PERFORMANCE OF MOBILE AD HOC NETWORKING ROUTING PROTOCOLS IN LARGE SCALE SCENARIOS

Julian Hsu, Sameer Bhatia, Ken Tang, Rajive Bagrodia
Scalable Network Technologies, Inc.,
Culver City, CA

Michael J. Acriche
US ARMY RDECOM CERDEC STCD

ABSTRACT

As wireless ad hoc networks grow in size and complexity, the need to study and comprehend the scalability and behavior of these systems and their protocol components becomes essential. This paper presents a comprehensive study on the performance of common MANET (mobile ad hoc network) routing protocols in large-scale networks. The routing protocols used in this study include AODV, DSR, OLSR, and ZRP, which represents a good mix of reactive, proactive and hybrid protocols.

INTRODUCTION

Mobile ad hoc networks are networks without a fixed infrastructure. Communications must be set up and maintained on the fly over mostly wireless links. Each node of an ad hoc network can both route and forward data. The exploding demand for computing and communication on the move has led to reliance for ad hoc networks.

Although substantial research efforts have been devoted towards the design and development of ad hoc routing protocols, there is relatively little understanding of their behavior as the system is scaled up. First, it is unclear that any existing ad hoc routing protocol can be used to successfully route data over the many hops that will be necessary in large scale wireless networks. Second, it is unknown which of the preceding techniques is likely to perform better for different types of traffic or mobility patterns as network size grows. This lack of understanding comes from the fact that previous simulation tools were not able to facilitate the analysis of large scale networks without compromising the fidelity, and thus the accuracy, of the ad hoc network models.

Our definition of scalability is that the routing protocol must provide applications with high throughput and low end-to-end delay, relative to the network scenario. Furthermore, the routing protocol itself must use minimal routing control overhead and exhibit high packet delivery ratio (which correlates to high application throughput). Finally, for a routing protocol to be considered completely scalable, these characteristics must hold true regardless of the network size, node density, number of hops, mobility, number of sources and destinations, and network load, and degrade gracefully if the network configuration has insufficient resources.

In this paper, we analyze the scalability of several unicast ad hoc routing protocols, including Ad Hoc On-Demand Distance Vector (ADOV), Dynamic Source Routing (DSR), Optimized Link State Routing (OLSR), and Zone Routing Protocol (ZRP), through simulation. We detail our general methodology for ascertaining the scalability of an ad hoc protocol and discuss the metrics used for our case studies. We then provide simulation results and analysis for the ad hoc unicast routing protocols under study.

A GENERAL METHODOLOGY FOR ROUTING SCALABILITY TESTING

Efficient, Scalable Network Simulation: Past simulation tools lack the ability to simulate large scale networks in an accurate manner. This is due to the fact that past simulation tools require an immense amount of memory and runtime that make such studies impractical. QualNet, on the other hand, is a scalable network simulation library that was designed with the primary goal of simulating large, high-fidelity models of wired, wireless, and mixed networks in an efficient manner. It was designed to achieve modular design for easy comparison of protocols under uniform conditions, detailed and accurate models, efficient execution, and transparent parallel execution for further scalability and runtime efficiency. These features make QualNet an excellent candidate for our study.

DESIGNING THE EXPERIMENTS

Our methodology for assessing the scalability of a routing protocol is to start with a carefully designed base configuration and network scenario for all experiments, and to vary one control parameter at a time to stress the network in different directions. Careful design of these control parameters enables us to assess and isolate the effect of each of the variables, such as mobility and network size, on the performance of the routing protocol for our desired metrics. In addition, design of the base condition, network topology, and traffic must take into account the real networks for which the results should be applicable.

Once simulation results are available, analyzing the results to determine the scalability of a routing protocol requires two distinct phases: one qualitative and the other quantitative as recommended in [12]. First, the qualitative phase consists of explaining why the results occurred for a given protocol. Second, one can develop a quantitative criterion for scalability and apply it as follows.

Given any metric, such as control overhead or throughput, to assess if the protocol is scalable according to the metric and criteria variable, such as mobility or network size:

1. Establish an upper and lower bound for the metric for the criteria variable depending on knowledge of network needs or existing literature. Example: a unicast routing protocol is considered scalable in terms of mobility if the control overhead remains between 20% and 50% for all experiments.
2. For each criteria variable, consider the metric to be a random variable in each experiment to be run. Example: For each mobility value, control overhead is the random variable to monitor.
3. Correspondingly, consider each experiment across all values for the criteria variable a random process indexed by the value for the criteria variable. Example: After running a set of experiments with varying mobility, there will be a value for the random variable of control overhead for 2 m/s, 9 m/s, 16 m/s etc.
4. Examine the values of the random variable for each value of the criteria variable. Find the expectation of the random variables for each value of the criteria variable. If the expectations all fall within the bounds, then the protocol is scalable.

As a quantitative approach, this method requires no specific knowledge of the protocols, but only computations on experimental result data. It can easily be catered to whatever variables are needed and to the specific performance requirements of metrics that are needed.

EXPERIMENTS

We have proposed a standard set of experiments as part of the methodology. The experiments are designed to ensure that the control values for the parameters of network size, node density, number of hops, mobility, number of sources and destinations, and network load can be assumed constant and as specified in this section unless the parameter's effect on the scalability of the MANET protocol is being investigated. For this reason, nodes are stationary, with the exception, of course, being the mobility experiments. The control values for parameters in experimental groups are:

Table1: Control Values for parameters for each experimental group

Parameters/ Group	**Size**	**Density**	**Hops**	**Load**	**Mobility**	**Sources**
Size (nodes)	varies	50	varies	50	50	50
Density (m/node)	253	varies	253	253	253	253
Max. Hops	10	10	varies	10	10	10
Load (bytes/s)	1460/src	23360	17520	varies	23360	17520
Mobility (m/s)	0	0	0	0	varies	0
Sources	1/3'rd	16	12	16	16	varies

For all experiments in this study, we use Constant Bit Rate (CBR) application traffic using UDP as the transport layer. We have experimented with AODV, DSR, OLSR, and ZRP as the wireless unicast routing protocols of interest. At the MAC layer, the IEEE 802.11 DCF protocol is utilized, with the IEEE 802.11b radio device with a maximum data rate of 2Mbps. A two-ray propagation path loss is modeled for the propagation where the free space path loss is used for near sight and plane earth path loss is used for far sight. The radio range is approximately 375.

Each experiment occurs within a square terrain dimension. Node placement within the topology is always random and uniform, meaning that the terrain is split into square cells and a node is randomly placed within each cell. The controlled node density was arbitrarily chosen to be 253 meters squared per node because this value is in the middle of the IEEE 802.11b radio range of 376.782 meters and does not permit all nodes in areas this size to communicate with its neighbors. The exception, of course, is for the node density group of experiments for which this value varies. The control values for all parameters in all experimental groups are summarized in table 1.

All experiments were run using the same set of 3 different random seeds for the simulator. The set of CBR applications for each set of experiments were chosen by randomly selecting the set of sources and destinations from the available nodes using 3 different random seeds. For each variable value for each experimental group, except the hops group described later, 9 experiments were run using the varying simulator and application generating seeds. The results are averaged over the number of multiple runs for each variable value in an attempt to reduce randomness. Consequentially, each point of the graphs of the results is the result of 9 separate simulation runs.

PROTOCOL PERFORMANCE METRICS

The metrics used to analyze protocol performance are mean throughput per node, control overhead, end-to-end delay or latency, and packet delivery ratio of the routing protocol. While all of these were collected for each experiment set, only select graphs are presented here due to space limitations.

Examining throughput, especially when it is considered relative to different network scenarios, helps determine how well the routing protocols permit applications to optimize the use of the available bandwidth given the aforementioned limitations. The latency gives a measure of the time to traverse the path and the time cost of the protocol's route discovery. The packet delivery ratio, the ratio of the number of packets actually received over the number of packets that are supposed to be received, quantifies how well applications are able to perform given the routing overhead. The packet delivery ratio, in conjunction with latency, assesses how efficiently best effort traffic performs. Finally, the control overhead, the ratio of the number of control packets processed divided by the total number of data and control packets processed, measures the efficiency of the routing protocol. The control overhead not only gives a metric for the amount of bandwidth available to data packets, but also indicates whether the latency and packet delivery ratio are compromised because of the network congestion and interference generated from control packets [2].

ROUTING PROTOCOL SCALABILITY

NETWORK SIZE

The network size experiments are designed to isolate the effects of adding additional nodes to the network on the routing protocols tested. We increased the network from a small 9 node network to 100, 225, 529, and finally a 1024 node network, keeping the node density fixed, and the traffic and simulation length proportional to the network size.

The following table summarizes the network size experiment parameter settings:

Table2: Network size experiment parameter settings

Network Sizes	10, 100, 225, 529, and 1024 nodes
Node Placement	Uniform density (avg. 1 node / 253m^2), random placement
Mobility	None
Traffic	1/3rd of the network randomly selected sources with randomly selected destinations
Simulation Time	Proportional to the number of nodes, varied from 160 seconds to 1850 seconds
Stabilized Application Load	Proportional to the number of nodes, varied from 4380 bytes/sec

The network size vs. throughput graph in figure 1 plots the per-node average of application level observations of bps data received. As you can see, DSR appears to scale best in response to increasing network size with the remaining protocols having much less success with this particular configuration.

OLSR shows a uniform, but fairly sharp decline. One of the issues with link state protocols is that, if the protocol is unable to converge, there are large disconnects in the known network topology and many packets are dropped due to lack of sufficient routing information. AODV, strangely, shows a steeper decline than DSR indicating potential lack of scalability in this direction and scenario. This behavior is likely to be explained by one or more of the optimization differences between DSR and AODV. ZRP has a very sharp decline between 9 and 200 nodes, before leveling off at less than half the throughput of DSR and AODV. This result is surprising, but given the parameter sensitivity of ZRP, much more investigation will need to be conducted to determine the cause. We used the default protocol designer's recommendations for the timer and zone size variables.

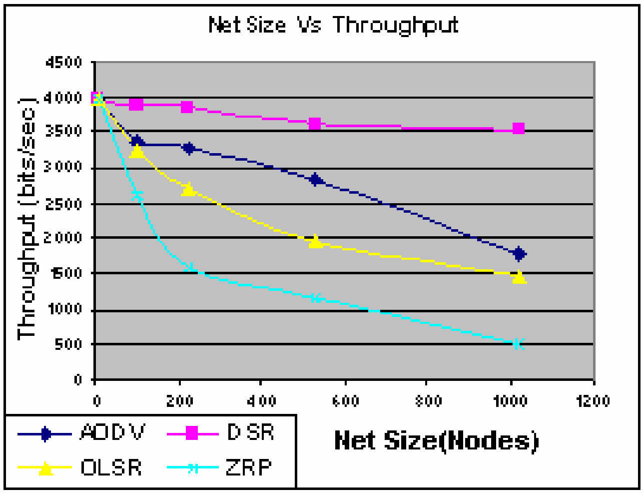

Figure 1: Throughput for Network Size

NODE DENSITY

Wireless routing protocols also have to scale to node density. These experiments analyze the effect of increasing node density by using a grid topology with fixed number of nodes, fixed number and identity of senders and receivers, and varying the distance between neighboring nodes in the grid. The following figure summarizes the node density experiment parameter settings:

Table3: Node density experiment parameter settings

Network Size	50 nodes
Node Placements	Grid placement: • Sparse: 1 node / 347m^2 • Moderate: 1 node / 185m^2 • Dense: 1 node / 132m^2
Mobility	None
Traffic	1/3rd of the network randomly selected sources with randomly selected destinations
Simulation Time	225 seconds
Stabilized Application Load	23360 bytes/sec

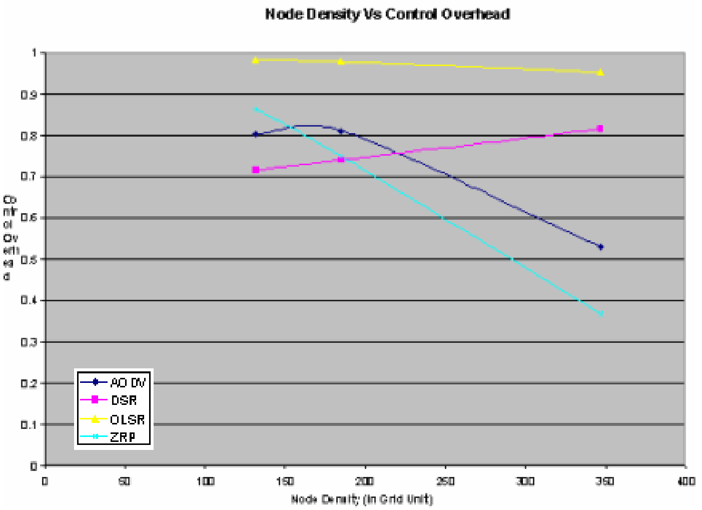

Figure 2**:** Node Density Vs Control Overhead

Figure 2 shows the Control Overhead curve for the Node Density experiments. The control overhead measurements have been normalized to enable easier protocol comparison. It assesses the efficiency of a protocol in terms of available bandwidth for data packets and the effect of control packets on latency and the packet delivery ratio.

DSR shows increased ratio in control overhead for the sparse networks. This is due to sparse networks having higher path lengths, and therefore more rebroadcasts of route requests, and more route reply packets. AODV shows a different trend. It is using fewer control packets as the network gets more sparse. ZRP shows a similar trend, indicating that these protocols have difficulty dealing with a network with few neighbors. The original route acquisition process depends on neighboring nodes overhearing and rebroadcasting route requests, and if route requests are lost, the entire process stalls. The hidden terminal problem can contribute to route request losses, and is more prevalent in sparse networks. This behavior is being observed here. OLSR, on the other hand, has fairly uniform control overhead, as expected from a proactive protocol. It trends downwards with sparse networks because there are fewer links to report. But since there are fewer links, route convergence takes longer, which was seen the latency graphs for this experiment set.

NUMBER OF HOPS

The performance of routing protocols should also be robust to the number of hops between the source and destination when handling route discovery. Routes that span multiple hops often increase the amount of routing protocol control traffic because routing packets must traverse numerous hops in order for the nodes to discover the routes. Furthermore, more than one path often exists for routes spanning multiple hops. Therefore, routing algorithms must also consider the control overhead required to discover these routes to select the appropriate paths from the multiple paths discovered. These experiments analyze the effect of increasing the distance between sources and destinations

The following figure summarizes the number of hops experiment parameter settings:

Table4: Number of hops experiment parameter settings

Network Size	121 nodes
Node Placements	Grid placement: 1 node / 252m^2
Mobility	None
Traffic	Approx 1/4^{rth} nodes randomly selected as sources with randomly selected destinations at a given hop count.
Simulation Time	225 seconds
Stabilized Application Load	17520 bytes/sec (802.11b)

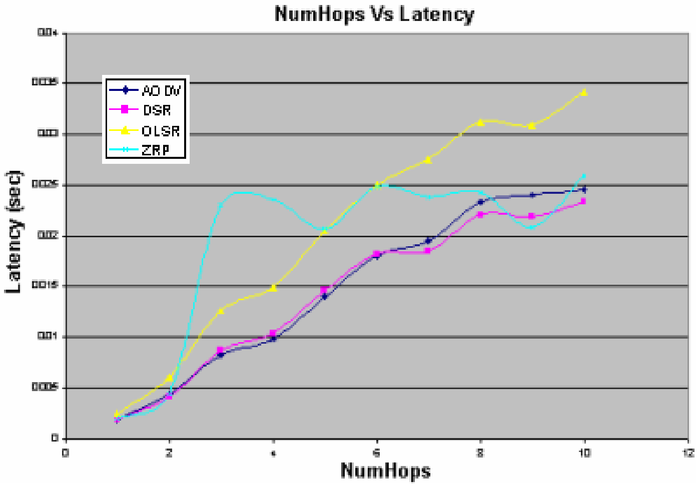

Figure 3: NumHops Vs Latency

Due to space constraint we are only enclosing the latency results for this experiment set(figure 3). The num hops vs. latency graph shows the relationship between the number of hops, and the amount of time it takes to acquire a route. The latency for OLSR has the highest values from 1 to 10 hops, and generally the highest slope. This indicates that OLSR has difficulty scaling to hop count in this scenario. As a proactive protocol, we would expect OLSR to have lower average latency than a reactive AODV or DSR, but the results for this set of experiments indicate otherwise. For OLSR to lose its innate advantage in latency, network route convergence would have to be slower than route acquisition, and given the high control overhead data that we collected for this experiment set, it is easy to see that this is the case. ZRP, on the other hand, is a tale of two graphs. In essence, the sharp line connecting the data point at 2 hops to the data point at 3 hops is misleading, because this is a break point in the algorithm rather than a continuous record of protocol behavior. At the 1 hop and 2 hop data points, ZRP actually has better latency than OLSR. This is logical, given that the proactive zone of interest is much smaller. At 3 hops and beyond, this result is indicative of the interzone routing, and it shows a fairly flat graph from 3 to 10 hops, with some oscillation caused by random number seeds not being completely filtered out. This would be a good result for scalability if only the protocol could pump up the throughput and PDR, although it may be a good result for precisely this reason: it is delivering the low hanging fruit.

MOBILITY

Routing protocols in the wireless ad hoc domain must also be scalable with respect to mobility. When mobility is introduced into the equation, link breaks and route rediscovery occur frequently. Wireless routing protocols thus need to account for link breakages and re-discoveries in their design. The mobility experiments are actually the ones least suited to this paper's form of variable isolation. Ideally, it would be best to combine mobility with each of the other axes in the scalability study, since it is likely that the effect will be more exponential than additive, and may even alter the perception of scalability as mobility increases. For example, if a protocol in this study has high control overhead, but is generally successful in delivering packets, it is quite possible that the additional overhead required for mobility updates would cause a chain reaction that overwhelmed the protocol's performance. The following figure summarizes the mobility experiment parameter settings:

Table5: Mobility experiment parameter settings

Network Size	50 nodes
Node Placements	Uniform density (avg. 1 node / 253m^2), random placement
Mobility	Random waypoint mobility, constant speed of 2m/s, 9m/s, 16m/s, 20m/s, and 30m/s with a 30 second pause when it reaches each randomly selected destination before choosing the next one.
Traffic	1/3rd nodes randomly selected as sources with randomly selected destinations at a given distance in hop count from the source
Simulation Time	225 seconds
Stabilized Application Load	17520 bytes/sec

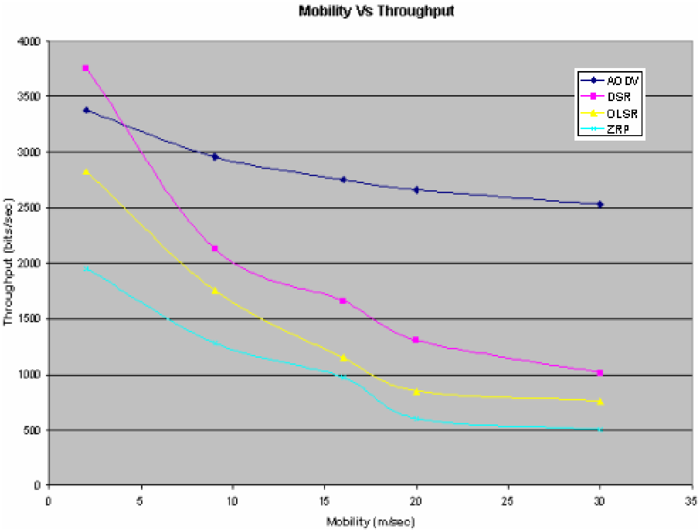

Figure 4: Mobility Vs Throughput

Figure 4 represents the mobility versus throughput data that we collected for this experiment set. It displays 802.11b throughput for mobility experiment set. AODV is the clear favorite in this roundup. DSR starts out with higher throughput in the lowest mobility case, but DSR optimizations seem less able to handle high mobility, but it still manages a second place finish. The only area of concern appears to be that, as car speed is approached, the rate of throughput decline seems to be accelerating more than that of OLSR, the third place finisher. OLSR is somewhat less scalable than DSR, but follows a roughly similar curve of decline. We do expect that link state protocols generally would have a tougher time dealing with random mobility, than with a group mobility that would cause fewer topology changes. Random mobility means that the entire network is in a constant state of flux, and it is actually arguable whether this represents an edge case, with much less of a change to be expected in the operational cases. ZRP is an interesting case. We notice that both ZRP and DSR have a slight bump upwards in throughput at the 16 m/s case, and it is possible that AODV has a flattening in this area for reactive protocols: OLSR does not show this behavior in this experiment set.

NUMBER OF SOURCES AND DESTINATIONS

The number of sources and destinations in the network has an effect on the routing algorithm. Large number of sources and destinations might require some routing protocols to maintain individual routes to most of the nodes in the network. Conversely, small numbers of sources and destinations only require the routing protocol to maintain fewer routes, even though the protocol may be maintaining routes to every destination. To evaluate the performance of routing protocols in this scenario, we varied the CBR sources between one-sixth that of the total network size to equal to the entire network size. The results from this experiment set are not included due to space constraints, however all protocols performed similarly and reasonably in this experiment set with the worst performer only being about 20kbps apart from the best in the throughput measurements.

NETWORK LOAD

Network load also plays a big part of the scalability of routing protocols. As an example, when the network load is high, the network load may affect the routing control packets and impede their delivery by competing for access to the channel, and increasing likelihood of collisions. Thus, route information may be slow to acquire or stabilize. In order to analyze the effect of network load, we increased the frequency of data packet transmissions from 0.2 packets per second (senders send 1 packet every 5 seconds), to 200 packets per second (senders send 1 packet every 5 milliseconds). Figure 5 is the network load versus PDR result for this experiment set.

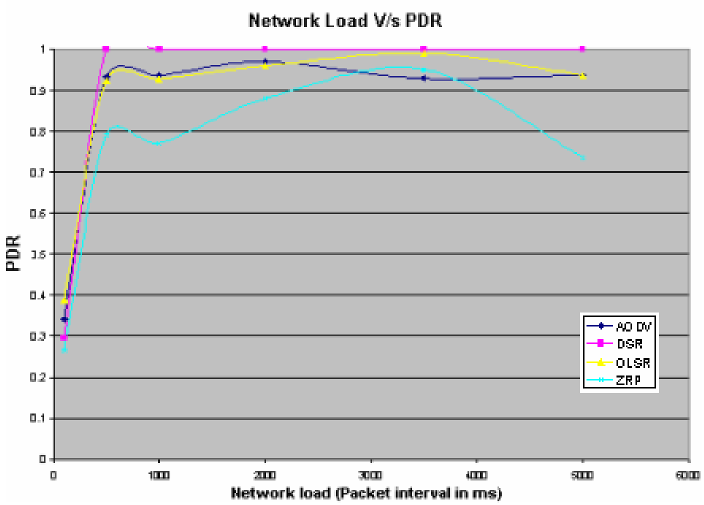

Figure 5: Network Load Vs PDR

DSR performed close to 100% all cases except for the last case where it declines. This last case illustrates protocol performance under load that is more than the physical medium can support. OLSR delivers more than 92% of packets in all cases except the last one, and delivers a class leading 38.9% there. The choice then is between the perfect delivery of DSR in adequately provisioned cases, and the increased robustness to high traffic that OLSR provides. AODV manages to slightly outperform OLSR on three of the cases, but trails by a wider margin near the high and low ends of the packet interval spectrum. ZRP trails overall, and is not well suited to this particular experiment set.

CONCLUSION

In summary, we have presented a methodology to assess the scalability of an ad hoc routing protocol. Our approach enables us to isolate the effect of a single variable on the scalability of a protocol by using carefully, designed control configurations for our simulations using QualNet. This method also insists on both quantitative and qualitative analysis of our simulation results. In our case studies, our qualitative approach details why each protocol performs differently relative to the specific values of criteria variables. Our methodology is flexible to any performance bounds, metrics, and criteria variables. Specific knowledge of desired network configurations and applications can motivate all performance bounds, metrics, and criteria variables to be used.

REFERENCES

[1] J. Broch, D. Maltz, D. Johnson, Y. Hu and J. Jetcheva. A Performance Comparison of Multi-hop Wireless Ad Hoc Network Routing Protocols. In *Proceedings of MOBICOM'98*, 1998.

[2] S. R. Das, C.E. Perkins, E. M. Royer, and M. K. Marina. Performance Comparison of Two On-demand Routing Protocols for Ad Hoc Networks. In *Proceedings of the IEEE INFOCOM 2000*, pages 3-12, Tel Aviv, Israel, March 2000. IEEE.

[3] A. Iwata, C.-C. Chiang, G. Pei, M. Gerla, and T.-W. Chen, "*Scalable Routing Strategies for Ad Hoc Wireless Networks*," IEEE Journal on Selected Areas in Communications, special issue on Wireless Ad Hoc Networks, vol. 17, no. 8, August 1999, pp. 1369-1379.

[4] C.E. Perkins and E. M. Royer. Ad hoc on-demand distance vector routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages, 90-100, Feb 1999.

[5] C.E. Perkins, E. M. Royer, and S. R. Das. Ad hoc on-demand distance vector (AODV) routing. http://www.ietf.org/internet-drafts/draft-ietf-manet-aodv-07.txt, November 2000. IETF Internet Draft (work in progress).

[6] E. M. Royer, P. M. Melliar-Smith, and L. E. Moser. An Analysis of the Optimum Node Density for Ad hoc Mobile Networks. To appear in *Proceedings of the IEEE International Conference on Communications*, Helsinki, Finland, June 2001.

[7] S. Bae, S. Lee, W. Su, and M. Gerla. The Design, Implementation, and Performance Evaluation of the On-Demand Multicast Routing Protocol in Multihop Wireless Networks. http://www.cs.ucla.edu/NRL

[8] S. Bae, S. Lee, and M. Gerla. Unicast Performance Analysis of the ODMRP in Mobile Ad hoc Network Testbed. http://www.cs.ucla.edu/NRL

[9] S.-J. Lee, W. Su, and M. Gerla, "On-*Demand Multicast Routing Protocol (ODMRP) for Ad Hoc Networks*," Internet Draft, draft-ietf-manet-odmrp-02.txt, January 2000, Work in progress.

[10] S. Lee, W. Su, J. Hsu, M. Gerla, and R. Bagrodia. A Performance Comparison Study of Ad Hoc Wireless Multicast Protocols. http://www.cs.ucla.edu/NRL/

[11] S.-J. Lee, W. Su, and M. Gerla. On-Demand Multicast Routing Protocol in Multihop Wireless Mobile Networks. http://www.cs.ucla.edu/NRL/

[12] M. S. Corson and J. Macker, “MobileAd hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations,” Requests for Comments 2501, Internet Engineering Task Force, Jan. 1999.