable–Size Packets in a Metro WDM Network. In *Proc., IEEE INFOCOM*, pages 1432–1441, New York, NY, June 2002.

[Muk92] B. Mukherjee. WDM–Based Local Lightwave Networks Part I: Single–Hop Systems. *IEEE Network*, pages 12–27, May 1992.

[Muk97] B. Mukherjee. *Optical Communication Networks*. McGraw–Hill, 1997.

[Muk00] B. Mukherjee. WDM Optical Communication Networks: Progress and Challenges. *IEEE Journal on Selected Areas in Communications*, 18(10):1810–1824, Oct. 2000.

[MW00] M. Maier and A. Wolisz. Single–Hop WDM Network with High Spectrum Reuse Based on an Arrayed–Waveguide Grating. In *Proc., Optical Network Workshop*, University of Dallas, TX, Jan./Feb. 2000.

[MW01] M. Maier and A. Wolisz. Demonstrating the Potential of Arrayed–Waveguide Grating Based Single–Hop WDM Networks. *Optical Networks Magazine*, 2(5):75–85, Sept. 2001.

[NR01] A. Neukermans and R. Ramaswami. MEMS Technology for Optical Networking Applications. *IEEE Communications Magazine*, pages 62–69, Jan. 2001.

[O'D00] Ronan O'Dowd. Tunable and Agile Laser Transmitter Developments for Future DWDM Optical Networks: Towards Managed Wavelength Control and Switching. *Photonic Network Communications*, 2(1):97–103, 2000.

[Ogu96] K. Oguchi. New Notations Based on the Wavelength Transfer Matrix for Functional Analysis of Wavelength Circuits and New WDM Networks Using AWG–Based Star Coupler with Asymmetric Characteristics. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1255–1263, 1996.

[OHI+97] K. Okamoto, T. Hasegawa, O. Ishida, A. Himeno, and Y. Ohmori. 32×32 arrayed–waveguide grating multiplexer with uniform loss and cyclic frequency characteristics. *Electronics Letters*, 33(22):1865–1866, 1997.

[OMS95] K. Okamoto, K. Moriwaki, and S. Suzuki. Fabrication of 64×64 arrayed–waveguide grating multiplexer on silicon. *Electronics Letters*, 31(3):184–186, 1995.

[OOOK92] K. Okamota, H. Okazaki, Y. Ohmori, and K. Kato. Fabrication of Large Scale Integrated–Optic $N \times N$ Star Couplers. *IEEE Photonics Technology Letters*, 4(9):1032–1035, 1992.

[opt] http://www.optospeed.com/.

[ORP00] Z. Ortiz, G. N. Rouskas, and H. G. Perros. Maximizing Multicast Throughput in WDM Networks with Tuning Latencies Using the Virtual Receiver Concept. *European Transactions on Telecommunications*, 11(1):63–72, Jan./Feb. 2000.

[OS96] K. Okamoto and A. Sugita. Flat spectral response arrayed–waveguide grating multiplexer with parabolic waveguide horns. *Electronics Letters*, 32(18):1661–1662, 1996.

[OSHT01] M. J. O'Mahony, D. Simeonidou, D. K. Hunter, and A. Tzanakaki. The Application of Optical Packet Switching in Future Communication Networks. *IEEE Communications Magazine*, pages 128–135, March 2001.

[OSS+01] A. Okada, T. Sakamoto, Y. Sakai, K. Noguchi, and M. Matsuoka. All–optical packet routing by an out–of–band optical label and wavelength conversion in a full–mesh network based on a cyclic–frequency AWG. In *Proc., OFC, paper ThG5*, 2001.

[OTS+91] K. Okamoto, H. Takahashi, S. Suzuki, A. Sugita, and Y. Ohmori. Design and Fabrication of Integrated–Optic 8×8 Star Coupler. *Electronics Letters*, 27(9):774–775, 1991.

[OYA01] T. Ohyama, T. Yamada, and Y. Akahori. Hybrid integrated multiwavelength photoreceivers consisting of photo–diodes and an arrayed–waveguide grating. In *Proc., IEEE Lasers and Electro–Optics Society (LEOS)*, pages 835–836, Nov. 2001.

[PF95] V. Paxson and S. Floyd. Wide Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.

[PM95] G. I. Papadimitriou and D. G. Maritsas. Self–adaptive random–access protocols for WDM passive star networks. *IEE Proc.–Comput. Digit. Tech.*, 142(4):306–312, July 1995.

[PM96] G. I. Papadimitriou and D. G. Maritsas. Learning Automata–Based Receiver Conflict Avoidance Algorithms for WDM Broadcast–and–Select Star Networks. *IEEE/ACM Transactions on Networking*, 4(3):407–412, June 1996.

[PS94] G. R. Pieris and G. H. Sasaki. Scheduling Transmissions in WDM Broadcast–and–Select Networks. *IEEE/ACM Transactions on Networking*, 2(2):105–110, April 1994.

[Qia00] C. Qiao. Labeled Optical Burst Switching for IP–over–WDM Integration. *IEEE Communications Magazine*, 38(9):104–114, Sept. 2000.

[RA94] G. N. Rouskas and M. H. Ammar. Multi–Destination Communication Over Single–Hop Lightwave WDM Networks. In *Proc., IEEE INFOCOM*, pages 1520–1527, 1994.

[RA95a] G. N. Rouskas and M. H. Ammar. Analysis and Optimization of Transmission Schedules for Single–Hop WDM Networks. *IEEE/ACM Transactions on Networking*, 3(2):211–221, April 1995.

[RA95b] G. N. Rouskas and M. H. Ammar. Minimizing Delay and Packet Loss in Single–Hop Lightwave WDM Networks Using TDM Schedules. In *Proc., IEEE International Conference on Communications (ICC)*, pages 1267–1271, 1995.

[RA97] G. N. Rouskas and M. H. Ammar. Multidestination Communication Over Tunable–Receiver Single–Hop WDM Networks. *IEEE Journal on Selected Areas in Communications*, 15(3):501–511, April 1997.

[RA01] F. Ruehl and R. Anderson. Cost–Effective Metro WDM Network Architectures. In *Proc., OFC, paper WL1*, 2001.

[Ram02] R. Ramaswami. Optical Fiber Communication: From Transmission To Networking. *IEEE Communications Magazine*, pages 138–147, May 2002.

[RBF+01] A. Richter, H. Bock, W. Fischler, P. Leisching, P. M. Krummrich, A. Mayer, R. E. Neuhauser, J. P. Elbers, and C. Glingener. Germany–wide DWDM field trial: Transparent connection of a long haul link and a multiclient metro network. In *Proc., OFC, paper ML3*, 2001.

[RH95a] I. Rubin and H. Hua. An All–Optical Wavelength–Division Meshed–Ring Packet–Switching Network. In *Proc., IEEE INFOCOM*, pages 969–976, Boston, MA, April 1995.

[RH95b] I. Rubin and H. Hua. SMARTNet: An All–Optical Wavelength–Division Meshed–Ring Packet–Switching Network. In *Proc., IEEE GLOBECOM*, pages 1756–1760, 1995.

[RHZ+88] M. H. Reeve, A. R. Hunwicks, W. Zhao, S. G. Methley, L. Bickers, and S. Hornung. LED Spectral Slicing For Single–Mode Local Loop Applications. *Electronics Letters*, 24(7):389–390, 1988.

[RIB+00] K. C. Reichmann, P. P. Iannone, M. Birk, N. J. Frigo, and C. F. Lam. Simultaneous Delivery of 1280 Video Channels over a WDM Passive Optical Network. In *Proc., ECOC*, volume 3, pages 79–80, Sept. 2000.

[RPR] http://www.rpralliance.org/.

[RPS+00] B. Rajagopalan, D. Pendarakis, D. Saha, R. S. Ramamoorthy, and K. Bala. IP over Optical Networks: Architectural Aspects. *IEEE Communications Magazine*, 38(9):94–102, Sept. 2000.

[RS92] M. N. Ransom and D. R. Spears. Applications of Public Gigabit Networks. *IEEE Network*, pages 30–40, March 1992.

[RS97] G. N. Rouskas and V. Sivaraman. Packet Scheduling in Broadcast WDM Networks with Arbitrary Transceiver Tuning Latencies. *IEEE/ACM Transactions on Networking*, 5(3):359–370, June 1997.

[RS98] R. Ramaswami and K. N. Sivarajan. *Optical Networks – A Practical Perspective*. Morgan Kaufmann, 1998.

[SB98] D. Sadot and E. Boimovich. Tunable Optical Filters for Dense WDM Networks. *IEEE Communications Magazine*, pages 50–55, Dec. 1998.

[SB99] T. E. Stern and K. Bala. *Multiwavelength Optical Networks*. Addison Wesley, 1999.

[SD95] K. M. Sivalingam and P. W. Dowd. A Multilevel WDM Access Protocol for an Optically Interconnected Multiprocessor System. *IEEE/OSA Journal of Lightwave Technology*, 13(11):2152–2167, Nov. 1995.

[SD96] K. M. Sivalingam and P. W. Dowd. A Lightweight Media Access Protocol for a WDM–Based Distributed Shared Memory System. In *Proc., IEEE INFOCOM*, pages 946–953, 1996.

[SGK91] G. N. M. Sudhakar, N. D. Georganas, and M. Kavehrad. Slotted Aloha and Reservation Aloha Protocols for Very High–Speed Optical Fiber Local Area Networks Using Passive Star Topology. *IEEE/OSA Journal of Lightwave Technology*, 9(10):1411–1422, Oct. 1991.

[SGK92] G. N. M. Sudhakar, N. D. Georganas, and M. Kavehrad. A Multi–Channel Optical Star LAN and Its Application as a Broadband Switch. In *Proc., IEEE International Conference on Communications (ICC)*, pages 843–847, 1992.

[SH93] G. Semaan and P. Humblet. Timing and Dispersion in WDM Optical Star Networks. In *Proc., IEEE INFOCOM*, pages 573–577, San Francisco, CA, March/April 1993.

[SH97] S.-T. Sheu and C.-P. Huang. An Efficient Multicasting Protocol for WDM Star–Coupler Networks. In *Proc., IEEE International Symposium on Computers and Communications*, pages 579–583, Alexandria, Egypt, July 1997.

[Sim98] A. Similjanić. An Efficient Channel Access Protocol for an Optical Star Network. In *Proc., IEEE International Conference on Communications (ICC)*, pages 514–519, 1998.

[SJ01] P. Streilein and J. John. Enabling Revenue–Generating Services — The Evolution of Next–Generation Networks. *Bell Labs Technical Journal*, 6(1):3–12, Jan. 2001.

[SK91] H. Shi and M. Kavehrad. ALOHA/Slotted CSMA Protocol For A Very High–Speed Optical Fiber Local Area Network Using Passive Star Topology. In *Proc., IEEE INFOCOM*, pages 1510–1515, Bal Harbor, FL, 1991.

[SKG91] G. N. M. Sudhakar, M. Kavehrad, and N. D. Georganas. Multi–Control Channel Very High–Speed Optical Fiber Local Area Networks and Their Interconnections Using a Passive Star Topology. In *Proc., IEEE GLOBECOM*, pages 624–628, 1991.

[SLBR01] D. Stoll, P. Leisching, H. Bock, and A. Richter. Metropolitan DWDM: A Dynamically Configurable Ring for the KomNet Field Trial in Berlin. *IEEE Communications Magazine*, 39(2):106–113, Feb. 2001.

[SM97] L. Sahasrabuddhe and B. Mukherjee. Probability Distribution of the Receiver Busy Time in a Multicasting Local Lightwave Network. In *Proc., IEEE International Conference on Communications (ICC)*, 1997.

[SNY+01] Y. Sakai, K. Noguchi, R. Yoshimura, T. Sakamoto, A. Okada, and M. Matsuoka. Management System For Full–Mesh WDM AWG–Star Network. In *Proc., ECOC*, pages 264–265, Sept./Oct. 2001.

[SS00a] K. M. Sivalingam and S. Subramaniam, editors. *Optical WDM Networks – Principles and Practice*. Kluwer Academic Publishers, 2000. Chapter 9.

[SS00b] M. J. Spencer and M. A. Summerfield. WRAP: A Media Access Control Protocol for Wavelength–Routed Passive Optical Networks. In *Proc., ECOC*, volume 3, pages 85–86, Sept. 2000.

[SSW+00] K. Shrikande, A. Srivatsa, I. M. White, M. S. Rogge, D. Wonglumson, S. M. Gemelos, and L. G. Kazovsky. CSMA/CA MAC Protocols for IP–HORNET: An IP over WDM Metropolitan Area Ring Network. In *Proc., IEEE GLOBECOM*, volume 2, pages 1303–1307, 2000.

[SW96] K. M. Sivalingam and J. Wang. Media Access Protocols for WDM Networks with On–Line Scheduling. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1278–1286, June 1996.

[SWR+01] K. Shrikhande, I. M. White, M. S. Rogge, F.-T. An, A. Srivatsa, E. S. Hu, S. S.-H. Yam, and L. G. Kazovsky. Performance Demonstration of a Fast–Tunable Transmitter and Burst–Mode Packet Receiver for *HORNET*. In *Proc., OFC, paper ThG2*, 2001.

[SWW+00] K. V. Shrikhande, I. M. White, D. Wonglumsom, S. M. Gemelos, M. S. Rogge, Y. Fukashiro, M. Avenarius, and L. G. Kazovsky. HORNET: A Packet–Over–WDM Multiple Access Metropolitan Area Ring Network. *IEEE Journal on Selected Areas in Communications*, 18(10):2004–2016, Oct. 2000.

[TBB+97] D. Trouchet, A. Béguin, H. Boek, C. Prel, C. Lerminiaux, and R. O. Maschmeyer. Passband flattening of PHASAR WDM using input and output star couplers designed with two focal points. In *Proc., OFC, paper ThM7*, 1997.

[TGI01] H. Takenouchi, T. Goh, and T. Ishii. 2x40–channel dispersion–slope compensator for 40–Gbit/s WDM transmission systems covering entire C– and L–bands. In *Proc., OFC, paper TuS2*, 2001.

[THWH96] H.-Y. Tyan, C.-J. Hou, B. Wang, and C.-C. Han. On Supporting Time–Constrained Communications in WDMA–based Star–Coupled Optical Networks. In *Proc., 17th IEEE Real–Time Systems Symposium*, pages 175–184, Los Alamitos, CA, 1996.

[TIIN96] Y. Tachikawa, Y. Inoue, M. Ishii, and T. Nozawa. Arrayed–Waveguide Grating Multiplexer with Loop–Back Optical Paths and Its Applications. *IEEE/OSA Journal of Lightwave Technology*, 14(6):977–984, 1996.

[TIK+93] Y. Tachikawa, Y. Inoue, M. Kawachi, H. Takahashi, and K. Inoue. Arrayed–waveguide grating add–drop multiplexer with loop–back optical paths. *Electronics Letters*, 29(24):2133–2134, 1993.

[TK98] W.-Y. Tseng and S.-Y. Kuo. A Combinational Media Access Protocol for Multicast Traffic in Single–Hop WDM LANs. In *Proc., IEEE GLOBECOM*, Sydney, Australia, Nov. 1998.

[TK00] W.-Y. Tseng and S.-Y. Kuo. Hybrid Scheduling for Unicast and Multicast Traffic in Broadcast WDM Networks. *IEICE Transactions on Communications*, E83–B(10):2355–2363, Oct. 2000.

[TMS94] S. Tridandapani, J. S. Meditch, and A. K. Somani. The MaTPi Protocol: Masking Tuning Times Through Pipelining in WDM Optical Networks. In *Proc., IEEE INFOCOM*, volume 3, pages 1528–1535, June 1994.

[Ton98] F. Tong. Multiwavelength Receivers for WDM Systems. *IEEE Communications Magazine*, pages 42–49, Dec. 1998.

[TOT96] H. Takahashi, K. Oda, and H. Toba. Impact of Crosstalk in an Arrayed–Waveguide Multiplexer on $N \times N$ Optical Interconnection. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1097–1105, June 1996.

[TOTI95] H. Takahashi, K. Oda, H. Toba, and Y. Inoue. Transmission Characteristics of Arrayed Waveguide $N \times N$ Wavelength Multiplexer. *IEEE/OSA Journal of Lightwave Technology*, 13(3):447–455, 1995.

[TSK99] W.-Y. Tseng, C.-C. Sue, and S.-Y. Kuo. Performance Analysis for Unicast and Multicast Traffic in Broadcast–And–Select WDM Networks. In *Proc., IEEE International Symposium on Computers and Communications*, Red Sea, Egypt, July 1999.

[TTH+00] H. Tsuda, H. Takenouchi, A. Hirano, T. Kurokawa, and K. Okamoto. Performance Analysis of a Dispersion Compensator Using Arrayed–Waveguide Gratings. *IEEE/OSA Journal of Lightwave Technology*, 18(8):1139–1147, 2000.

[TYI95] K. Takada, H. Yamada, and Y. Inoue. Origin of channel crosstalk in 100 GHz–spaced silica–based arrayed–waveguide grating multiplexer. *Electronics Letters*, 31(14):1176–1177, July 1995.

[vA94] H. R. van As. Media access techniques: The evolution towards terabit/s LANs and MANs. *Computer Networks and ISDN Systems*, 26:603–656, 1994.

[VBG+88] M. P. Vecchi, R. M. Bulley, M. S. Goodman, H. Kobrinski, and C. A. Brackett. High–bit–rate measurements in the LAMBDANET multiwavelength optical star network. In *Proc., OFC, paper WO2*, 1988.

[VCR00] S. Verma, H. Chaskar, and R. Ravikanth. Optical Burst Switching: A Viable Solution for Terabit IP Backbone. *IEEE Network*, pages 48–53, Nov./Dec. 2000.

[WC92] S. S. Wagner and T. E. Chapuran. Multiwavelength Ring Networks for Switch Consolidation and Interconnection. In *Proc., IEEE International Conference on Communications (ICC)*, pages 1173–1179, 1992.

[Wes91] R. J. Westmore. SYMFONET: Interconnect Technology For Multinode Computing. *Electronics Letters*, 27(9):697–698, 1991.

[WFS+00] I. M. White, Y. Fukashiro, K. Shrikande, C. Wonglumson, M. S. Rogge, M. Avenarius, and L. G. Kazovsky. Experimental Demonstration of a Media Access Protocol for HORNET: A WDM Multiple Access Metropolitan Area Ring Network. In *Proc., OFC, paper WD3*, volume 2, pages 50–52, 2000.

[WKR+88] S. S. Wagner, H. Kobrinski, T. J. Robe, H. L. Lemberg, and J. S. Smoot. Experimental Demonstration of a Passive Optical Subscriber Loop Architecture. *Electronics Letters*, 24(6):344–346, 1988.

[WLK+88] S. S. Wagner, H. L. Lemberg, H. Kobrinski, L. S. Smoot, and T. J. Robe. A Passive Photonic Loop Architecture Employing Wavelength–Division Multiplexing. In *Proc., IEEE GLOBECOM*, pages 1569–1573, 1988.

[WLL+01] J. Y. Wei, C. Liu, K. H. Liu, J. L. Pator, and A. Roy. IP over WDM Network Traffic Engineering Demonstration and Experimentation. In *Proc., OFC, paper PD33*, 2001.

[WMW02] H. Woesner, M. Maier, and A. Wolisz. Comparison of Single–Hop and Multihop AWG–Based WDM Networks. In *Optical Network Design and Modelling*, Feb. 2002.

[Woe98] H. Woesner. Primenet – A Concept for a WDM–based Fiber Backbone. In *Optical Network Design and Modelling*, pages 98–106, April 1998.

[WRH+02] I. M. White, M. S. Rogge, Y.-L. Hsueh, K. Shrikhande, and L. G. Kazovsky. Experimental Demonstration of the HORNET Survivable Bi–directional Ring Architecture. In *Proc., OFC, paper WW1*, 2002.

[WRRW00] P. J. Williams, D. J. Robbins, F. O. Robson, and N. D. Whitbread. High Power and Wide Quasi–Continuous Tuning, Surface Ridge SG–DBR Lsers. In *Proc., ECOC*, volume 2, pages 163–164, Sept. 2000.

[WRS+00] I. M. White, M. S. Rogge, K. Shrikande, Y. Fukashiro, D. Wonglumson, F.-T. An, and L. G. Kazovsky. Experimental Demonstration of a Novel Media Access Protocol for HORNET: A Packet–Over–WDM Multiple–Access MAN Ring. *IEEE Photonics Technology Letters*, 12(9):1264–1266, 2000.

[WSR+00] I. M. White, K. Shrikhande, M. S. Rogge, S. M. Gemelos, D. Wonglumsom, G. Desa, Y. Fukashiro, and L. G. Kazovsky. Architecture and Protocols for HORNET: A Novel Packet–over–WDM Multiple–Access MAN. In *Proc., IEEE GLOBECOM*, volume 2, pages 1298–1302, Nov./Dec. 2000.

[WWG+99] D. Wonglumson, I. M. White, S. M. Gemelos, K. Shrikande, and L. G. Kazovsky. HORNET — a Packet–Switched WDM Metropolitan Area Ring Network: Optical Packet Transmission and Recovery, Queue Depth, and Packet Latency. In *Proc., IEEE Lasers and Electro–Optics Society (LEOS)*, volume 2, pages 653–654, 1999.

[XQ01] C. Xin and C. Qiao. A Comparative Study of OBS and OFS. In *Proc., OFC, paper ThG7*, 2001.

[YDA00] Y. Ye, S. Dixit, and M. Ali. On Joint Protection/Restoration in IP–Centric DWDM–Based Optical Transport Networks. *IEEE Communications Magazine*, 38(6):174–183, June 2000.

[YGK96] A. Yan, A. Ganz, and C. M. Krishna. A Distributed Adaptive Protocol Providing Real–Time Services on WDM–Based LAN's. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1245–1254, June 1996.

[YHA87] Y.-S. Yeh, M. G. Hluchyi, and A. S. Acampora. The Knockout Switch: A Simple, Modular Architecture for High–Performance Packet Switching. *IEEE Journal on Selected Areas in Communications*, SAC–5(8):1274–1283, Oct. 1987.

[YLG01] T.-W. Yeow, K. L. E. Law, and A. Goldenberg. MEMS Optical Switches. *IEEE Communications Magazine*, pages 158–163, Nov. 2001.

[YMD00] S. Yao, B. Mukherjee, and S. Dixit. Advances in Photonic Packet Switching: An Overview. *IEEE Communications Magazine*, 38(2):84–94, Feb. 2000.

[YOH97] A. Yu, M. J. O'Mahony, and A. M. Hill. Transmission limitation of all–optical network based on $N \times N$ multi/demultiplexer. *Electronics Letters*, 33(12):1068–1069, 1997.

[YQD00] M. Yoo, C. Qiao, and S. Dixit. QoS Performance of Optical Burst Switching in IP–over–WDM Networks. *IEEE Journal on Selected Areas in Communications*, 18(10):2062–2071, Oct. 2000.

[YQD01] M. Yoo, C. Qiao, and S. Dixit. Optical Burst Switching for Service Differentiation in the Next–Generation Optical Internet. *IEEE Communications Magazine*, pages 98–104, Feb. 2001.

[YYMD01] S. Yao, S. J. B. Yoo, B. Mukherjee, and S. Dixit. All–Optical Packet Switching for Metropolitan Area Networks: Opportunities and Challenges. *IEEE Communications Magazine*, pages 142–148, March 2001.

[ZDJ92] M. Zirngibl, C. Dragone, and C. H. Joyner. Demonstration of a 15×15 Arrayed Waveguide Multiplexer on InP. *IEEE Photonics Technology Letters*, 4(11):1250–1253, 1992.

[Zir98] M. Zirngibl. Multifrequency Lasers and Applications in WDM Networks. *IEEE Communications Magazine*, pages 39–41, Dec. 1998.

[Zit99] E. Zitzler. *Evolutionary Algorithms for Multiobjective Optimization: Methods and Application.* PhD thesis, Swiss Federal Institute of Technology, Zurich, 1999.

[ZJS+95] M. Zirngibl, C. H. Joyner, L. W. Stulz, C. Dragone, H. M. Presby, and I. P. Kaminow. LARNet, a Local Access Router Network. *IEEE Photonics Technology Letters*, 7(2):215–217, 1995.

[ZLG+02] B. Zhu, L. Leng, A. H. Gnauck, M. O. Pedersen, and *et al.* Transmission of 3.2 Tb/s (80 x 42.7 Gb/s) over 5200 km of UltraWave$^{\text{TM}}$ fiber with 100–km dispersion–managed spans using RZ–DPSK format. In *Proc., ECOC*, Sept. 2002.