An Adaptive Approach to Group Communications in Multi-Hop Ad Hoc Networks

Kumar Viswanath, Katia Obraczka
University of California, Santa Cruz
Computer Engineering Department
E-mail: {kumarv,katia}@cse.ucsc.edu

The diverse nature of mobile wireless ad hoc networks, or MANETs, makes it almost impossible for a single routing protocol to perform well under a wide range of operating conditions. We propose an adaptive approach to group communication that allow nodes to switch among different routing protocols on the fly in order to adapt to current network conditions. Considering that MANET's are generally deployed in mission critical applications, our goal is to provide high reliability and timeliness guarantees in the presence of a wide range of network conditions (e.g., mobility, network load). To this end we develop an *adaptive flooding protocol* in which nodes can dynamically switch among different flooding variations, namely *scoped-*, *plain-*, or *hyper* flooding based on their perception of current network conditions. We employ relative velocity and perceived network load as the criteria nodes use to switch among protocols. Simulation results comparing our adaptive protocol against two of the best performing MANET multicast protocols, namely ODMRP and M-AODV, show considerable performance benefits, under various MANET scenarios. We thus argue that the proposed protocol can be used as the basis for developing adaptive, integrated routing techniques for the integrated media networks of the future.

1 Introduction

In the past five years, routing in wireless, mobile multi-hop ad hoc networks, or MANETs, has received considerable attention from the network research community. More recently, as it became clear that group-oriented services are one of the primary classes of applications targeted by MANETs, a number of multicast routing protocols for MANETs have been developed. The On Demand Multicast Routing protocol (ODMRP) [1] and Multicast-Ad hoc On Demand Distance Vector (MAODV) [2] are examples of *on-demand* multicast routing protocols where a route is established only when a source has data to send. Although these protocols are known to perform well in constrained mobility MANETs, it has been shown that their performance degrades under more stringent network conditions such as high mobility and traffic load [3].

In general, we believe that no single multicast protocol is optimal for all MANET scenarios. We envision that future internetworks will consist of a wired backbone and a collection of wired, fixed-infrastructure mobile, and ad hoc networks as leaves. We argue that a "global" multicast solution for the future internet will include specialized solutions for each type of network, as well as mechanisms for integrating these solutions.

Our long-term goal is to provide seamless integrated multicast service whereby a single multicast group can span all network types (fixed, fixed mobile, and different types of MANETs). This will allow a given host to partake in multicast communication regardless of the underlying network type. Therefore hosts will have to dynamically switch among different multicast routing mechanisms as they move from one network to another. To our knowledge, there is little or no experience in the network research community in multicast protocol interoperation (albeit, some proposals have been floated in the IETF) or adaptation [4,5,6].

2 Focus

In this paper we concentrate on the problem of providing adaptive multicast for different types of MANET scenarios. Important MANET applications, including battlefield, disaster and emergency rescue operations, are mission critical in nature and require MANET protocols to provide high delivery and timeliness guarantees in the presence of mobility and permanent or temporary outages.

We propose an *adaptive flooding* protocol which, besides plain flooding, uses two other flooding variations: *scoped*- and *hyper* flooding. While scoped flooding tries to reduce re-broadcasts and consequently overhead and is thus well suited to constrained mobility environments, hyper flooding's goal is to provide high delivery guarantees at the routing layer[a]. It does so by re-broadcasting packets based on some heuristics (e.g., the acquisition of new neighbors). The price hyper flooding pays is its high overhead.

In the adaptive routing protocol we propose, nodes switch among flooding variations based on their own perception of current network conditions. We implement two distinct switching criteria: the first one is based on relative node velocity and the second on perceived network load. A detailed description of the different operating modes as well as the criteria for switching between modes is explained in detail in Section 3.

The primary motivation to use flooding as the basis for our adaptive multicast routing framework, is that flooding and its variations perform considerably better than other protocols such as MAODV and ODMRP [3] over a wide range of mobility and traffic load conditions. Moreover, since flooding and its variations inter-operate easily, it made sense to choose to integrate variants

[a]We do not employ any other reliability mechanism such as positive- or negative acknowledgements.

of flooding into a single adaptive protocol as our first step towards designing adaptive multicast routing protocols.

This paper analyzes the performance of our adaptive flooding protocol and compares it with ODMRP and MAODV. Through simulations, we study the performance of the protocol for multicast communication and consider "synthetic"- as well as "more typical" MANET scenarios such as emergency rescue operations and conferencing. Our results show that in some scenarios, adaptive flooding achieves packet delivery ratios on the order of 20% higher than ODMRP and MAODV.

While we propose flooding variations as the first step in investigating the merits of adaptive- over non-adaptive routing mechanisms, we are also currently investigating other adaptive protocols which are not based on flooding.

3 Protocol Overview

Each node running our adaptive flooding protocol is capable of operating in any of three modes: *scoped*-, *plain*-, and *hyper* flooding. Individual nodes dynamically switch among the different operating modes according to their own perspective of current network conditions. The different criteria for switching between modes is explained in detail later in this section.

3.1 Scoped Flooding Mode

The basic principle behind scoped flooding is the reduction of re-broadcasts to avoid collisions and minimize overhead. Scoped flooding is suitable for constrained mobility environments (e.g., conference scenarios) where nodes do not move much and thus plain flooding will likely yield unnecessary redundant re-broadcasts. In fact, S. Ni et. al.[7] show that the coverage area of subsequent retransmissions reduces drastically and drops down to 0.05% when the number of retransmissions is greater than 4.

Different heuristics can be used in deciding whether to re-broadcast a packet. In our scoped flooding implementation, each node periodically transmits *hello* messages which also contain the node's neighbor list. Nodes use *hello* messages to update their own neighbor list and add received lists to their neighbor list table. When a node receives a broadcast, it compares the neighbor list of the transmitting node to its own neighbor list. If the receiving node's neighbor list is a subset of the transmitting node's neighbor list, then it does not re-broadcast the packet. In our simulations we did not require neighbor lists to be strict subsets of one another. An 85% overlap was considered sufficient to prevent re-broadcasts and this was obtained after extensive simulation-based analysis of scoped flooding.

3.2 Hyper Flooding Mode

Hyper flooding is suitable for highly mobile scenarios where high reliability is required. The price to pay for the additional reliability is of course higher overhead.

Nodes in hyper flooding mode periodically transmit *hello* messages. When a neighbor receives a *hello* message, it adds the *hello* message originator to its neighbor list (if the list does not already contain that node). Similarly to plain flooding, when a node receives a data packet, it simply re-broadcasts the packet and also queues it in its packet cache. Additionally, re-broadcasts are triggered by two other events: receiving a packet from a node which is not in the current neighbor list or receiving a *hello* message from a new node. In these cases, nodes transmit all packets in their cache. The rationale behind re-broadcasts is that "newly acquired" nodes could have possibly missed the original flooding wave on account of their mobility. This increases overall reliability by ensuring that new nodes entering the transmission region of a node receive data packets which they otherwise would have missed. Nodes periodically purge their packet cache to prevent excess re-flooding of older packets.

3.3 Switching Among Protocols

One fundamental issue in the design of adaptive integrated multicast is deciding when a node should switch protocols and which protocol to switch to. Every node needs to make its own decision based on its perception of current network conditions. For the current version of adaptive flooding, we chose *relative velocity* and *network load* as the preliminary criteria nodes use to switch among the different flooding variations. The rationale for using relative velocity and network load in deciding when to switch operating modes is based on some of our previous studies. Our results show that multicast routing performance is highly dependent on mobility and network traffic load [3].

The proposed relative-velocity based switching criterion works as follows. Nodes send velocity (speed and direction) information as part of *hello* messages. Each node is then able to compute its velocity relative to all its neighbors. We use only immediate neighbor information to calculate a node's relative velocity. Each node maintains a running average, as well as the minimum and maximum value of relative velocity for the past five time windows. Based on the current value of relative velocity and its past history, each node adaptively chooses a *low_threshold* and a *high_threshold* value for the current time window. If the current value of relative velocity is higher than *high_threshold*, the node switches to hyper flooding mode. If the relative velocity is below *low_threshold*, scoped flooding is used. Otherwise, the node switches to plain flooding. We assume that nodes will be able to obtain information on their own velocity (e.g., from an external device such as a tachometer).

Switching based on network load uses MAC-layer collisions as an indicator of network traffic. We chose collisions instead of nominal network load because it is possible that certain nodes may have very sparse neighbor sets, allowing those nodes to communicate even at high loads with low collision. In this switching method each node computes the total number of collisions that have occurred in the current time window. Similar to the velocity criterion, each node adaptively computes a *low_threshold* and a *high_threshold* value for the current time window. If the current number of collisions is lower than *low_threshold*, the node switches to hyper flooding mode. If the number of collisions is greater than *high_threshold*, scoped flooding is selected.

4 Methodology

In our study we compare the performance of adaptive flooding against ODMRP [8] and M-AODV [2], two of the best performing protocols, the first in the mesh-based and the second in the tree-based category of MANET multicast routing, respectively [3]. We use the network simulator `ns-2` for our simulations. Some of the MANET scenarios we simulated were generated using a scenario generator for ad hoc networks [9] and will be described in greater detail below.

4.1 MANET Scenarios

We use two types of MANET scenarios in our simulations. In “synthetic” scenarios, parameters such as mobility, traffic load and multicast sender and receiver population are varied over an arbitrary range of values. We also define more “concrete” environments reflecting specific MANET applications, namely impromptu teleconferencing and disaster relief/battlefield scenarios. The scenario generator [9] was used to generate conferencing and rescue scenarios for our experiments.

For the synthetic scenarios, 50 nodes are placed in a 1000 m^2 field. Each node transmits a maximum of 1000 packets (256 bytes each) at various times during the simulations. Nodes’s channel bandwidth is set to 2 Mbit/sec and their transmission range is 225 meters. The mobility model used is a modified version of the random-waypoint model (also known as the bouncing ball model). A constant bit rate (CBR) traffic generator was used for adaptive flooding based on relative velocity, and an ON-OFF traffic generator was used for adaptive flooding based on network load.

4.2 Concrete Scenarios

The conference scenario consisted of a total of 50 nodes in a 1000 m^2 field with one speaker node and three groups of audiences, i.e., *audience1*, *audience2* and the *wanderers*. Both audience groups consisted of 20 members moving with speeds between 2-5 m/s. The movement of the audience groups was modeled using brownian motion and node movement was restricted to a limited area within the field. *Wanderers* consisted of 9 nodes who were capable of moving over the entire topology. The speeds for these nodes were randomly chosen between 1-5 m/s with pause times between 0-1 sec. The random waypoint model was used as the mobility model for all the wanderer nodes. The speaker node and 20 randomly chosen audience nodes acted as sources of data.

The second scenario was that of a disaster-rescue operation with a total of 75 nodes in a 2000 m^2 field. This scenario consisted of 2 helicopters, 2 rescue teams of foot soldiers and 2 teams on vehicles. The helicopters moved with speeds ranging between 0-50 m/s according to the random waypoint model. The first vehicle team consisted of 25 nodes while the second team consisted of 8 nodes. The members of both vehicle teams moved according to the random waypoint model with speeds ranging between 5-15 m/sec. The team of foot soldiers consisted of 20 nodes moving with speeds ranging between 0-5 m/s and pause times between 0-2 secs. Each team covered well-defined areas within the field with sufficient overlap to ensure that information could be relayed among the different teams. Two helicopters and 20 other nodes were randomly chosen nodes as data sources for this scenario.

We used *CBR* as well as *ON-OFF* traffic for both scenarios. In CBR, each source transmitted one pkt/s, while the traffic rate was set to 5 Kb/s for ON-OFF traffic with a burst period of 3 secs and idle time of 3 secs.

5 Results

We investigate both sides of the protocol reliability (i.e., delivery ratio) versus efficiency (overhead) trade-off. We compute packet delivery ratio as the ratio of total number of packets received by the nodes to the total number of packets transmitted times the number of receivers. Routing overhead is computed as the ratio between the number of control bytes to the number of data bytes received. In adaptive flooding, control overhead includes *hello messages*, retransmits in hyper flooding and all data header bytes forwarded by network nodes. In ODMRP, control bytes account for *Join Request* and *Join Table* packets. It also includes data packet header bytes forwarded by forwarding group members. In MAODV control bytes account for the *Rreq, Rrep, Mact, Hello* and *Grp_Hello* packets. It also includes the data packet headers forwarded by intermediate nodes.

5.1 Relative Velocity Based Switching

The graphs in Figure 1 show how protocol reliability varies with node mobility which is expressed in terms of average relative velocity. The average relative velocity is computed as follows. The relative velocity of each node with respect to its neighbors is calculated throughout the duration of the simulation and is then averaged over all nodes.

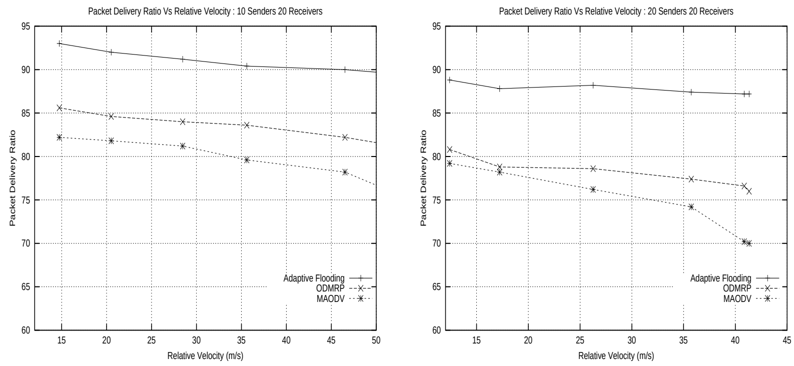

(a) 10 Senders 20 Receivers

(b) 20 Senders 20 Receivers

Figure 1: Packet Delivery Ratio as a function of Node Mobility

It can be observed from Figure 1 that packet delivery ratio decreases with increase in relative velocity. The increased mobility of the nodes causes them to move outside the radio range of their neighbors more frequently resulting in lower packet reception. In case of the adaptive flooding protocol, nodes rely on neighbor information to decide if they retransmit packets. Neighbor information may become stale as the mobility of the nodes increases resulting in lower packet delivery ratio at higher speeds. It can be seen from the graphs that adaptive flooding performs better than ODMRP or MAODV in terms of packet delivery ratios delivering around 90% of the packets at a relative velocity of 50 m/s. At lower speeds adaptive flooding switches to scoped flooding mode in an attempt to reduce redundant retransmissions. As the relative velocity increases it switches to flooding and hyper flooding modes resulting in consistent packet delivery ratios. Comparing adaptive flooding to ODMRP we notice that at lower speeds the difference in packet delivery ratio is only within 5%. However at higher speeds the difference in packet delivery ratio starts widening. For instance, in the case of 20 senders and 20 receivers we observe packet delivery ratio differences of up to 12% in favor of adaptive flooding. This is because with increased mobility the forwarding group members need to be updated more frequently. This requires that sources send out *Join-Requests* more frequently resulting in higher control overhead and greater packet loss due to contention.

Comparing ODMRP with MAODV we observe that ODMRP has better packet delivery ratios (around 7-10%) at higher speeds. Since ODMRP maintains meshes, it has multiple redundant paths to receivers and is not affected by mobility as greatly as MAODV. In the case of MAODV increased mobility causes frequent link changes and requires tree reconfiguration to prevent stale routing information. This in turn requires higher control traffic which can result in greater packet loss due to contention.

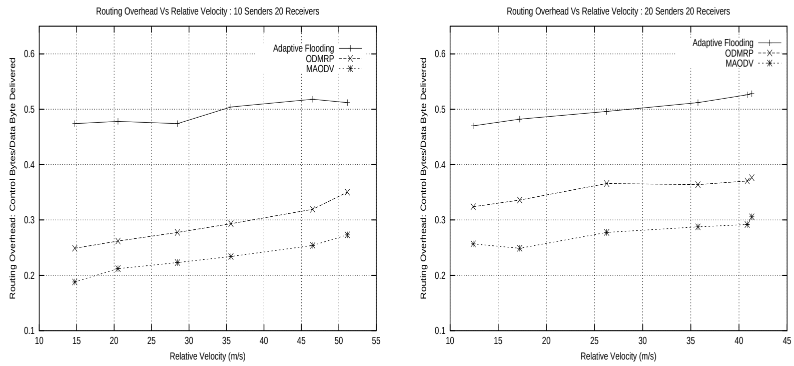

(a) 10 Senders 20 Receivers

(b) 20 Senders 20 Receivers

Figure 2: Routing Overhead as a function of Node Mobility

The graphs in Figure 2 plot control overhead per data byte delivered as a function of node mobility. It can be seen from the graphs that the routing overhead/data byte delivered increases with increase in node mobility. This is due to the fact that the routing overhead remains almost constant with increase in mobility but fewer data packets are delivered. Adaptive flooding has the highest overhead among the protocols on account of the redundant transmissions. At high velocities, adaptive flooding switches to the hyper flooding mode in an attempt to increase reliability, thereby increasing the routing overhead.

5.2 Network Load Based Switching

In this section we present results for adaptive flooding based on the network load switching criteria. The graphs in Figure 3 show how the reliability varies with network load. Although we ran simulations for different node velocities we only include results for node speeds of 20 m/s (72 km/hr). An ON-OFF traffic generator was used for the simulation results presented below. The overall traffic rate was obtained by averaging the data rate of all senders.

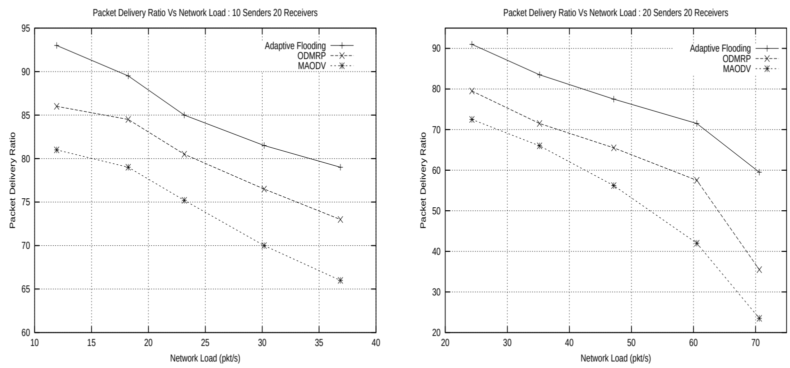

(a) 10 Senders 20 Receivers

(b) 20 Senders 20 Receivers

Figure 3: Packet Delivery as a function of Network Load

From Figure 3 it can be observed that for the 20 sender, 20 receiver case ODMRP and MAODV deliver around 40-60% packets at network load of 60 packets/sec. Both ODMRP and MAODV are affected to a greater extent, than adaptive flooding. As the network load increases, adaptive flooding switches to the scoped flooding mode in an attempt to reduce collisions. In case of adaptive flooding, the losses are mostly due to collisions. In case of ODMRP and MAODV, node mobility results in stale routing information. The bursty nature of the traffic causes a large number of packet drops before the routes are refreshed.

5.3 Conference and Rescue Scenarios

Tables 1, 2 and 3 present results for the conference and disaster rescue scenarios using relative velocity and network load as switching criteria.

Conference Scenario			
	Protocol	Delivery Ratio %	Routing Overhead (Bytes Xmit/Data byte recvd)
CBR Traffic	Adaptive Flooding (NL)	86.32	0.138
	ODMRP	81.38	0.136
	MAODV	79.87	0.081
ON-OFF Traffic	Adaptive Flooding (NL)	81.41	0.140
	ODMRP	67.78	0.112
	MAODV	64.58	0.087

Table 1: Conference Scenario: Network Load (NL)

Disaster Scenario			
	Protocol	Delivery Ratio %	Routing Overhead (Bytes Xmit/Data byte recvd)
CBR Traffic	Adaptive Flooding (RV)	82.47	0.170
	ODMRP	65.24	0.164
	MAODV	60.81	0.108
ON-OFF Traffic	Adaptive Flooding (RV)	76.79	0.180
	ODMRP	60.56	0.147
	MAODV	56.47	0.101

Table 2: Disaster Scenario: Relative Velocity (RV)

Disaster Scenario			
	Protocol	Delivery Ratio %	Routing Overhead (Byte Xmit/Data byte recvd)
CBR Traffic	Adaptive Flooding (NL)	81.73	0.150
	ODMRP	65.24	0.164
	MAODV	60.80	0.108
ON-OFF Traffic	Adaptive Flooding (NL)	76.44	0.150
	ODMRP	60.56	0.147
	MAODV	56.47	0.101

Table 3: Disaster Scenario: Network Load (NL)

For the conference scenario, adaptive flooding (using network load as the switching criterion) performs better than ODMRP and MAODV for both CBR and ON-OFF traffic. In particular, for ON-OFF traffic, adaptive flooding's delivery ratio is around 14% higher than ODMRP and around 17% higher than

MAODV. In this scenario, node density was sufficiently large and average node mobility was quite low. The low mobility of nodes coupled with high traffic load results in a large number of collisions. This triggers adaptive flooding to mostly operate in the scoped flooding mode in an effort to reduce re-transmissions. However, in case of ODMRP the number of forwarding group members was quite large, resulting in a large number of redundant transmissions. This effect is compounded in the case of bursty traffic resulting in lower packet delivery ratios for ON-OFF traffic. Adaptive flooding's routing overhead is comparable to ODMRP in case of CBR traffic and slightly higher than ODMRP in case of ON-OFF traffic.

For the disaster scenario, both versions of adaptive flooding performed considerably better, delivering around 16-22% more data packets than ODMRP or MAODV. This scenario consisted of several groups of nodes which were restricted to move within a subset of the total topology. The groups had sufficient overlap to ensure that data packets could be relayed from one group to another. In case of ODMRP only forwarding group members can relay data, whereas in MAODV only multicast tree members can forward data traffic. At the time of route setup, nodes in the overlap region are incorporated as forwarding group members (ODMRP) or multicast tree members (MAODV). However node mobility causes the forwarding group members and multicast tree members to move outside the overlap region resulting in a large number of packet drops until the route is refreshed at the end of the *Active_Route_Interval*. This effect is more severe for bursty traffic as compared to CBR traffic. However in the case of adaptive flooding all nodes can forward data traffic and thus adaptive flooding delivers around 20% more data than ODMRP or MAODV. Applications that require high delivery guarantees will likely trade adaptive flooding's slightly higher overhead for its considerably higher delivery rate.

6 Conclusions

The diverse nature of MANETs make it impossible for any one protocol to be be optimal under all scenarios and operating conditions. This calls for specialized multicast solutions for each type of network and the means for integrating those solutions. To this end we have proposed an adaptive approach to routing where nodes dynamically switch routing mechanisms based on their perception of current network conditions. This paper investigated the performance of an adaptive approach to group communication that tries to achieve reliable delivery for a wide range of MANET scenarios, including high mobility and traffic load. Using the proposed adaptive protocol, which incorporates different variations of flooding, nodes can switch from one mode of flooding to another using relative velocity and traffic load as switching criteria.

We reported simulation results comparing the adaptive protocol with ODMRP and MAODV for "concrete" MANET scenarios. The results demonstrate that the adaptive protocol performs consistently well in terms of both packet delivery ratios and routing overhead. For the disaster-rescue scenario, adaptive flooding's delivery ratio was about 15-20 % higher than ODMRP and MAODV for CBR and ON-OFF traffic, which was achieved at a comparable routing overhead. In case of the conference scenario, adaptive flooding's delivery ratio was 15-17 % higher than ODMRP and MAODV for ON-OFF traffic. The routing overhead in this case was comparable to ODMRP and slightly higher than MAODV.

We should point out that given the diversity of MANETs, adaptive protocols are capable of providing consistent performance benefits over a wide range of operating conditions. Our simulation results highlight these performance benefits and lay the foundations for other adaptive routing mechanisms which are not based on flooding.

References

1. S.J. Lee, W. Su, J. Hsu, M. Gerla, and R. Bagrodia. A performance comparison study of ad hoc wireless multicast protocols. *Proceedings of the IEEE Infocom 2000*, March 2000.
2. E. Royer and C. Perkins. Multicast operation of the ad-hoc on-demand distance vector routing protocol. *Proceedings of the ACM Mobicom'99*, pages 207–218, August 1999.
3. Katia Obrazcka, Gene Tsudik, and Kumar Viswanath. Pushing the limits of multicast in ad hoc networks. *Proceedings of the 21st. IEEE ICDCS*, pages 719–722, April 2001.
4. D. Estrin et. al. Pim border router specification for connecting pim-sm domains to a dvmrp backbone. Internet draft draft-ietf-idmr-PIM-SM-spec-09.ps, October 1996.
5. D. Thaler. Interoperability rules for multicast routing protocols. Internet draft draft-thaler-interop-00.ps, November 1996.
6. D. Meyer. Some issues for an inter-domain multicast routing protocols. Internet draft draft-ietf-mboned-issues-03.txt, November 1997.
7. S. Ni, Y. Tseng, Y. Chen, and J. Sheu. The broadcast storm problem in a mobile ad hoc network. *Proceedings of IEEE/ACM MOBICOM'99*, pages 151–162.
8. M. Gerla and S.J. Lee. On-demand multicast routing protocol for mobile ad-hoc networks. Available from http://www.cs.ucla.edu/NRL/wireless/.
9. Li Quiming. Scenario generator for manets. Available from http://www.comp.nus.edu.sg/ liqiming/fyp/scengen/index.html, April 2001.