Desirable Trends in Mobile Communication[1]

Adam Wolisz

Technische Universität Berlin, TKN
and University of California, Berkeley, BWRC

Accommodation of the predicted rapid growth of mobile data traffic is creating serious challenges. In this paper, we separate the notion of mobile data service and the data transport needed for the service. We argue that it is desirable and feasible to mitigate the traffic volume needed for achieving satisfactory services. In addition, we argue that traffic volume is not a proper measure reflecting the "transportation effort" in mobile communication. Various means to reduce this "transportation effort" and measures to quantify it are discussed in this paper. We conclude that more effort should be devoted to mitigation of the traffic growth as well as efficient traffic organization in contrast to accommodating the traffic explosion by purely blowing up the communication capacity..

I. Introduction

Communication is one of the fundamental human needs. It seems undisputed that in the future the "last hop"- communication that enables information transmission to the end-user equipment - will be dominated by wireless technologies, supporting the highly appreciated freedom of movement beyond the "leash of cable". So it is the duty and the pleasure of researchers active in this area to look how the communication requirements can be served in the best way.

In order to consider alternative architectures and solutions one requires on the one hand, a vision of the future needs on the other hand metrics making it possible to compare candidate solutions.

So far, the vision of future communication is dominated by the predictions of traffic growth. The forecast of mobile traffic published by CISCO [1] is: *"Global mobile data traffic will increase 26-fold between 2010 and 2015. Mobile data traffic will grow at a compound annual growth rate (CAGR) of 92 percent from 2010 to 2015, reaching 6.3 exabytes per month by 2015".*

In a recent prediction [2], the traffic growth for Korea is estimated even more aggressively with a CAGR of 102 % for the coming years. The UMTS Forum [3] expects the mobile traffic on licensed bands (excluding the WiFi traffic, etc.) to reach 127 exabytes in 2020 and 351 exabytes in 2025.

If the mobile traffic is really to increase by factor of ca. 130 within 10 years, a significant investment in the infrastructure will be needed. And somebody will have to pay for it.

This is not a question for the far future, bottlenecks have already been observed – e.g. in [4] the overload of the AT&T infrastructure due to wider deployment of smartphones has been reported.

Let us look somehow critically at the above predictions starting with a discussion of the basic "philosophy" justifying this trend. In fact, it is interesting to understand what do people use communication for? Looking at the historical development of human culture, we can summarize the human needs for communication as belonging to the following three areas. (Table I)

Exchanging Information / getting information…
- Messengers, Flags, Mirrors….
- E-Mail, WWW…

Maintaining inter-human relations
- Letters (of affection!)
- Phone …
- SMS, Chat, Social networks….

Entertainment
- Passive, e.g. Live Theater/Circus/ Movies/Radio/TV /Video on demand
- Active, e.g. Games – network games – virtual world

Table I: What purpose do people use communications for?

[1] This paper is an extended version of the invited talk "Do we have a vision of the Future Wireless Networks" given at the UMIC Day, RWTH Aachen in October 2011.

Interestingly enough, the reasons to use communication remain; it is only the technologies, the means of communication, which evolve.

So the question we have to ask is: Does better satisfaction of the human communication needs, as presented above, really have to imply the dramatic growth of mobile traffic volume?

The above cited predictions seem be rather convinced that it is the case: In [1], it is assumed that an average smartphone will generate 1.3 GB of traffic per month in 2015, a 16-fold increase over the 2010 average of 79 MB per month. In [3], the estimation of the future traffic is based on the statement that "the traffic per device dongle is expected to grow from 26.7 MB per day in 2010 to 503 Mbyte / day in 2020…" – a factor of almost 20!

II. What can we learn from other types of networks?

Before looking at this question from the technical point of view, we would like to start with a kind of "general observation". Communication networks are just one of the network infrastructures needed to satisfy human needs in access to a specific resource, and in fact, the youngest one. So it might be educative to have a look at the development of other types of the infrastructure like electricity networks, water provision networks etc.

It is easy to verify, that not so long ago in the '50s and '60s, it has been believed that

- the technological maturity of individual societies is measured by the (continuously growing) electricity consumption per (inhabitant * Year)
- the level of hygiene within the society is measured by the (continuously growing) water consumption.

It is pretty clear that both the above claims have not been confirmed by long term development. In contrary, we do work actively in direction of reducing the consumption of both electricity and water per person-year. Let us illustrate these trends by some data about electricity consumption derived from [5] (see Table II.)

In addition to the data presented in Table II, the prediction for 2008- 2035 features marginal increase of consumption in the highly developed OECD countries, while the less developed non-OECD countries will have a predicted increase by moderate 50%. It can be assumed that the later will be achieved by usage of modern, energy efficient technologies rather than copying the path of development of the OECD countries.

Primary Energy Usage in Quadrillion Btu (British Thermal Units)

	1981	**1991**	**2001**	**2006**
World	281	348	402	472
Europe	70	76	83	86
Germany	14.4	14.3	14.6	14.6
US	76	84	96	100

World Energy Intensity in Btu per 2000 USD of GDP using purchasing power parity

	1991	**1996**	**2001**	**2006**
World	--	10.2	9.2	8.9
Europe	--	7.8	7.1	6.5
Germany	7.8	7.3	6.7	6.4
US	11.9	11.3	9.7	8.8

Per Capita Primary Energy Consumption in million Btu per capita

	1981	**1991**	**2001**	**2006**
World	62.2	65.1	65.5	72.4
Europe	131	136	142	146
Germany	--	179	177	177
US	331	333	338	334

Table II: Energy consumption statistics (excerpt from [5].

Btw. a similar trend pertains also to usage of another basic resource delivered via distributions networks – Water. For the sake of brevity, let us just mention some data about water consumption trends in Germany in the period 1991-2007 [6]. In fact, the usage of water per person per day has decreased by 17%, and the total usage of water by the private users and small businesses (considered jointly) decreased in Germany by 21%. In parallel, the leakage of water in the water distribution system has been reduced by 23%.

Historically, the above changes in usage pattern have been enforced by really observed or envisioned **supply shortages**. Notably, there is neither feeling of less comfort in energy usage/delivery nor is there a feeling of decreased comfort/hygiene accompanying these changes.

There is one more lesson, we can learn from the electricity networks and other infrastructures. It is by far not only the total volume of demand which matters, but also the variability of demand, and the issue of when and where demand is present are critical for the infrastructure dimensioning. In last year, a lot of effort has been devoted to influencing the demand elasticity by shifting the demand in time with deployment of smart metering to create incentives for such shifts.

In this context, it is reasonable to state that we should learn from previous experience and possibly not simply improve the communication services by "just shuffling more traffic", but rather **consider in the early phases of new service development the potential of reduction of the traffic volume involved, as well as potential for shaping the traffic demand**. And, we should start doing it now, not after the bottlenecks in the infrastructure become really critical.

III. How Many Bytes per Service?

In the case of energy usage and water usage it has long ago become natural to look at the consumption per specific service. We specify, e.g. the amount of electricity and water used per cycle of dish washing.

Let us look from a similar perspective at the amount of bytes to be transported for the sake of a specific service; in fact we believe that there are numerous possibilities to reduce this value. So far only one of the approaches: development of more efficient coding seems to be widely recognized. This began with source coding (MP3 being a great example) and is continued within the network coding research.

In the following section, we will present some examples illustrating how the design of a service can contribute to dramatic reduction of amount of transported data while preserving or even improving user's experience.

III.A. Attention while watching TV

Let assume that Bob behaves like many of his friends: he comes home from work, switches on the IPTV in the living room, and starts preparing his dinner in the kitchen, while glimpsing only occasionally on the screen.

Does Bob really need the HD/3D transmission with XXX Megabits/s? In fact, a rather low resolution (easily provided by the basic layer of a modern SVC codec) would perfectly suffice in the period until his pasta will be cooked, and he will sit in front of the TV set to enjoy the food, and the newsreel.

Later, when Bob feels tired after the tough day, he switches on his IPTV set in his bedroom, selects a video, watches it a few minutes and falls to sleep afterwards. A completely superfluous traffic keeps on flowing for possibly an hour or two. This is a very common scenario in most of the homes.

This scenario could, however, be easily changed. What about using activity sensors like [7], possibly in cooperation with a camera in the room, in order to detect Bob's activities? In particular, recognition that Bob has fallen to sleep could trigger switching the TV set off thereby reducing both the data traffic and the power consumption in homes?

Of course, the sleep detector could also trigger numerous other actions like blocking the ringer on the cellular phone [8], tuning the temperate in the room etc.

III.B. Babyphone

An even more significant reduction in volume of transmitted data can be achieved by proper engineering of a **babyphone.**

The goal of this popular device is well known: Bob would like to be sure that the 3 month old baby whom Alice and he love so much is sleeping well while he is working in the garden. The classical babyphone monitors the baby with a voice channel; the technical enthusiast like Bob might prefer a more sophisticated version with an HD video being transmitted to his new smartphone.

But, does Bob really need a continuous transmission of this video?

Wouldn't it be much more convenient for him to instrument babies´ room with some microphones and cameras connected to the computing unit running activity recognition functionalities? And be sure that he will be alerted on his smartphone, whenever his baby starts being uneasy? With the possibility to use the video and possibly also get some additional information (see e.g. recent announcement [9])

III.C. Data in the cloud?

Finally, we should address one more issue: not always the data needed for the service has to be moved at all.

In fact, recently, there is a big push towards the mobile cloud [10]. And, yes, the model of “keeping my data” in the cloud and being able to access them from any device, any time is definitely a very attractive one. Unfortunately, keeping data in the cloud is frequently considered to imply that devices and appliances supporting the user will have only minimal memory, and **all** the data will continuously be shifted back and forth between the user interfaces and the cloud.

While such scenario might be reasonable in **some cases**, the long term trend in memory cost reduction makes aggressive caching of data copies within the devices and appliances quite feasible. Thus, it is reasonable to re-visit the trade-off between storage of data and transportation of data! Keeping the master data copy (or back-up copy) “in the cloud” is great, but in numerous usage scenarios mainly the locally stored copies are sufficient. Improving the speed of service delivery and reducing the amount of transmitted data to incremental changes.

III.D. The Economy of Traffic Growth

We believe that these examples have well demonstrated the potential to reduce the traffic without influencing the value of the services delivered to the user. In other words, the user does not loose anything. In fact, we argue that frequently the user even GAINS: potentially a user needs to absorb less stress factors (noise, light), less information (second example) and can get quicker service. To conclude: It is not at all true that more traffic means better service.

In fact the push toward reduction of the traffic needed per service is supported by the well known fact, expressed e.g. in [2]: “ARPU (average revenue per user) does not increase as quickly as overall data traffic”. Expressing it more directly:

- the revenue is generated according to the functionality and quality of the service (as well as competition on the market)
- the traffic transportation necessary to offer the service (with a given functionality and quality) contribute to costs. Higher traffic ⇔ Higher Service provisioning cost

Btw. a frequent question is: What about the stakes of the operators? In order to answer this question, we have to observe that operators make their gains via selling services, and the traffic transportation contributes to operational COSTS.

So there is, in fact, no good reason for operators NOT to push for reduction of the volume of data transmitted per service. Raising the awareness about the amount of bytes necessary for a given service could also be supported through the introduction of new traffic charging models. In [11] the so called "dual charging" model, in which "byte transport" is paid by the service providers has been suggested. Interestingly, similar ideas have been recently floated by major European telecommunication companies [12].

IV. Optimized data delivery

Notably, all the predictions of mobile data growth cited in section I address only the amount of bytes to be transported. This does not seem to be the proper figure of merit to drive the development of data transportation infrastructure!

The spectrum used for wireless communication is a valuable resource which usage is by no means simply proportional to the transported amount of data. The distance among the transmitter and receiver as well as the utilized technology influence significantly the area over which usage of the proper radio resources has to be kept exclusive for the considered data transmission (the general rule is to economize resources by using as short distance and as high transmission bit rate as possible).

In addition we have to remember that mobile data traffic uses to big extend also fixed transmission lines ☺

IV.A. Heterogeneous Networks (HetNets)

Following these principles since ca. 10 years the complementarity of shorter range technologies (WLAN) and long distance technologies has been recognized (see e.g. [13]), however the detailed assessment of the real resource usage by individual data flow is still an interesting issue (e.g. [14]).

Harmonized usage of the different technologies (HetNets) in supporting the actual data traffic has recently got more attention (see. e.g. [15]).

Due to the lack of efficient mobility support the usage of short–range technologies has been considered most efficient in case of static users (within home, office, coffee shops etc.), and it has been frequently argued that therefore the share of mobile services these technologies can support is limited. But recent studies [16] prove that the potential of WLAN in support of mobile services without creating additional access delays for everyday usage is as high as 65%. Interestingly, in addition, up to 55% of battery saving can be observed.

IV.B. Time-shifting the data transport

We believe that even more exciting are results of recent work investigating the potential of time-shifting the data transmission supporting mobile services. The basic question to ask is: Do mobile services REALLY always need instant connectivity? To this point a very interesting study [17] has demonstrated the INCREASE of users' satisfaction with mobile services, if the activation of the service could be prepared in advance.

In the reported study, ca. 70% of users appreciated the idea of actively preparing data for some services that might be used later. Referring to the examples used in this study, it is indeed convincing to prepare in advance an itinerary or additional information of possible relevance to the trip (to be consulted on the way, rather than generated during the trip, under time pressure, mobility conditions, with unclear connectivity quality).

And, obviously, there is high potential of pre-fetching of data not only in course of manual preparation process, but also in using multiple available context data.

In [18], the use of electronic calendar in combination with social networks as means to predict future activities has been advocated. In fact, knowing the activities planned for the next day, and the patterns of previous behavior in similar situations, it seems to be quite possible to predict numerous information items which might be needed for the next day.

E.g. knowing that Alice's habits (vegetarian, preference for a bigger lunch) and the schedule of her full day business trip, it is pretty easy to predict where and when she might take lunch. Thus, it is very reasonable to have a list of possible suggestions prepared in advance rather than generate it upon the request of a hungry Alice. Similarly, it might be very reasonable to upload in advance Alice's presentation on the server of the host of tomorrow's meeting rather than struggle with this immediately before the presentation is about to start.

Let us note, that this approach allows being pretty flexible in when the preparing searches (and associated data movements) might take place with consideration of the available connectivity and actual traffic load.

In fact, such preparations if done at home, office or hotel, might use connectivity with even smaller range and higher speed than the usually considered and frequently also overloaded WLAN. Like the Wireless USB/UWB solutions - e.g. [19].

User need prediction and pre-loading of information expected to be requested, which results in both improved network usage and higher user satisfaction has been also described in [20].

The actual research in context estimation and data mining promise a high potential for this direction. Most importantly, there is no real discomfort caused by introduction of such pre-loading of needed information.

If some information is not prepared in advance the user will get a "usual" level of service with data transmission on demand. The only overhead is created by unavoidable, preparation of excessive information and therefore excessive data transfers. But considering that this will take place during periods of "efficient" connectivity, the cost of such over provisioning with information will be really low.

Finally, let us notice that the flexibility in data transmission has also been already observed in case of user-generated content, contributing increasingly to mobile traffic. In [21], the analysis of FLICKR picture uploads led to the conclusion that 40% of content is uploaded after 10 hours or longer since it has been generated. This finding has motivated the authors to conjecture that it might be feasible to arrange for selected, very efficiently connected **Drop Zones**. In fact, they claimed that around 1000 such Drop Zones placed properly across the US territory might accommodate up to 50% of user generated content.

The idea of time delayed uploads in efficiently connected locations would become especially interesting if the expected time to pass close to one of such zones in the everyday routine were predictable. But, in fact, it **is** the case! In [22], due to use of context information including time, history, cellular network conditions and device motion the availability of Wi-Fi within one or ten hours into the future has been predicted with 95 and 90 percent accuracy, respectively!

IV.C. Device-to-Device Communication

While the potential benefits of usage of different shapes of heterogeneous networks (HetNets) for improving the efficiency of traffic delivery have already become a hot topic, the potential of direct device-to-device delivery is still by far not widely recognized.

How frequently do we drop an e-mail with a LOT of attachments to a friend sitting in the same room? How rarely do we use the ad-hoc modus of communication? And how frequently another user in our proximity might have already available the content we are interested in? Let us consider the examples of tourists in interesting places, theater or cinema viewers, commuters in a delayed train etc..

The very reluctant usage of direct device-to-device communication is a known fact. This is not the issue of pure communication technologies - WLAN is widely used, but mostly in the infrastructure mode. Highly efficient short-range technologies, e.g. UWB are available (however, indeed, not widely used yet).

The key issue is, in fact, the weak support for SERVICE discovery in the ad- hoc communication.

Today the process of establishing access to a proper service in ad-hoc mode is rather cumbersome. The available technologies support establishing of communication, but not much the management, advertising and binding of services.

A breakthrough might be the concept of **service advertisement** directly in the Physical Layer (PHY), as presented recently by the Qualcom´s FlashLinq [20]. Individual OFDM subcarriers are directly used to advertise a set of WELL DEFINED services. Such cross-layer approach increases significantly the efficiency and scalability of the service discovery process.

This innovation can be described using an analogy form everyday life. At the New York street, a free Taxi can be recognized at a glance: it is Yellow, and carries a burning "free" signal. We do NOT stop each car asking, if it is a taxi. We recognize the fire Brigade; it is RED.

Well, one could argue that the alphabet of service descriptors will become vary large, not necessarily so. We could use provider supported infrastructure to broadcast local service dictionaries – e.g. for the swimming pool? Las Vegas Casino? University? Today we also have to understand that in London taxis are black, not yellow.

Before this – or another similar in functionality – solution is widely deployed, some other approaches, based on already deployed technologies, might help.

But, in fact, the establishment of device-to-device connectivity is also complex due to the amount of wireless channels available in some of the technologies, time shifts in the activity schedule of the devices, as well as the necessity to assure agreement on different transmission parameters and access privileges.

Rather than pushing us into usage of **their** resources for communication, providers could offer the support of the discovery of potentially interesting communication partners; support for establishing direct device-to-device connectivity as a NEW SERVICE.

Providers might use their knowledge of "rough" users' location to enhance such matching; they might know (in advance) on which channel a given device is listening (or what is its listening schedule), they might distribute (in advance) some certificate of trustworthiness, etc. In any case some money could be charged for such support – without the necessity to use the provider's resources for the resulting "Bulk data transfers".

V. Conclusions

The main goal of this paper has been to trigger discussion about the trends we might like to create rather than following the "natural cycle" of development. In particular we advocate following some patterns of thinking, while developing new services, or improving the existing ones:

- We should consider decreasing the amount of information which has to be transported in order to offer the satisfactory level of a given service: it is NOT necessarily true that more bits lead to a better service
- The amount of resources used for information delivery is NOT simply proportional to the amount of transported bits. It is critical to consider how we transport these bits.
- Mobile services do not necessarily need "immediate data transportation"; services could be designed to benefit from the available time flexibility in order to use "more efficient" or just "less loaded" means of data transportation.

Besides of keeping in mind these rules while designing services, we will need also to develop proper environments supporting flexible selection of the most efficient means of communication for the sets of services, executed in parallel in distributed environment. One attempt to develop a set of concepts supporting such operational environment is presented in [24].

We are not really able to predict what types of mobile service might be developed and deployed in the future. But if we design them keeping in mind the above discussed principle, the traffic growth can be reduced. Decreasing the necessary investments in the infrastructure and helping in reducing the amount of energy consumed by the communication.

VI. Acknowledgements

The considerations presented in this paper resulted, among others, from work done within EC funded Projects COAST (FP7-ICT-248036), CONET (NoE, FP7-ICT-2007-2-224053), MEVICO (CELTIC CP7-011) and TREND (NoE, FP7-ICT-257740). Numerous motivating discussions with the members of Berkeley Wireless Research Center, the Telecommunication Networks Group at the TU Berlin, as well as project partners are gratefully acknowledged. Finally, the encouragement of the editor and his comments are highly appreciated.

References

[1] CISCO "Global Mobile Data Traffic Forecast Update, 2010- 2015", February 2011 http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html

[2] T. Jasny, "New approaches in the new chapter of telecommunications", IEEE Communications Magazine, vol. 49, no. 10, pp. 20 – 22, Oct. 2010.

[3] UMTS Forum: *"Mobile traffic forecasts 2010-2020"* Report, No. 44, Jan 2011. http://www.umts-forum.org/component/option,com_docman/task,doc_download/gid,2537/Itemid,213/

[4] J. Wortham, "Customers Angered as iPhones Overload AT&T". http://www.nytimes.com/2009/09/03/technology/companies/03att, Sept 2009.

[5] U.S. Energy Information Administration "International Energy Outlook", Sept. 19, 2011. http://www.eia.gov/forecasts/ieo/ieo_tables.cfm

[6] Umwelt Bundesamt, *"Öffentliche Wasserversorgung"* , June 2009. http://www.umweltbundesamt-daten-zur-umwelt.de/umweltdaten/public/theme.do;jsessionid=56960E1296724A6D41AF7B45FE6638A8?nodeIdent=2302

[7] The Body Media Armband http://www.bodymedia.com/Shop/Armband-Packages , featured in MIT Technology Review, July/Aug 2011.

[8] S. Phithakkitnukoon and R. Dantu, "*ContextAlert: context-aware alert mode for a mobile phone*." In: International Journal of Pervasive Computing and Communications, vol. 6, no. 3, pp. 1–23, 2010.

[9] „Smart Baby Monitor," http://www.withings.com/babyphone/

[10] O. Ingthorsson "Mobile Cloud Computing Will Soar in 2012", Jan 2012. http://www.datacenterknowledge.com/archives/2012/01/12/mobile-cloud-computing-will-soar-in-2012/

[11] A. Wolisz, "Information Access is Fine, But Who is Going to Pay?", In K. Zielinski, K. Geihs, and A. Laurentowski, editors, *New Developments in Distributed Applications and Interoperable Systems*, pp. 149-160, Kluwer, September 2001.

[12] Rik Myslewski: EU Telecoms to Apple, Google: "Pay up!", The Register, December 8th, 2010 http://www.theregister.co.uk/2010/12/08/eu_infrastructure_payment_demands/

[13] M.Stemm, R.Katz, "Vertical handoff in wireless overlay networks" Mobile Networks and Applications, vol. 3, no., 1998

[14] S. Wiethölter, M. Emmelmann , R. Andersson, and A. Wolisz ,"Performance evaluation of selection schemes for offloading traffic to IEEE 802.11 hotspots" accepted for IEEE ICC´12, Wireless Networks Symposium, Ottawa, Canada, June 2012.

[15] IEEE Wireless Communication Magazine, vol. 18, no. 3, 2011, Special Issue on HetNets.

[16] K. Lee, I. Rhee, J. Lee, S. Chong, Y. Yi "Mobile data offloading: how much can WiFi deliver?" In Proc. Con-NEXT´10 '10, pp. 26.1-26.12, Philadelphia, Nov 2010.

[17] M. Harding, O. Storz, N. Davies and A. Friday, "*Planning ahead: techniques for simplifying mobile service use*" In: HotMobile '09: Proceedings of the 10th workshop on Mobile Computing Systems and Applications, pp.13.1 – 13.6, Santa Cruz, US, Feb 2009.

[18] T. Lovett, E. O'Neill, J. Irwin and P. David, *"The Calendar as a Sensor: Analysis and Improvement Using Data Fusion with Social Networks and Location."* In: Proceedings of Ubicomp '10, pp. 3–12, Copenhagen, Denmark, Sept 2010.

[19] Samsung´s Wireless USB/Ultra wideband Solution; http://pdf.directindustry.com/pdf/samsung-semiconductor/samsung-s3c2680/34290-184447.html

[20] P. Lungaro, Z. Segall and J. Zander, "ActiveCast- a network and user aware mobile content delivery system" Second International Conference on Ubiquitous and Future Networks (ICUFN), pp. 309-313, Jeju Island, Korea (South). June 2010.

[21] I. Trestian, S. Ranjan, A. Kuzmanovic, A. Nucci "Taming user-generated content in mobile networks via Drop Zones, In Proceedings of IEEE INFOCOM 2011, pp. 2840 – 2848, Shanghai, April 2011.

[22] A. Rahmati, Z.Lin "Context-Based Network Estimation for Energy-Efficient Ubiquitous Wireless Connectivity" IEEE Trans on Mobile Computing, vol. 10, No.1, pp.54– 66, Jan 2011.

[23] M. Corson, R.Laroia, J. Li, V. Park, T.Richardson, and G. Tsirtsis "Toward Proximity-Aware Internetworking", IEEE Wireless Communications, vol. 17, no. 6, pp. 26-33, Dec 2010.

[24] J. Rabaey, A. Wolisz, A. Ercan, A. Araujo, F. Burghardt, S. Mustafa, A. Parsa, S. Pollin, I-H. Wang and P. Malagon, "*Connectivity Brokerage - Enabling Seamless Cooperation in Wireless Networks*" White Paper), Sept. 2010. http://bwrc.eecs.berkeley.edu/php/pubs/pubs.php/1484.html