



HAL
open science

Towards Scalable MANETs

Emmanuel Baccelli, Jochen Schiller

► **To cite this version:**

Emmanuel Baccelli, Jochen Schiller. Towards Scalable MANETs. 8th IEEE International Conference on ITS Telecommunications, Oct 2008, Phuket, Thailand. pp.133 - 138, 10.1109/ITST.2008.4740243 . hal-00651438

HAL Id: hal-00651438

<https://hal.science/hal-00651438>

Submitted on 13 Dec 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards Scalable MANETs

E. Baccelli, *INRIA** J. Schiller, *Freie Universität Berlin*

Abstract— In the near-future, self-organized networking is expected to become an important component in ITS, and in the Internet architecture in general. An essential challenge concerning the integration of this new component is the accomplishment of scalable and efficient mobile ad hoc routing. This paper overviews considerations relative to the design of such MANET protocols inside the framework provided by the IETF, stating the need for new hybrid protocols and architecture which offer a gradual transition from "traditional" MANET routing towards scalable MANET routing integrated in the Internet. This paper also proposes a tentative solution in this domain: DHT-OLSR, based on OLSR enhanced with dynamic clustering and distributed hash table routing.

Index Terms – Ad hoc, Scalability, IP, Routing, Network, Wireless, Standardization, IETF

Submission Area – Policy, Standardization & Service Applications in ITS

I. INTRODUCTION

Self-organized networking is foreseen as an important component in the Internet's near-future architecture, which is managed by the IETF [11], the organization in charge of elaborating standard Internet protocols published as RFC documents [18]. Such spontaneous networks are deemed to be of various shapes and nature, from smart dust to autonomous community networks, and from Intelligent Homes to vehicular networks and ITS. A substantial challenge concerning the integration of these new components is the accomplishment of scalable and efficient mobile ad hoc routing.

A. The IETF Framework

Within the IETF, mobile ad hoc routing protocols are standardized in a specific suborganization, the MANET [12] working group. Through this working group, the IETF is currently standardizing two ad hoc routing protocols: OLSR [4] [3] and DYMO [5] (the heir to AODV [7]). However, while these protocols perform well in small topologies, they fail to provide a functional ad hoc network when the number of nodes involved in the ad hoc network increases. Indeed, proactive approaches (such as OLSR), and reactive approaches (such as DYMO), both require too much network-wide control signalling, which quickly saturates the available bandwidth when nodes participating

in the ad hoc network are too numerous or too mobile (see for instance [21]).

In order to overcome this scalability issue within the IETF framework, two approaches can be envisioned: (i) design an entirely new protocol, or (ii) extend and improve the scalability of an existing protocol, DYMO or OLSR. One can easily argue against a design "from scratch" in a context where interoperability and legacy are of the utmost importance, such as in the IETF. Moreover, the time that was necessary to finalize the existing protocols (approximately a decade) suggests that a solution following approach (i) would experience quite some delay while it matures within this organization. On the other hand, standardizing an extension of an existing protocol has proven to be a much quicker process. This paper therefore takes on approach (ii), or in other words: focus is on specific mechanisms that can be used in conjunction with "usual" ad hoc routing in order to scale better.

B. Constraints and Considerations

Some requirements are to be considered for the compatibility of new mechanisms with legacy protocols and the Internet architecture in general. For instance, IP routers are basically supposed to either (i) directly forward packets as they are received, or else (ii) drop them if no appropriate route is currently available. This philosophy disqualifies some mechanisms a priori, including some reactive approaches which rely too heavily on buffering traffic before discovering (and forwarding along) an appropriate route. We note that buffering is also not an option in ad hoc networks that may encompass small, memory-challenged devices.

Further considerations to be taken into account include control traffic confinement. For instance, most of Internet's scalability is based on the clear definitions of natural bounds between networks, which "filter" inter-network control traffic. Such natural limits are not found in MANETs, and this lack is thus an obstacle towards scalability.

In the past, several "hybrid" solutions have been proposed in this space (such as ZRP [19]), that aim at improving scalability by relying on simple combinations of proactive routing (OLSR) with reactive routing (DYMO or AODV) while dividing the network into more or less independent routing zones. However, such solutions still require too

* Corresponding author: emmanuel.baccelli@inria.fr, LIX - Ecole Polytechnique, 91128 Palaiseau Cedex, FRANCE

much network-wide control signalling (most of them are also fragile in face of node mobility), and thus do not scale as desired.

Another category of solution has been proposed with protocols such as FOLSR [21] (see also FSR [22]), that combines proactive routing (OLSR) with a mechanism that allows to keep the amount of network-wide control signalling bounded when the number of node increases, by gradually diminishing the frequency of long range routing information updates. While this type of approach scales as desired, convergence time may be an issue in some scenarios, as the decrease of update frequency may slow down the discovery of some part of the network.

A novel and interesting approach is the cross-fertilization between the fields of mobile ad hoc routing and peer-to-peer networking. Indeed, both fields are dealing with similar issues, including decentralized network formation and maintenance. For instance, specific distributed hash table (DHT) approaches such as VRR [13], MADPastry [14], Ekta [16], SSR [17], can provide routing solutions in large scale ad hoc networks. However, while such routing approaches are very interesting in theory, they do not take into account legacy and compatibility with currently deployed Internet protocols. On the contrary, our approach bridges between theory and Internet reality.

C. Appropriate Protocol Design

This paper thus introduces an instance of hybrid solution designed with respect to the above described framework. This new protocol is called DHT-OLSR, and is based on dynamic OLSR clustering enhanced with distributed hash tables routing based on MADPastry. The latter is chosen over other DHT approaches for its convenient property of explicitly considering physical locality in the construction of its routing tables for efficient key-based unicast routing. On the other hand, OLSR is chosen over other ad hoc routing approaches as it is the legacy proactive MANET routing standardized by the IETF, and thus provides natural integration with the Internet infrastructure.

The remainder of the paper is organized as follows. The next two sections will briefly recall the principles of OLSR routing on one hand, and MADPastry's key-based routing on the other hand. The following sections will then describe the functioning of DHT-OLSR, followed by a section evaluating this solution, providing experimental results.

II. PROACTIVE AD HOC ROUTING

DHT-OLSR builds on a legacy ad hoc routing protocol: OLSR (Optimized Link State Routing [3]). Therefore in this section we briefly outline OLSR, giving the essential facts about the protocol that are interesting in the context of this paper.

Periodic control messaging. As a proactive link-state routing protocol, OLSR employs the periodic exchange of control messages in order to accomplish topology discovery and maintenance. This exchange results in a topology map being present in each node in the network, from which a routing table can be constructed. Basically, OLSR employs two types of control messages: HELLO messages and TC (Topology Control) messages. HELLO messages have local scope and are exchanged periodically between neighbor nodes only, essentially tracking the status of links between neighbors. On the other hand, TC messages have larger scope and are emitted periodically to diffuse link-state information throughout the entire network.

Multipoint relaying. This operation of diffusing a message to the entire network – also called flooding – is optimized in OLSR with a mechanism called Multi-Point Relaying [9] which drastically reduces the cost of flooding operation, by having each node select a minimal set of “relay neighbors” responsible for relaying flooded packets.

III. DHT KEY-BASED ROUTING

Aside from OLSR, DHT-OLSR also uses a key-based routing mechanism: MADPastry [14], a DHT substrate specifically designed for mobile ad hoc networks. In this section, we will thus give a brief overview of this mechanism.

Overlay Space. Each node in a MADPastry network assigns itself a unique virtual ID, for instance via hashing its IP address, which defines its logical position on the virtual overlay ID space [15]. Each MADPastry message contains a message key (an ID from the virtual overlay ID space) in its header. MADPastry then routes the message to the node in the network that is currently responsible for the message key, *i.e.* the node whose virtual ID is currently numerically the closest to the message key among all nodes in the network.

Clusters. In order to reduce potential route stretching, MADPastry exploits the physical locality of the nodes in the construction and maintenance of its overlay by using random landmarking [23]. Well-known landmark *keys* divide the virtual overlay ID space into equal sections, *e.g.* 16 keys with hexadecimal IDs "0800...000", "1800...000", ..., "F800...000", *etc.* The nodes whose virtual IDs are currently numerically closest to the landmark keys temporarily become landmark nodes and periodically issue beacon messages. Nodes overhear these beacon messages and periodically determine the physically closest landmark node, *e.g.* in terms of hops. If need be, a node assigns itself a new virtual ID sharing the same prefix with the closest landmark node and (re-)join the network under its new ID. Thus nodes that are close to each other in the virtual overlay ID space are also likely to be close to one another physically, as demonstrated in Fig. 1.




Fig. 1. Spatial distribution of virtual ID prefixes in a 250 node MADPastry network. Equal symbols of equal shades represent equal virtual ID prefixes.

Multiple Routing Tables. MADPastry relies on AODV to maintain a routing table for local physical routes. On top of this, MADPastry maintains a sparse Pastry routing table and a standard Pastry leaf set for indirect routing. The Pastry routing table only needs to contain one entry for each distinct cluster’s virtual ID prefix, and each node simply pings its “left” (*i.e.* the node who has the largest virtual ID smaller than the node’s own) and “right” (*i.e.* the node who has the smallest virtual ID larger than the node’s own) leaf periodically, while all other routing entries are gained or updated passively by overhearing packets. Overlay hops are then determined by the Pastry tables, while physical routes between overlay hops are determined by the AODV table.

IV. THE DHT-OLSR PROTOCOL

In this section, we introduce a novel OLSR extension named DHT-OLSR. The principle of this hybrid protocol is that each node runs OLSR locally within the cluster of nodes that are currently within a certain scope (*i.e.* within a given number of hops). When a node needs to forward a packet towards a destination that is not currently listed in the routing table maintained by OLSR, the node uses DHT-based unicast to immediately forward the packet towards this remote destination. In other words, nodes forward data packets without buffering them due to route discovery latency, a feature that is highly desirable as MANET nodes’ capacity may be limited.

For this purpose, each node running DHT-OLSR maintains a regular OLSR routing table, providing efficient and low-delay routing. However, in order to avoid larger amount of maintenance traffic as the network size increases,




Fig. 2. Visualization of DHT-OLSR’s routing concept seen from the point of view of node 17. The grey-shaded area around node 17 shows the 2-hop radius in which node 17’s TC messages are propagated. The dotted circle represents the radio range of node 17.

DHT-OLSR nodes confine their OLSR signalling to a local scope. This is achieved by simply limiting the TTL of OLSR messages. For instance, nodes might set the TTL of the TC messages that they issue to 2, which effectively places each node at the center of its own OLSR cell with a diameter of 4 hops. Hence, whenever a node forwards a data packet, it first tries to lookup the route in its OLSR routing table. If a valid route is found, the data packet is forwarded to the next hop on the path towards the destination. In case no route could be found in the OLSR routing table, the node engages into the low maintenance overhead DHT routing scheme shown in the following. Fig. 2 visualizes the base of the DHT-OLSR routing protocol. It depicts node 17 (whose virtual MADPastry ID is 75A1FF) at the center of its local scope (2 hop radius) OLSR cell. Assuming equal radio ranges, the OLSR routing table of node 17 will contain routes to all nodes within this 2-hop radius and to some neighbors three hops away (dark grey nodes around the OLSR cell in Fig. 2), those which are included in the TC messages received from 2-hop neighbors.

Whenever a node wants to forward a data packet to a destination for which it cannot find a valid route in its OLSR routing table, the node engages in DHT-based routing. For this purpose, each DHT-OLSR node also runs a special instance of MADPastry to communicate with more remote nodes in the network. Deviating from the original MADPastry, DHT-OLSR do not run AODV, and use instead the OLSR routing table. This modified instance of MADPastry provides key-based routing, *i.e.* whereby packets are no longer routed based on an IP destination address but rather based on a virtual ID from the MADPastry ID space.

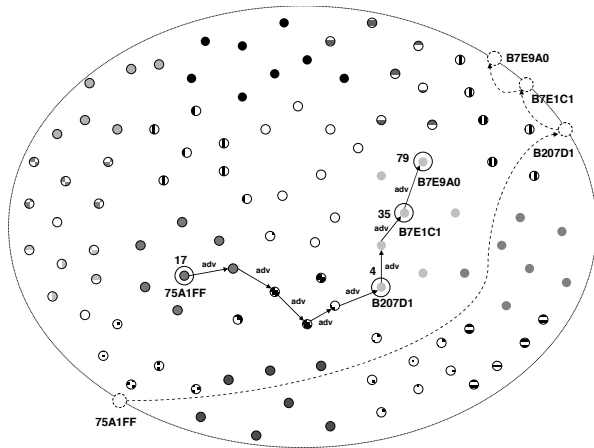


Fig. 3. DHT-OLSR virtual ID advertising. Equal symbols of equal shade represent equal virtual ID prefixes. Overlay hops are dashed, physical hops are plain.

The question now arises how DHT routing can be used to deliver a packet to its given IP destination address. To resolve IP destination addresses, DHT-OLSR uses a unicast scheme based on the approach presented in [1]. In order for nodes to be able to resolve node addresses to their corresponding current virtual IDs, each DHT-OLSR periodically advertises its current virtual ID to the network. For this purpose, a node x hashes its IP address into the virtual overlay ID space to obtain its advertisement key. Using this key, the node will then send an advertisement packet containing its current virtual ID towards the node currently responsible for this advertisement key. That node will, upon reception of the advertisement packet, store node x 's address along with its current virtual ID, *i.e.* its key.

Conversely, node address resolution works in an analogous fashion. Whenever a node y wants to send a data packet to a given destination using DHT unicast, it needs to resolve the destination node address in order to obtain its current virtual ID. If node y does not know the destination's current virtual ID, it will simply hash the destination node address and acquire the same advertisement key that the destination used to advertise its current virtual ID. Next, node y will send the data packet towards that advertisement key, which by definition (assuming consistent DHT routing), will be delivered to the node that has previously received the destination node's advertisement. That node will then forward the data packet using the key provided in the advertisement and the data packet will be delivered to the given destination node.

The remainder of this section will depict virtual ID advertising, resolution and data forwarding via concrete examples. Fig. 3 depicts an instance of DHT-OLSR's virtual address advertising. Node 17 intends to advertise its current virtual ID 75A1FF. For this purpose it simply hashes its

IP address, which yields the following advertisement key: $\text{hash}(17) = \text{B7E97D}$. Using MADPastry, node 17 will then send an advertisement packet containing its current virtual ID towards the node currently responsible for its advertisement key. In the first overlay hop, MADPastry routing delivers the packet to node 4 with virtual ID B207D1. In the next overlay hop, node 4 forwards the packet on to node 35 (virtual ID B7E1C1) who, then, forwards the packet to node 79 whose virtual ID B7E9A0 is numerically closest to node 17's advertisement key (B7E97D) among all other nodes' virtual IDs. Thus, node 79 is responsible for the packet and stores the tuple (17, 75A1FF) in its cache.

Fig. 4 visualizes how DHT-OLSR nodes use DHT unicasts to forward data packet to remote nodes. In this example, node 51 wants to send a data packet to node 17. After discovering that there is no valid route to destination 17 in its OLSR routing table, node 51 engages in a DHT unicast to deliver the data packet. Since node 51 does not know node 17's current virtual ID, it hashes the destination address into the MADPastry virtual ID space and, by definition, acquires the same advertisement key as node 17 did for its advertisement in the example above: $\text{hash}(17) = \text{B7E97D}$. Next, node 51 forwards the data packet towards that key. Assuming consistent DHT routing, this will deliver the packet to the same node that has previously received node 17's advertisement: in the first overlay hop, the packet is routed to node 42 with virtual ID B7E705, who then forwards the packet on to node 79 (virtual ID B7E9A0). Obviously, node 79 can now resolve node 17's address and then sends the packet towards node 17's current virtual ID. As node 17 is trivially responsible for its own virtual ID, the data packet will, thus, be delivered over node 92 (virtual ID 7FF05C) to the original destination (node 17).

V. EVALUATION OF THE DHT-OLSR PROTOCOL

In this section we evaluate DHT-OLSR by means of ns-2 simulations comparing the performance of this protocol with those of AODV, OLSR.

In order to evaluate the general feasibility of our approach, we have implemented DHT-OLSR in the network simulator ns-2. In order to put the performance of DHT-OLSR into perspective, its results were compared to that of a conventional proactive (OLSR) and reactive routing protocol (AODV, and more precisely the AODV-UU implementation 0.9.1 for ns-2.). For all conducted simulations, all nodes were moving around according to the Random Waypoint model at a constant velocity of 1.4 m/s (quick walking pace) with 0s pause time (constant node mobility).

Nodes were communicating with each other in an ad hoc mode using the 802.11 communication standard with a

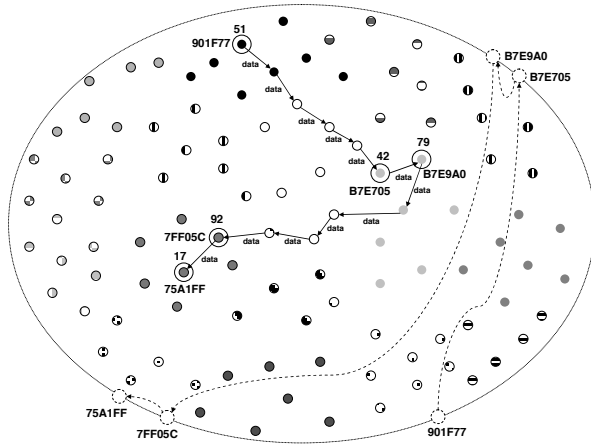


Fig. 4. DHT OLSR virtual ID resolution and data forwarding. Equal symbols of equal shade represent equal virtual ID prefixes. Overlay hops are dashed, physical hops are plain.

transmission range of 250m. We evaluated different network sizes of 50, 100, 150, 200, and 250 nodes while keeping the node density fixed at 100 nodes per square kilometer. We used a simple traffic pattern whereby each node sends a data packet to a randomly picked destination every 10 seconds. Each run simulated the course of one hour. We have analyzed the results following the two metrics listed below:

Packet Delivery Ratio (PDR). This figure represents the percentage of data packets that are successfully delivered to their respective destinations.

Total Amount of Traffic. Since many different packet types (e.g. AODV route requests, MADPastry packets, OLSR control packets, etc.) of various packet lengths are transmitted during a simulation run, we are not evaluating the total packet count. Instead, we are considering the total network traffic (in *kbytes*) that is created during the simulated hour. Whenever a node forwards a packet, this figure is increased by the packet size. This figure includes all routing and application level packet types (data packets, control packets, etc.).

Fig. 5 shows the PDR achieved by the respective routing protocols for various network sizes. As can be seen, DHT-OLSR's PDR decreases only slightly with an increasing network size and remains at or above 90% in all scenarios considered. Both the conventional reactive and the proactive routing protocol, however, do not scale well to large network sizes and see their respective PDRs fall drastically as the number of nodes in the network increases. The reason for this is that DHT-OLSR mostly uses local routes (also for MADPastry overlay routing, see [1]) that are updated through local control traffic (as explained in the previous section) whereas both AODV's and OLSR's




Fig. 5. Packet delivery ratios vs. number of nodes. Diamond for DHT-OLSR, triangle for OLSR, square for AODV.




Fig. 6. Total traffic vs. number of nodes. Diamond for DHT-OLSR, triangle for OLSR, square for AODV.

control traffic affects the entire network, which consumes more and more bandwidth as the network size increases. This effect also becomes clear in Fig. 6 which shows that DHT-OLSR also produces significantly less overall traffic than the two reference protocols do. Again, this is due to the fact that DHT-OLSR nodes generally are only affected by local control traffic whereas both OLSR's and AODV's control packets are broadcast.

VI. CONCLUSION

In this paper, we have stated the need for novel hybrid approaches to ad hoc routing in order to provide scalability in MANETs. The considerations taken into account in this discussion were with respect to the framework provided by the IETF. The paper also introduced a novel protocol

which fits this framework, DHT-OLSR, designed for large and mobile MANET topologies, based on a hybrid scheme mixing ideas coming from both the mobile ad hoc routing field, and the peer-to-peer networking field. The protocol uses key-based unicast routing to scale to very large MANET topologies, along with OLSR routing for natural compatibility with IP infrastructure, offering an architecture that may introduce a gradual transition towards integrating scalable MANET routing in the Internet.

Simulations so far show that DHT-OLSR outperforms the protocols standardized by the IETF in a variety of large MANET topologies. By using the hybrid mechanisms described in this paper, a particular feature of DHT-OLSR is that it drastically reduces the impact of remote control traffic on the bandwidth available locally. Since DHT-OLSR organizes the nodes into virtual clusters of physically close nodes, control traffic generated by DHT-OLSR will generally not affect nodes in other clusters. Therefore, a node will mainly have to bear only local control traffic, limited to a radius of a few hops. Such control traffic confinement is the key challenge towards scalability in MANETs.

Future work in this direction will include trying out different hybrid schemes using alternative ad hoc routing components or alternative key-based routing components. Potential enhancements to DHT-OLSR are also envisioned such as mechanisms to dynamically adjust OLSR parameters or to adapt the number of MADPastry clusters. Moreover, further evaluations of DHT-OLSR will be conducted, both through simulations and real physical testbed implementations. Here, the performance of DHT-OLSR should also be compared against a broader variety of ad hoc routing protocols, such as ZRP or F-OLSR.

REFERENCES

- [1] T. Zahn, J. Schiller, "DHT-based Unicast for Mobile Ad Hoc Networks," MP2P Proceedings, 2006.
- [2] E. Baccelli, "OLSR Scaling with Hierarchical Routing and Dynamic Tree Clustering," NCS Proceedings, 2006.
- [3] T. Clausen, P. Jacquet Ed. "RFC 3626: Optimized Link State Routing Protocol." Internet Engineering Task Force (IETF) Request For Comments, <http://ietf.org/rfc/rfc3626.txt>, 2003.
- [4] T. Clausen, C. Dearlove, P. Jacquet, "The Optimized Link State Routing Protocol version 2," Internet Engineering Task Force (IETF) Internet Draft (work in progress), draft-ietf-manet-olsrv2-04, 2007.
- [5] I. Chakeres, C. Perkins, "Dynamic MANET On-demand (DYMO) Routing," Internet Engineering Task Force (IETF) Internet Draft (work in progress), draft-ietf-manet-dymo-10, 2007.
- [6] S. Corson, J. Macker, "RFC 2501: Mobile Ad hoc Networking (MANET), Routing Protocol Performance Issues and Evaluation Considerations," Internet Engineering Task Force (IETF) Request For Comments, <http://ietf.org/rfc/rfc2501.txt>, 1999.
- [7] C. E. Perkins, E. M. Royer, S. R. Das, "RFC 3561: Ad Hoc On-Demand Distance Vector Routing," Internet Engineering Task Force (IETF) Request For Comments (experimental), <http://ietf.org/rfc/rfc3561.txt>, July 2003.
- [8] D. Johnson, D. Maltz, Y. Hu, "The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks," Internet Engineering Task Force (IETF) Request For Comments (experimental), <http://ietf.org/rfc/rfc4728>, 2007.
- [9] A. Qayyum, L. Viennot, A. Laouiti, "Multipoint Relaying: An Efficient Technique for Flooding in Mobile Wireless Networks," INRIA Research Report RR-3898, March 2000.
- [10] P. Jacquet, A. Laouiti, P. Minet, L. Viennot, "Performance Evaluation of Multipoint Relaying in Mobile Ad Hoc Networks," Networking 2002 Proceedings, 2002.
- [11] Internet Engineering Task Force (IETF), <http://www.ietf.org>
- [12] The Mobile Ad-Hoc Networks Working Group (MANET), <http://www.ietf.org/html.charters/manet-charter.html>
- [13] M. Caesar, M. Castro, E. Nightingale, G. OShea, and A. Rowstron, "Virtual Ring Routing: Network routing inspired by DHTs," in Proc. of ACM SIGCOMM'06, September 2006.
- [14] T. Zahn and J. Schiller, "MADPastry: A DHT Substrate for Practicably Sized MANETs," in Proc. of ASWN, June 2005.
- [15] A. Rowstron, P. Druschel. "Pastry: Scalable, distributed object location and routing for large scale peer-to-peer systems". In Proc. of Middleware, November 2001.
- [16] H. Pucha, S. M. Das, and Y. C. Hu, "Ekta: An Efficient DHT Substrate for Distributed Applications in Mobile Ad Hoc Networks," in Proc. of WMCSA 2004, December 2004.
- [17] T. Fuhrmann, "Scalable Routing for Networked Sensors and Actuators," in Proc. of SECON 2005, September 2005.
- [18] S. Bradner, "The Internet Standards Process," Internet Engineering Task Force (IETF) Request For Comments, <http://ietf.org/rfc/rfc2026>, 1996.
- [19] Z. J. Haas, M. R. Pearlman, and P. Samar, "The Zone Routing Protocol (ZRP) for Ad Hoc Networks," Internet Engineering Task Force (IETF) Internet Draft, draft-ietf-manetzone-zrp-04, 2002.
- [20] M. Jiang, J. Li, Y. Tay, "Cluster Based Routing Protocol," Internet Engineering Task Force (IETF) Internet Draft, draft-ietf-manet-cbrp-spec-00, 1998.
- [21] C. Adjih, E. Baccelli, T. Clausen, P. Jacquet, R. Rodolakis, "Fish Eye OLSR Scaling Properties," published in the IEEE sponsored Journal of Communication and Networks (JCN), Special Issue on Mobile Ad Hoc Wireless Networks, 2004.
- [22] M. Gerla, G. Pei, X. Hong, T. Chen, "Fisheye State Routing Protocol for Ad Hoc Networks," Internet Engineering Task Force (IETF) Internet Draft, draft-ietf-manet-fsr-00, 2001.
- [23] R. Winter, T. Zahn, J. Schiller, "Random Landmarking in Mobile, Topology-Aware Peer-to-Peer Networks". In Proc. of FTDCS, May 2004.