This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2010 proceedings

Inter-Domain Routing and Data Replication in Virtual Coordinate Based Networks

Falko Dressler*, Abdalkarim Awad* and Mario Gerla†
*Dept. of Computer Science, University of Erlangen, Germany
{abdalkarim.awad,dressler}@informatik.uni-erlangen.de
†Dept. of Computer Science, University of California, Los Angeles
gerla@cs.ucla.edu

***Abstract*—In recent work it has been shown that the use of virtual coordinates or identifiers for efficient routing and data management has several advantages compared to the use of predefined addresses or geographical coordinates. However, these advantages only hold for single domain networks with limited mobility. In this paper, we discuss the challenges arising from using virtual coordinates for routing (to a particular sensor ID or to indexed data or resources) in ad hoc and sensor networks in multi-domain network scenarios. We show the feasibility of inter-domain routing by exploiting a concept that is central to most virtual coordinate approaches: the availability of data management operation using a DHT like mechanism. Based on the Virtual Cord Protocol (VCP), we show how inter-domain routing can be realized using appropriate indirections. Furthermore, we investigate the possibility of replicating data among different networks, or DHTs, to provide seamless data access in multi-domain environments. Our simulation results clearly show that both functions can be realized with only marginal overhead.**

I. Introduction

Routing in sensor network scenarios is still a challenging issue, especially taking into account the unreliability of nodes and links, the geographical spreading of the networks, and the need to completely self-organize all routing functions. The main objective is usually to access produced sensor data in a seamless manner, independent of the current topology [1]. Routing in sensor networks is dominated by two basic approaches: either based on conventional, routing table driven Mobile Ad Hoc Network (MANET) solutions or by content based routing, where the route is discovered step by step from content we are looking for. The most common example is georouting, where the content is represented by geographic coordinates of the destination. In this case, all nodes have geographic position identifiers (learned from GPS for example). For most scenarios, MANET based routing turned out to be less scalable due to routing table storage overhead as well as traffic update overhead to adjust to topology changes. The traffic overhead is due to periodic updates in proactive routing schemes or to flood search in reactive schemes. Position-based routing solutions inherently improve the situation as simple greedy routing towards the destination can be employed. However, such approaches only work well if the network is dense, as routing holes cause geographic routing to rely on inefficient face routing methods. Additionally, GPS is required for all the nodes have to be able to precisely obtain their geographic locations, and a Geo Location Service (GLS) is necessary to find the destination coordinates [2].

Recently, a number of improvements to overcome geo routing holes have been proposed. One idea is to "re-arrange" the nodes' positions appropriately to prevent routing holes. A typical example of this strategy has been described in [3]. The main idea is to use either location transformation or additional virtual location identifiers. The second, more innovative concept is to rely on virtual coordinates only and to create an "overlay" that connects the nodes and guides the search so that it does not get stuck in a hole. Recall that in geo-routing there is no overlay. Rather, a node must rely on the information heard from neighbors. Protocols like Virtual Cord Protocol (VCP) [4] and Virtual Ring Routing (VRR) [5] build their own coordinate system, which is completely independent of the geographic node positions. Current work on virtual coordinate based approaches focuses on two aspects: The provided quality of service, which is mainly an issue of optimizing the delivery ratio or even providing guarantees [6], [7], and the reliability of the system as a whole, using data replication and other redundancy increasing techniques [8]. Many of these approaches exploit the availability of the virtual ordering (numbering) of the nodes on an overlay that allows the use of Distributed Hash Table (DHT) like operations for data management. Both VRR and VCP inherently integrate a DHT that is used for routing as well as for data storage and lookup. Similarly, the recently published Prefix Routing Over Set Elements (PROSE) approach is based on distributed hashing for scalable MANET routing [9]. The previous examples use an one dimensional overlay (e.g., a ring or cord). There are also examples of multidimensional overlays like CAN [10].

However, the use of virtual coordinates, which are usually managed locally in each domain in a way that optimized routing, data management, or both, seems to make inter-domain routing extremely complicated. Many scenarios can be envisioned in which multiple (virtual coordinate based) networks are established and maintained separately, yet with a strong demand to support routing among these different networks in case of connectivity. The problem is illustrated in Figure 1. We assume that the network integrity (in terms of an ordered overlay) for a single domain is almost always ensured, which happens in typical group mobility scenarios [11]. Whenever two networks meet each other, i.e. they are overlapping,

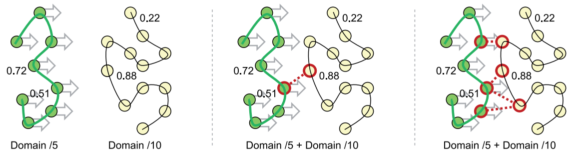

Fig. 1. VCP inter-domain routing example

physical connectivity between nodes of different domains is established. Figure 1 shows the scenario at three different points in time. As can be seen, different nodes will have to provide gateway functionality as soon as physical connection is available. The inter-domain routing is responsible for establishing adequate paths. There might be several other possible reasons that two networks are operated separately, e.g. for security, privacy, or search efficiency reasons, different applications in the same geographic area may be kept in two virtual networks. However, inter-domain routing might be necessary for efficient communication.

The main problem of MANET inter-domain routing has been first discussed in [12]. Four challenging issues have been identified: addressing, membership management, handling domain-level topology changes, and routing between the networks. As Internet-based protocols have been considered, the addressing and membership management basically targeted the IP address assignment procedure and the resulting routing problems. A cluster-based solution for inter-domain routing in MANETs has been described in [13]. Here, especially the issue of domain-level topology changes has been addressed. Using bloom filters, the effort for necessary topology updates was greatly reduced.

Motivated by this work, we investigated the issue of inter-domain routing for virtual coordinate based routing protocols, in particular focusing on our VCP approach [14]. In this paper, we show that inter-domain routing in virtual coordinate environments can be established exploiting the available DHT based data management operations. Inter-domain routing becomes feasible with only marginal overhead. Based on our VCP protocol, which we briefly introduce in Section II, we outline the basic idea of inter-domain routing in Section III. We not only consider routing issues but also provide support for inter-domain data replication. In Section IV, we present selected simulation results demonstrating the feasibility and the performance of our approach. Finally, we conclude the paper in Section V and discuss further open challenges.

II. Virtual Cord Protocol

The Virtual Cord Protocol (VCP) protocol [4], [8] has been developed keeping in mind two objectives. First, efficient routing in sensor networks should be supported and, secondly, data management should be integrated based on the concept of DHTs. In previous work, we demonstrated the capabilities of VCP and showed that VCP outperforms MANET based solutions as well as other virtual coordinate protocols such as VRR. We were also able to demonstrate that the integrated mechanisms to support failure tolerance make the protocol suitable for highly dynamic application scenarios [8].

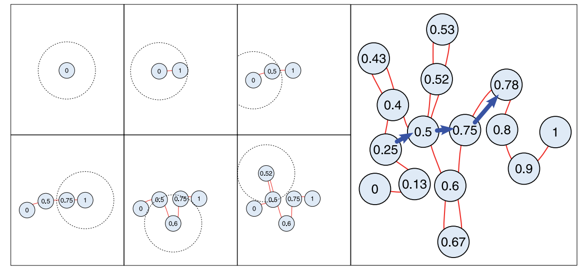

Fig. 2. Basic join operation and the resulting greedy routing process

VCP uses the concepts of DHTs to combine data management with efficient routing in sensor networks. The main idea is to arrange all the nodes in the network in the form of a virtual cord. The topology of this cord must not be "optimal" in any sense, because routing is organized by exploiting information about the physical neighbors for greedy forwarding. Nevertheless, the cord ensures the availability of at least one path between any two nodes in the network.

The cord is established using periodic `hello` messages. Besides the assigned virtual address, these messages carry all relevant information including the physical and the virtual neighbors. One node must be pre-programmed as the initial node, i.e. it gets the start position $S = 0$. Based on received `hello` messages (at least one is required) in the last time interval, a new node can determine its position in the cord. A cord is formed according to a number of simple rules. Basically, new nodes either join at one end of the cord, or get integrated, if at least two other nodes that are virtual neighbors in the cord are detected. A special rule is applied if the node has connectivity to a non-end node but not to its virtual neighbors. Then, a *virtual position* is generated at the discovered potential neighbor that is close to its virtual coordinate. This address allows the new to join between the real and the virtual position in the cord, i.e. to extend the cord without disrupting it. Figure 2 outlines the join process for a normal join (steps 1+2) and for the use of a virtual position (steps 3+4). On the right hand side, the greedy routing process is depicted. The next hop is chosen that is closest to the final destination.

An application-dependent hash function is used for associating data items to nodes; thus, both pushing to a node and pulling data from a node are supported. The same mechanism can be used for service discovery as well. Not yet considered was the case of inter-domain routing.

III. Inter-Domain Routing in VCP

We investigate inter-domain routing in the context of virtual coordinate based routing protocols [14]. This is especially challenging as no fixed intra-domain addressing exists. We target to exploit the available DHT capabilities to solve inter-domain routing using stored gateway identifiers together with the concept of indirections. The approach is designed to work with VCP, however, other virtual coordinate based routing schemes can be used as well if they support storing identifiers in an internal DHT.

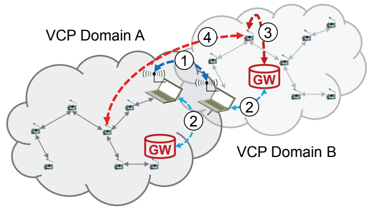

Fig. 3. Inter-domain routing using VCP

Conceptually, we associate each network with a unique domain identifier. This can be performed in VCP during the cord setup phase by assigning this ID to the start node. Then, all nodes joining the cord also obtain the domain ID. The concept is depicted in Figure 3. The periodically exchanged `hello` messages also contain the domain ID (step 1 in Figure 3). If two networks are getting into each others communication range, a node receiving `hello` messages from another domain automatically becomes a gateway node. It then stores this information into the local DHT by hashing the well-known identifier of the gateway service and storing the information at the node closest to the resulting hash value (step 2). If the gateway no longer receives `hello` messages from the detected neighbor, it removes the gateway information from the DHT. This way, the local DHT always contains the most recent gateway information and routing between neighboring domains becomes possible using simple indirections. Whenever a nodes wants to transmit a packet to another domain, it pulls the gateway address from the DHT (step 3) and then forwards the message via the gateway node (step 4).

An example is depicted in Figure 1. Two VCP domains are becoming interconnected, first by a single gateway, at a later point in time, multiple gateway nodes become available. In the following, we use the notation `node id/domain id`, e.g. 0.22/10, to identify nodes within a particular domain. Let's assume that node 0.72/5 needs to transmit a message to node 0.22/10. It first looks up an appropriate gateway. The first available gateway in our example is 0.51/5. Thus, 0.72/5 forwards the data to 0.51/5, which, in turn, forwards it to the other domain, i.e. to gateway node 0.88/10, to finally reach 0.22/10. If more than a single gateway is available, any one can be used to reach the other domain. The routing tables are distributed (and replicated, if necessary) by the DHT. An excerpt of a typical routing table is shown in Table I.

Direct communication between two nodes in arbitrary domains requires global topology information, i.e. the gateway information needs to be distributed into all VCP domains. Inter-domain routing can be supported using a shortest path algorithm together with source routing on domain level. VCP's greedy routing is only used within a domain.

Besides inter-domain routing, we integrated inter-domain data replication as an important mechanism to dynamically update information bases within different virtual domains. We enabled on-demand data replication based on the VCP data management that already supports replication within a single domain. The concept is based on a publish-subscribe mechanism. The replicator subscribes itself to the gateway service. Whenever a gateway becomes available to a yet unconnected domain, the publish-method is called, enabling the replicator to synchronize with the other domain as quickly as possible by inserting all new data items.

TABLE I
SAMPLE ROUTING TABLE

Gateway	Src Domain	Dst Domain
0.27	VCP /5	VCP /10
0.51	VCP /5	VCP /10
0.90	VCP /5	VCP /10
0.88	VCP /10	VCP /5
	...	

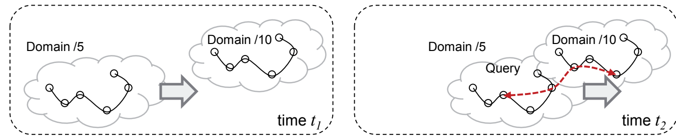

Fig. 4. Experiment setup for inter-domain routing

IV. Simulation Results

We investigated the feasibility and the performance of the inter-domain routing concept for VCP in a simulation scenario. We used our implementation of VCP for the simulation tool OMNeT++ [4] to analyze the behavior of the dynamic gateway configuration and the performance of the inter-domain routing using indirections. The basic simulation parameters are summarized in Table II.

A. *Inter-domain routing*

We configured depicted in Figure 1. We first used two networks consisting of 10 nodes each. We allow an initial setup time of 100 s to establish two VCP networks, one for each group. Within this time, a node in the mobile network creates and inserts 20 data items in this VCP domain. After the initialization, the mobile group moves towards the stationary group. After some time, the first nodes get into the radio range of the other group and they start to set up gateway information, at the end of the simulation, most of the nodes already passed the stationary group. In the second scenario, we used network 100 instead of 10 nodes in both networks to evaluate the impact of a larger number of gateway nodes

TABLE II
SIMULATION PARAMETERS

Input Parameter	Value
Number of Nodes	10+10, 100+100
Speed	$1\,\mathrm{m\,s^{-1}}$, $3\,\mathrm{m\,s^{-1}}$, $6\,\mathrm{m\,s^{-1}}$
Query period	$1\,\mathrm{s^{-1}}$
Initialization time	100 s
mac.bitrate	$2\,\mathrm{Mbit\,s^{-1}}$
mac.maxQueueSize	14 packets
mac.rtsCts	false

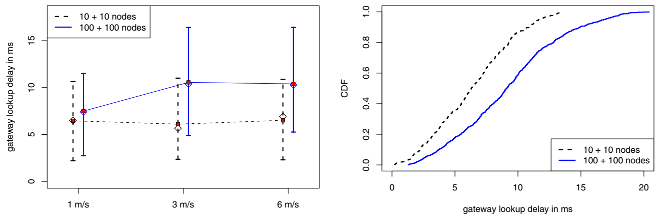

Fig. 5. Gateway lookup delay

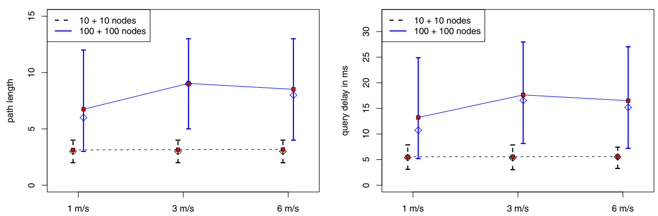

Fig. 6. Path lengths and query delay

and longer communication paths. We also varied the speed at which the mobile network is moving. We selected $1\,\mathrm{m\,s^{-1}}$, $3\,\mathrm{m\,s^{-1}}$, and $6\,\mathrm{m\,s^{-1}}$ according to the experimental setup in [13] (we adjusted the speed to the used transmission range of about $50\,\mathrm{m}$). We performed at least 5 runs for all the simulation experiments. Most of the plot shows the median as a small diamond together with error bars indicating the 10% and 90% quantiles, as well as the mean value depicted as a small red marker.

The first metric that we investigated is the gateway lookup delay. As inter-domain routing is performed using an indirection to identify a gateway to the neighboring VCP cord, this gateway ID needs to be looked up in the local DHT. Figure 5 shows the simulation results. In the plot, the lookup delay for three different speeds is drawn. The observed delays in the smaller group (10 nodes) are obviously smaller compared to the larger group (100 nodes). This trend is confirmed in the plotted CDF: a single run is plotted that is representative for all simulations. However, the node speed seems to have less impact for the smaller network compared to the large one. This effect is reasonable as the paths are longer in this case.

The path lengths are depicted in Figure 6. As can be seen, the measured number of hops corresponds to the gateway lookup delay. In the small network with only 10 nodes, the typical path is about 3 hops, whereas 10 up to 20 hops are observed in the 100 node network. Depicted are again the statistical measures for multiple runs. The higher variance in the larger network is a result of the random initial topology. We also measured the overall communication delay for inter-domain routing as depicted in Figure 6. Here, the communication between the source node and the gateway as well as the additional delay from the gateway to the destination node is measured. Again, the observed delay corresponds to the path length.

Finally, we studied the success ratio for queries in the inter-domain routing experiments. We only consider queries in the time period, where at least one gateway is available. We

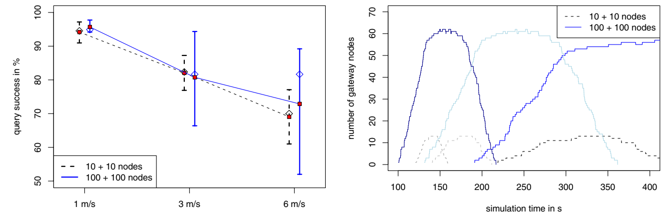

Fig. 7. Success ratio of queries and available gateways over time

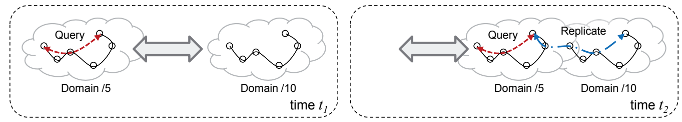

Fig. 8. Experiment setup for inter-domain data replication

expected a strong impact of the speed at which the nodes move, because this results in frequent path changes and the selection or removal of gateway nodes. Figure 7 shows the results. If the nodes move at a low speed of $1\,\mathrm{m\,s^{-1}}$, the success ratio is clearly exceeding the $90\,\%$ mark. However, at higher speeds, the success ratio drops to $80\,\%$ for $3\,\mathrm{m\,s^{-1}}$ and even to less than $70\,\%$ for $6\,\mathrm{m\,s^{-1}}$. In order to show that the nodes' speed does not influence the gateway establishment, Figure 7 also shows the number of available gateways over time. As can be seen, the contact time between the networks clearly depends on the speed the domains pass by each other, however, the distribution of the total number of available gateways is matching for the different velocities.

B. Inter-domain data replication

In a second experiment, we evaluated the feasibility of inter-domain data replication. We first used the same simulation setup as before, but executed queries only locally. Whenever two VCP domains become connected, the publish-subscribe mechanism is used to initiate data replication between both domains via the just established gateway node. As a second simulation setup, we used the setting depicted in Figure 8. The main difference is that the two networks become connected and disconnected several times to simulate the typical scenario of two independently operating networks that have to synchronize some data whenever possible, i.e. whenever getting connected.

We again evaluated both experimental setups using the same set of settings as described before, i.e. for two network sizes and three node velocities. The simulation results show the expected trends (data not shown). The query delay becomes less variable as only local queries have to be performed. Furthermore, the success ratio increases, again, because the path length is shorter compared to the inter-domain routing scenario and because the paths do not change during the experiment. Interestingly, the gateway lookup delay become more variable in the 100 node case. Obviously, the performed data replication caused a significant network load that delayed individual gateway information. The CDF shown in Figure 9 outline these effects; again, selected runs are used to demonstrate the trend.

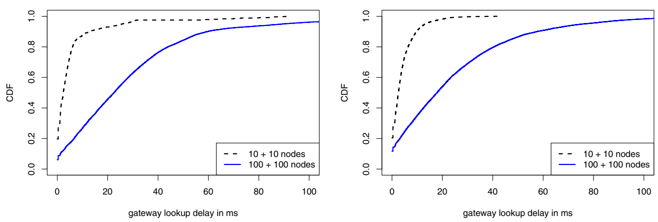

Fig. 9. Gateway lookup delay with data replication

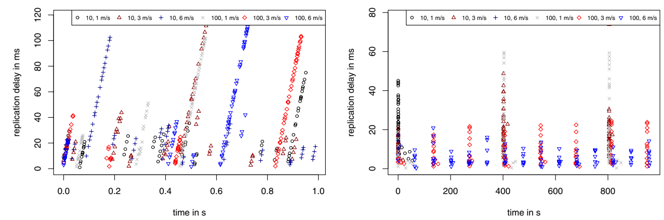

Fig. 11. Temporal behavior of the data replication

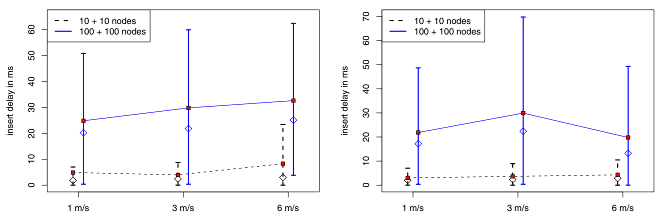

Fig. 10. Insertion delay as experienced for data replication

More prominent metrics for evaluating the inter-domain data replication quality are the delay for inserting items into the neighboring domain and the temporal behavior of this replication process. Figure 10 shows the results for the insert delay for both configurations. Shown on the left is the setup with only a single connection during the experiment. The experienced delay is quite small for the 10 node network (mostly less than 10 ms). For the 100 node network, the average delay increases as well as the variance of the insertion delay. The trend can be explained as before. The setup with multiple connections between both networks, shown on the right, shows the same effects and roughly the same delays.

Finally, the temporal behavior of the data replication is studied based on the plots shown in Figure 11. Both plots depict the replication delay over the time, each marker representing a single replication message, for the single connection scenario (left) and the multiple connection scenario (right), respectively. As can be seen, the complete replication takes roughly 1 s, with replication messages being sent in bursts. The delay-time pattern indicates single bursts with queued update messages. The different scenarios (network size, node speed) seem only marginally influencing the replication behavior. As shown in the multiple connection scenario, the main difference between the scenarios is the total number of replication messages. The slower the nodes move, the more new data items have been created to be replicated. Thus, the total delay is significantly higher. The faster the nodes move, the more connections per time can be realized, and less data needs to be transmitted.

V. Conclusion and Further Challenging Issues

In conclusion, it can be said that both inter-domain routing and inter-domain data replication can be integrated with virtual coordinate based routing and data management solutions. Based on the case study using our VCP protocol, we have shown that the routing and replication performance is well suited even in networks with significant dynamics, i.e. established and broken connections between multiple domains due to node movements. The used indirections for gateway lookup and the implemented publish-subscribe update mechanism allow a timely reaction on global topology changes.

Open issues to be studied include the scalability of the approach w.r.t. the supported number of connected network domains. Our solution performs well with a few connected domains. For larger networks, either a bloom filter based approach similar to [13] may be used or even a global DHT. Furthermore, domain splitting is not yet considered. Split detection can be supported based on the neighborhood management. The main problem is to distinguish between a node failure and a domain split.

References

[1] J. Yick, B. Mukherjee, and D. Ghosal, "Wireless sensor network survey," *Elsevier Computer Networks*, vol. 52, no. 12, pp. 2292–2330, August 2008.

[2] Y.-J. Kim, R. Govindan, B. Karp, and S. Shenker, "Geographic Routing Made Practical," in *USENIX/ACM NSDI 2005*, San Francisco, CA, 2005.

[3] R. Flury, S. V. Pemmaraju, and R. Wattenhofer, "Greedy Routing with Bounded Stretch," in *IEEE INFOCOM 2009*, Rio de Janeiro, Brazil, April 2009.

[4] A. Awad, C. Sommer, R. German, and F. Dressler, "Virtual Cord Protocol (VCP): A Flexible DHT-like Routing Service for Sensor Networks," in *IEEE MASS 2008*, Atlanta, GA, September 2008, pp. 133–142.

[5] M. Caesar, M. Castro, E. B. Nightingale, G. O'Shea, and A. Rowstron, "Virtual Ring Routing: Network routing inspired by DHTs," in *ACM SIGCOMM 2006*, Pisa, Italy, September 2006.

[6] C.-H. Lin, B.-H. Liu, H.-Y. Yang, C.-Y. Kao, and M.-J. Tasi, "Virtual-Coordinate-Based Delivery-Guaranteed Routing Protocol in Wireless Sensor Networks with Unidirectional Links," in *IEEE INFOCOM 2008*, Phoenix, AZ, April 2008.

[7] M.-J. Tsai, F.-R. Wang, H.-Y. Yang, and Y.-P. Cheng, "VirtualFace: An Algorithm to Guarantee Packet Delivery of Virtual-Coordinate-Based Routing Protocols in Wireless Sensor Networks," in *IEEE INFOCOM 2009*, Rio de Janeiro, Brazil, April 2009.

[8] A. Awad, L. R. Shi, R. German, and F. Dressler, "Advantages of Virtual Addressing for Efficient and Failure Tolerant Routing in Sensor Networks," in *IEEE/IFIP WONS 2009*, Snowbird, UT, February 2009, pp. 111–118.

[9] D. Sampath and J. Garcia-Luna-Aceves, "PROSE: Scalable Routing in MANETs Using Prefix Labels and Distributed Hashing," in *IEEE SECON 2009*, Rome, Italy, June 2009, pp. 1–9.

[10] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker, "A Scalable Content-Addressable Network," in *ACM SIGCOMM 2001*, San Diego, CA, 2001, pp. 161–172.

[11] B. Zhou, K. Xu, and M. Gerla, "Group and Swarm Mobility Models for Ad Hoc Network Scenarios Using Virtual Tracks," in *MILCOM 2004*, Monterey, CA, 2004.

[12] C.-K. Chau, J. Crowcroft, K.-W. Lee, and S. H. Y. Wong, "Inter-Domain Routing Protocol for Mobile Ad Hoc Networks," in *ACM MobiArch 2008*, Seattle, WA, 2008, pp. 61–66.

[13] B. Zhou, Z. Cao, and M. Gerla, "Cluster-based Inter-domain Routing (CIDR) Protocol for MANETs," in *IEEE/IFIP WONS 2009*, Snowbird, UT, February 2009, pp. 19–26.

[14] F. Dressler, A. Awad, R. German, and M. Gerla, "Enabling Inter-Domain Routing in Virtual Coordinate Based Ad Hoc and Sensor Networks," in *ACM MobiCom 2009, Poster Session*, Beijing, China, September 2009.