



NRL/MR/5523--00-8521

Protocol Considerations For Distributed Sensor Networks

VINCENT D. PARK
JOSEPH P. MACKER

*Protocol Engineering and Advanced Networking (Protean) Research Group
Communication Systems Branch
Information Technology Division*

December 18, 2000

20010129 021

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE December 18, 2000	3. REPORT TYPE AND DATES COVERED Final	
4. TITLE AND SUBTITLE Protocol Considerations For Distributed Sensor Networks		5. FUNDING NUMBERS PE - 0601153N TA - 02113	
6. AUTHOR(S) Vincent D. Park and Joseph P. Macker		8. PERFORMING ORGANIZATION REPORT NUMBER NRL/MR/5523--00-8521	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory Washington, DC 20375-5320		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research 800 N. Quincy Street Arlington, VA 22217		11. SUPPLEMENTARY NOTES This work was supported by the Surveillance, Communications, and Electronic Combat Division (Code 313) of the Office of Naval Research under contract number N0001401WR20010.	
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) In this report, we consider the problem of developing large-scale distributed sensor networks. We briefly describe the motivation behind the development of such networks and then consider the expected characteristics of the networking environment. We define a concept of operations for sensor network applications that allows expected characteristics to be exploited for maximizing the scalability of networking solutions. Finally, we discuss the available and emerging networking protocols, provide insight into their applicability for supporting distributed sensor networks, and identify areas for further research.			
14. SUBJECT TERMS Communication Sensor Protocol Network Wireless		15. NUMBER OF PAGES 11	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED		16. PRICE CODE	
18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	
20. LIMITATION OF ABSTRACT UL			

PROTOCOL CONSIDERATIONS FOR DISTRIBUTED SENSOR NETWORKS

INTRODUCTION

Recently, there has been considerable research interest in the development of large-scale, distributed sensor networks. This interest has been fueled in part by the anticipated future availability of small, inexpensive, low-power network computing and wireless communication devices. Conceptually, the combination of such devices and traditional sensor technology provide the physical components needed to form the foundation of the envisaged sensor networks. A small sensor network node with sensing, processing, storage and communication capabilities could either be completely self-contained or integrated into other equipment, and a large number of such nodes could be deployed in an area of interest and networked together to collect information, Figure 1.

Ideally, deployment of sensor networks should be easy, rapid and flexible. The networking protocols should be unassuming and designed to accommodate virtually any method of deployment. Once deployed, the network would provide the capability to detect, track, and report on data of interest. Sensor networks have wide potential use for both military and commercial applications. In a military context, distributed sensor networks could be used to gather information regarding potential threats—such as enemy personnel or vehicles, biological or chemical agents, or radioactive fallout. Sensor networks could also be used to monitor environmental conditions, gather information in high risk areas, or maintain physical grounds security. In each case, a similar commercial application can also be envisaged. In addition to relatively direct applications such as the examples above, the development of this technology will likely lead to other novel uses and currently unforeseeable applications.

A desire to provide coverage for a given area coupled with the need for communication efficiency and power conservation may push the limits of network scalability. This creates challenges both in the design of suitable networking protocols and in network configuration and management. While recent research on the topic of mobile ad hoc networks may have some relevance, the applicability of proposed protocols may be limited due to the differences in the salient characteristics of the respective networking environment. Additionally, to achieve the desired network scale and ease of deployment, auto-configuration and self-organization techniques will be of paramount importance. Currently, this remains a largely unexplored research topic.

SENSOR NETWORK ATTRIBUTES

Certainly, distributed sensor networks and prior notions of mobile ad hoc networks [CM 99] have much in common, but there are also some significant differences. Prior work on mobile ad hoc networks should be considered and applied to sensor networks

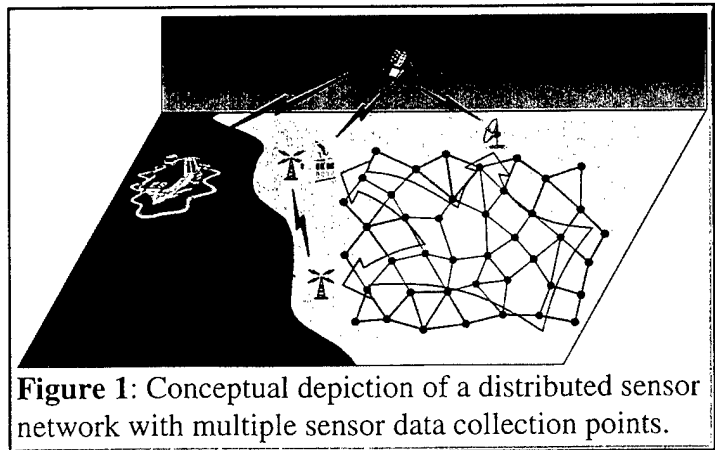


Figure 1: Conceptual depiction of a distributed sensor network with multiple sensor data collection points.

where appropriate. However, given the challenging nature of the networking environment, it is important to recognize the unique aspects of problem and identify or develop approaches that are not only compatible with respect to any underlying assumptions but also which potentially exploit any expected characteristics.

Distributed sensor networks and mobile ad hoc networks are both examples of multihop, wireless networks. In either case, the network nodes are assumed to be equipped with wireless receiver/transmitters and the status of the communication links between the nodes at any given time is a function of their positions, transmitter power levels, antenna patterns, interference levels, and other time-varying factors. Wireless links inherently have significantly lower capacity than their hardwired counterparts. This *bandwidth-constrained* nature of wireless networks may dramatically alter the trade-offs to be made in protocol design. Additionally, the nodes in a wireless network are often assumed to have limited available power. In such cases, the *power-constrained* nature of the nodes is likely to also be a significant factor in network and protocol design. Together these constraints emphasizes the need for communication-efficient and energy-efficient techniques.

Network Size and Connectivity

The expected size and connectivity of a network can have a significant impact on the suitability of various networking protocols. For example, the scalability of a specific routing protocol may be primarily limited by the number of network nodes, the diameter of the network, or the nodal degree (i.e., the number of neighbors a node has). Thus, characterizing the expected network size and connectivity can provide insight into the suitability of different protocols. The size and connectivity of a sensor network may be determined, in part, by the following factors:

- Area to be covered by sensor network
- Sensing capability/range of network nodes
- Cost of network nodes
- Transmission range of network nodes
- Power constraints of network nodes

Based on these factors, sensor networks may potentially comprise a very large number of nodes. It is fair to assume that the desire to build a sensor network, implies that the sensing capability/range of a single node is less than the area to be covered. Thus, some density of sensor nodes in the area is required to collect the desired information in a distributed fashion. By our initial assumptions, we also expect the nodes to be inexpensive—thus, allowing for the potential use of a large number of nodes to increase the sensor sampling density within the area to be covered and making them relatively expendable (i.e., throwaway or leave-behind devices).

The power conservation requirements of the nodes, advantages of spectrum reuse, and the potential need for low probability of detection favor minimizing the transmission power levels used for communication between nodes. Thus, direct connectivity between a node and other nodes in the network may be limited to a small set of geographically close neighboring nodes. This suggests that the expected nodal degree may be much less than the number of nodes in the network—resulting in a mesh-like network topology with a relatively large network diameter. It is logical to conclude that protocol scalability with respect to the number of network nodes and the diameter of the network are more critical than with respect to the nodal degree.

Topology Dynamics

Connectivity changes between nodes in a wireless network are likely to be more frequent than in hardwired networks simply due to the time-varying nature of wireless communications. However, for sensor network applications there are several other factors that may have a more significant effect on the dynamics of the network topology.

- Mobility of network nodes
- Failure/replacement of network nodes
- Transmission power control of network nodes
- Power-conservation/sleep-mode of network nodes

Depending on the specific design and application of a sensor network, any combination of the above attributes may be present. Node mobility may be used as a means of deploying the sensor nodes, adapting the coverage after deployment, more thoroughly sampling the area to be covered, or avoiding detection/capture of sensor nodes. Given that connectivity between nodes is determined in part by their positions, node mobility can significantly increase the frequency of network topological changes.

If node failure/replacement is frequent, the network topology may be highly dynamic even when nodes are essentially static. This may be more significant for military applications, where sensor nodes may be intentionally removed or destroyed by unfriendly forces. Node replacement at an equivalent rate may be required to maintain a desired level of performance.

Dynamic control of transmission power levels and power conservation techniques, such as allowing nodes to enter a hibernation or sleep mode, are areas of ongoing research. However, their obvious applicability to sensor networks warrants consideration of their effect on other networking protocols. It should be clear that these mechanisms can also increase the dynamics of the network topology.

Collectively, these factors indicate that the topology of sensor networks may be moderately to highly dynamic even if not mobile. Such dynamics emphasize the need for protocols that are robust, highly adaptive, and efficient in their reactions to topological changes.

Traffic Patterns

The traffic patterns in distributed sensor networks may be somewhat unique—allowing for the design of networking approaches and protocols that are tailored for efficient operation under the expected conditions. In the following discussion we attempt to categorize types of traffic and consider the potential sources and destinations of the different traffic types. Based on a concept of operation where the sensor data being gathered by the nodes in the network is forwarded to a small set of nodes for processing, analysis, or fusion (or to a small set of gateway nodes that provide access to an external network), we define the following categories of multihop traffic:

- Sensor Data
- Sensor Network Control
- Sensor Node Control
- Sensor Node Feedback

While other concepts of operation and traffic categorizations are possible, the above seems to be an intuitive and logical approach.

We define *sensor data* as traffic being generated by the sensor nodes and forwarded to the fusion (or gateway) nodes. This traffic has a potentially large number of sources but a small number of destinations. Depending on the specific sensor network application, this traffic may be periodically generated by all sensor nodes, or asynchronously generated by sensor nodes upon specific triggering events.

By *sensor network control* we refer to traffic used for modifying or controlling the operation of the entire sensor network as a whole. For example, enabling/disabling data collection or adjusting the rate of periodic reporting. There would likely be only a small number of sources for this type of traffic (perhaps the fusion nodes or external network nodes), but the traffic would, by definition, be destined for all sensor nodes.

There may also be a need to send control information to specific individual sensor nodes. This *sensor node control* traffic would presumably be generated by the same nodes that generate the sensor network control traffic, but would be destined for specific sensor nodes. Some sensor network applications may not require support for this type of traffic and those which do, may only need it infrequently.

Finally, *sensor node feedback* refers to traffic that may be sent by the recipient of sensor data traffic back to the source of the sensor data traffic. This type of traffic may be present if traditional transport-layer protocols like the Transport Control Protocol (TCP) are used to ensure reliable end-to-end transfer of sensor data. While the source/destination pairs for this type of traffic may be the same as for sensor node control traffic, the conditions under which it is required may differ.

For the assumed concept of operations we view the first two categories of traffic as essential for basic operation, but the need for the latter two may be more application specific. Thus, in some cases, routing support for the later two categories of traffic may not be required. Also note that we did not include a category for multihop traffic to be sent from one sensor node to another. Eliminating the need to support such *peer-to-peer* traffic provides another opportunity to reduce the amount of control overhead generated to support multihop routing.

NETWORKING PROTOCOLS

Network protocols available in commercial products today are typically designed for operation in quasi-static hardwired networks. To operate effectively under the resource constraints and dynamics of wireless networks, protocols must be designed to be scalable, adaptive, efficient and robust. While this topic has attracted renewed research interest in recent years, the focus has been limited to a few isolated “building blocks”—e.g., developing unicast routing protocols for mobile ad hoc networks or optimizing the performance of the TCP for use over satellite links. Synergistic, interoperable solutions that provide a unified approach to networking in a heterogeneous wireless infrastructure and offer an equivalent capability to that currently available in hardwired networks are far from complete. Overall, the commercial technology has yet to meet the public expectations of today; thus, realizing the expectations of future wireless networks will require innovative research and additional technological advancements.

Auto-configuration and Self-Organization

Techniques for auto-configuration and self-organization of the network infrastructure will be critical for rapid deployment of large sensor networks. The potential scale and dynamics of the scenario are unparalleled in existing networking applications. Manual configuration of hundreds, thousands, or possibly even greater numbers of nodes that are to be rapidly deployed in potentially hostile or high-risk areas is extremely impractical.

Commercial advances in auto-configuration have focused primarily on end system configuration, since the number of end hosts typically far exceeds the number of routers in the overall system and end hosts are more often added or removed. Recently developed and emerging network protocols such as the Dynamic Host Configuration Protocol (DHCP) and the Service Location Protocol (SLP) are easing the burden of end system configuration. However, auto-configuration of network routers and protocols and self-organization of the network infrastructure remains a largely unexplored research area.

While near-term solutions could focus on auto-configuration and mobile adaptation of existing protocols, longer-term solutions should consider new architectures, systems, and information handling mechanisms that are better suited for auto-configuration and self-organization. Research should not be limited to initial protocol configuration performance, but should investigate adaptive protocol operation and robustness as the infrastructure behavior changes over time.

Routing

There has been a significant amount of research on the topic of multihop routing in dynamic wireless networks and much of that research has applicability to distributed sensor networks. However, the design space is large and many of proposed approaches have focused on solving different problems within the design space—e.g., a small group of business associates with laptops using a single type of wireless technology to form an ad hoc network during a meeting, as opposed to a large group of soldiers and vehicles using a heterogeneous mix of wireless technologies to form a battlefield network. As a routing solution for sensor networks, the scalability of the various routing protocols with respect to the number of network nodes and the diameter of the network should be carefully considered.

A relatively common feature of many wireless network routing proposals is the *on-demand* behavior for constructing routes. This design choice is based on the notion that it may not be necessary (nor desirable) to maintain routes between all source/destination pairs at all times. In a dynamic topology, the control overhead expended to establish a route between a source/destination pair will be wasted if the source does not require the route prior to its invalidation due to topological changes. While on-demand behavior may reduce routing protocol control overhead in some situations, it may also increase route acquisition times. In some sensor network applications, nodes may remain silent for long periods of time—only waking up to send data upon detection of triggering event. On-demand routing behavior may be applicable to sensor networks for applications where the sensor data traffic is largely asynchronously triggered by relatively infrequent events and the route acquisition times are tolerable.

Destination-oriented Unicast Routing

In some unicast routing approaches (such as those based on link-state [BG 92, Moy 94] or path-finding [CRKG 89, RF 89] algorithms), the basic underlying algorithm inherently computes paths between *all* source/destination pairs. However, in other unicast routing approaches (such as those based on distance-vector [BG 92, Hedrick 88] or link-reversal [PC 97] algorithms), a logically separate version of the basic underlying algorithm is executed for each destination to which routing is required. These *destination-oriented* approaches may exhibit better scalability for scenarios where the number of destinations is small.

If we consider sensor networks where routing support for sensor node control traffic and sensor feedback traffic is not required, then a destination-oriented unicast routing protocol seems ideal for supporting delivery of the sensor data traffic—since the number of destinations to which routing is required is small. Even in cases where routing support for sensor node control traffic and sensor feedback traffic is required, supporting these traffic types through destination-oriented approaches or alternative mechanisms that build on the destination-oriented routing structure established for the sensor data traffic may be preferable to using a routing approach that provides routes between all source/destination pairs. Recall, that we have assumed there is no driving need to support multihop peer-to-peer traffic between sensor nodes.

In a dynamic topology, link-reversal approaches such as the Temporally-Ordered Routing Algorithm (TORA) [PC 97] may provide reduced routing control overhead and increased scalability when compared with destination-oriented approaches that support a shortest-path computation [PC 98]. TORA attempts to minimize communication overhead by localizing algorithmic reaction to topology changes. The scope of control messaging following a topological change is typically localized to a small set of nodes near the topological change. The design of TORA makes it potentially well-suited for supporting the delivery of sensor data traffic in large, dynamic, bandwidth-constrained sensor networks. TORA can be adapted to run in either an on-demand or a more tradition proactive routing mode. However, the preferred mode of operation for a particular sensor network would likely be dependent on the frequency and distribution of

sensor data traffic and the performance trade-offs (e.g., communication efficiency vs. route acquisition latency).

Anycast Routing

The anycast communication paradigm functionally provides the capability to locate and forward network traffic to *any one* of a set of distributed receivers within a network [PMM 93]. While there are many possible approaches to providing an anycasting capability, the use of anycast routing algorithms is perhaps the best-suited approach for the mobile wireless networking environment. In earlier works [PM 99a, PM 99b], we described how the link-state, distance-vector and link-reversal classes of unicast routing protocols can be extended to provide efficient construction and maintenance of anycast routes. We also presented simulation results that demonstrate how anycast routing techniques can provide a *one-to-any* communication capability with greater efficiency and robustness than traditional unicast based techniques. Anycast routing also provides a major benefit by allowing the network routing to do most of the work in locating and tracking the set of receivers—thus, easing configuration and management burdens of the end systems.

In the context of a distributed sensor network, anycast routing extensions may be useful when sensor data traffic need only be forwarded to any one of set of fusion (or gateway) nodes. Sensor nodes can send data to the appropriate anycast address and the dynamic anycast routing will forward this data to one of the set of fusion (or gateway) nodes. Under these conditions, a *single version of a destination-oriented unicast routing protocol with anycast extensions* could be used to support delivery of sensor data traffic, as opposed to running a separate version of the unicast routing protocol for each fusion (or gateway) node.

Multicast Routing and Network Broadcast

Efficiently supporting delivery of sensor network control traffic suggests the need for multicast routing protocols or network broadcast mechanisms (e.g., flooding). There are many alternatives and the performance trade-offs are likely to be very complex. In the following discussion we identify several potential approaches that are compatible with the destination-oriented unicast routing approach to supporting sensor data traffic.

If TORA is used for supporting delivery of sensor data traffic, then the Lightweight Adaptive Multicast (LAM) [JC 98] protocol (or an adaptation thereof) may be an appropriate solution for multicasting sensor network control traffic. LAM constructs a multicast routing tree on top of the unicast routing structure established by TORA. Thus, by design, LAM takes advantage of the properties of TORA (e.g., loop freedom) and benefits from its adaptability and efficiency.

There is also the option of using a multicast routing approach that is completely independent. Several multicast routing protocols have been proposed for use in dynamic wireless networks [GM 99, LGC 99]. Many of these approaches attempt to exploit the broadcast nature of a wireless transmission to improve the communication efficiency or traffic delivery.

Note that for this particular application we do not require support for joining or leaving a multicast group or for supporting more than one group. The requirement is simply for a network broadcast capability. This suggests the possibility of simpler (perhaps less adaptive) techniques for essentially flooding the traffic throughout the network. However, efficiency of the flooding mechanism is still important.

One approach to efficiently flooding traffic throughout a network is based on the concept of *multipoint relays* [QVL 00]. In this techniques, each network node attempts to identify a minimal set of *one-hop* neighboring nodes that can be used to forward traffic, such that the forwarded traffic will be received by the entire the entire set of two-hop neighbors. The multipoint relay sets of the individual nodes can collectively be used to flood traffic throughout the network. The approach is distributed, operates with *only* local topology information, and can improve the efficiency of network broadcast over basic flooding techniques.

Another area of research that merits consideration is that of energy-efficient broadcasting in wireless networks [WNE00]. The work considers node-based techniques that exploit the broadcast nature of wireless communications and has demonstrated performance improvements over adaptation of link-based techniques. A present limitation of this work is the need for full topological knowledge; the development of distributed approaches that provide similar performance benefits is the subject of ongoing research.

Reverse-Path Unicast Routing

While sensor node control traffic and sensor feedback traffic could be supported through the use destination-oriented unicast routing protocols, there are potentially more efficient approaches that take advantage of the routing structure established for delivery of the sensor data traffic.

One possible approach is to develop a protocol that establishes the reverse routes between the sensor nodes and the fusion (or gateway) nodes based on the the routing structure used to deliver sensor node traffic. These reverse routes could be either established on-demand or proactively depending on the frequency and distribution of sensor node control and sensor feedback traffic. Depending on the frequency of sensor data traffic, it may even be possible to cache the reverse routes upon forwarding of sensor data traffic and maintain them using soft-state mechanisms.

Transport Reliability

There are several expected characteristics of sensor networks (and dynamic wireless networks, in general) that limit the viability of transport-layer protocols designed for hardwired networks. Primarily, a significant amount of packet loss may not be due to congestion. Thus, the validity of a fundamental assumption upon which prior techniques for end-to-end congestion control are based must be challenged. Secondly, the dynamics of a sensor network will likely result in more frequent changes in both bandwidth and routes between source/receiver pairs. This introduces additional transients in the system and will likely also lead to an increase in out-of-order delivery and delay variance. All of these aspects can have a significant impact on the performance of transport-layer protocols.

Research and investigation of transport-layer issues in dynamic wireless networks and further developments are essential for providing end-to-end service comparable to that available in hard-wired networks. The use of *network-assisted* congestion indicators—e.g., Explicit Congestion Notification (ECN)—may provide an improved capability for sensing network congestion. With techniques such as ECN packets are tagged with congestion information as they are forwarded through the network. This provides a potentially less ambiguous indication of congestion than packet loss in the dynamic wireless networking environment. However, ECN tagged packets may be subsequently lost prior to reaching the receiver, end-to-end measures of delay and loss may still be important indicators. A hybrid of both network-assisted and end-to-end congestion indicators may provide the best solution. The challenge is in determining when and how to tag datagrams with congestion information and how to combine the network-assisted indicators with end-to-end indicators to improve performance over traditional techniques.

CONCLUSIONS

The development of large distributed sensor networks comprising hundreds, thousands or even greater numbers of small, inexpensive, low-power network computing and wireless communication devices is an extremely challenging and exciting research topic. The potential scale and dynamics of the scenario are unparalleled in existing networking applications and required innovative research both in the design of suitable networking protocols and in the development of techniques for auto-configuration and self-organization of the networking protocols and infrastructure.

We have presented a concept of operations for distributed sensor networks, considered the attributes of a network based in part upon this concept, and discussed the application of available and emerging networking

technologies. In defining the concept of operations, we limited the requirements for routing support to only supporting certain types of traffic. While it can be argued that greater routing functionality may be required for some scenarios, we conjecture that such sacrifices can significantly increase the scalability of the networking protocols and thus may also be required for some scenarios. Thus, there is likely a trade-off between functionality and scalability, and we focused on techniques for maximizing scalability.

There are several areas of research where additional developments are needed to better support the envisaged distributed sensor data networks—e.g., auto-configuration and self-organization of networking protocols and infrastructure. In other areas, such as multihop routing, we have identified several different approaches that may be readily adaptable for sensor network applications; however, additional investigation is required to assess relative performance and benefits of different approaches. The design space is complex and performance trade-offs will likely be partially driven by assumptions regarding the concept of operations. As the work progresses it will become increasingly important to refine these concepts.

REFERENCES

- [BG 92] D. Bertsekas and R. Gallager, *Data Networks*, (Prentice-Hall, 1992).
- [CRKG 89] C. Cheng, R. Riley, S. Kumar, and J.J. Garcia-Luna-Aceves, "A Loop-Free Extended Bellman-Ford Routing Protocol without Bouncing Effect." *Proc. ACM SIGCOMM '89*, Austin, TX, September 1989.
- [CM 99] S. Corson and J. Macker, "Mobile Ad hoc Networking (MANET) Routing Protocol Performance Issues and Evaluation Considerations," *Internet RFC 2501*, January 1999.
- [GM 99] J.J. Garcia-Luna-Aceves and E.L. Madruga, "The Core Assisted Mesh Protocol," *IEEE Journal on Selected Areas in Communications*, Special Issue on Ad-Hoc Networks, Vol. 17, No. 8, pp. 1380-1394, August 1999.
- [Hedrick 88] C. Hedrick, "Routing Information Protocol," *Internet RFC 1058*, June 1988.
- [JC 98] L. Ji and M.S. Corson, "A Lightweight Adaptive Multicast Algorithm," *Proc. IEEE Globecom '98*, Sydney, Australia, November 1998.
- [LGC 99] S.-J. Lee, M. Gerla and C.-C. Chiang, "On-Demand Multicast Routing Protocol," *Proc. IEEE WCNC '99*, New Orleans, LA, September 1999.
- [PC 97] V. Park and M.S. Corson, "A Highly Distributed Routing Algorithm for Mobile Wireless Networks." *Proc. IEEE INFOCOM '97*, Kobe, Japan, April 1997.
- [PC 98] V. Park and M.S. Corson, "A Performance Comparison of the Temporally-Ordered Routing Algorithm and Ideal Link-State Routing," *Proc. of IEEE Symposium on Computers and Communications (ISCC) '98*, Athens, Greece, June 1998.
- [PM 99a] V. Park and J. Macker, "Anycast Routing for Mobile Services," *Proc. Conference on Information Sciences and Systems (CISS) '99*, Baltimore, MD, January 1999.
- [PM 99b] V. Park and J. Macker, "Anycast Routing for Mobile Networking," *Proc. IEEE MILCOM '99*, Atlantic City, NJ, October 1999.
- [PMM 93] C. Partridge, T. Mendez and W. Miliken, "Host Anycasting Service." *Internet RFC 1546*, November 1993.
- [QVL 00] A. Qayyum, L. Viennot and A. Laouiti, "Multipoint Relaying: An Efficient Technique for Flooding in Mobile Wireless Networks," INRIA research report, 2000.

[Moy 94] J. Moy, "OSPF Version 2," *Internet RFC 2328*, April 1998.

[RF89] B. Rajagopalan and M. Faiman, "A New Responsive Distributed Shortest-Path Routing Algorithm," *Proc. ACM SIGCOMM '89*, Austin, TX, September 1989.

[WNE 00] J. Wieselthier, G. Nguyen and A. Ephremides, "On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks," *Proc. IEEE INFOCOM '00*, Tel Aviv, Israel, March 2000.