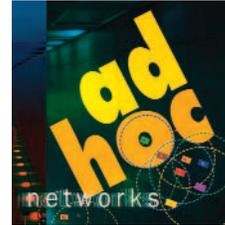


Group Communications in Mobile Ad Hoc Networks



This survey of approaches to group communications in mobile ad hoc networks explores several potential solutions to the unique problems of wireless mobile communications, which have variable and unpredictable characteristics due to mobility and signal strength fluctuations with respect to time and environment.

Prasant
Mohapatra
Chao Gui
Jian Li
University of
California, Davis

A mobile ad hoc network (manet) comprises a set of wireless devices that can move around freely and cooperate in relaying packets on behalf of one another. A manet does not require a fixed infrastructure or centralized administration. Because mobile nodes have limited transmission range, distant nodes communicate through multihop paths.

Their ease of deployment makes manets an attractive choice for a variety of applications. Examples include battleground communications, disaster recovery efforts, communication among a group of islands or ships, conferencing without the support of a wired infrastructure, and interactive information sharing. Unlike typical Internet applications, most applications of manets involve one-to-many and many-to-many communication patterns.

Efficient support of group communications is critical for most ad hoc network applications. However, manet group communications issues differ from those in wired environments for the following reasons: The wireless communications medium has variable and unpredictable characteristics and the signal strength and propagation fluctuate with respect to time and environment. Further, node mobility creates a continuously changing communication topology in which routing paths break and new ones form dynamically.

Because manets have limited bandwidth availability and battery power, their algorithms and protocols must conserve both bandwidth and energy. Wireless devices usually use computing compo-

nents—processors, memory, and I/O devices—that have low capacity and limited processing power. Thus, their communications protocols should have lightweight computational and information storage needs.

MULTICASTING

The multicasting communications model can facilitate effective and collaborative communication among groups. *Flooding* and *tree-based routing* represent two ends of the multicasting spectrum. Flooding is a simple approach that offers the lowest control overheads at the expense of generating very high data traffic in the wireless environment. The tree-based approach, on the other hand, generates minimal data traffic in the network, but tree maintenance and updates require many control-traffic exchanges. Both flooding and tree-based approaches scale poorly.

Multicast routing protocols for manets vary in terms of route topology, state maintenance, reliance on unicast routing, and other attributes. Instead of using a taxonomic approach to previously proposed multicasting protocols, our approach emphasizes the schemes' salient features.

Most proposed multicasting protocols primarily exploit one or more specific characteristics of the manet environment. These characteristics include variable topology, soft-state and state aggregations, knowledge of location, and communication pattern randomness. For example, mesh-based protocols exploit variable topology, stateless multicasting

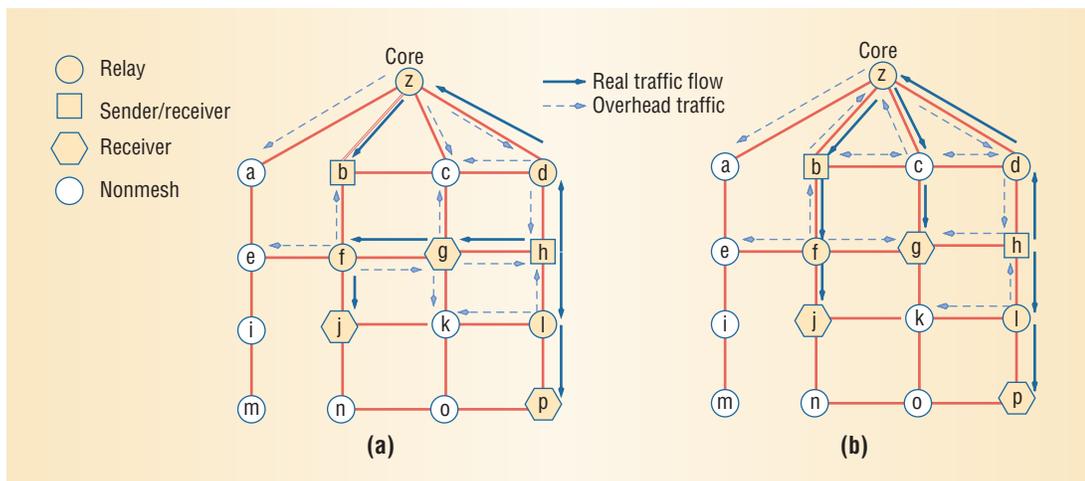


Figure 1. CAMP multicast routing. (a) Data-packet forwarding from node *h* using the core-assisted mesh protocol. (b) An equivalent shared tree, which uses a receiver-initiated router method. The solid arrows indicate the flow of actual traffic and the dashed arrows indicate the broadcast traffic due to the broadcast nature of wireless links.

exploits soft-state maintenance, location-aided multicasting exploits knowledge of location, and gossip-based multicasting exploits randomness in communication and mobility.

Mesh-based protocols

The addition of redundant paths between on-tree nodes converts a multicast tree into a mesh topology. The availability of alternative paths lets nodes deliver multicast packets regardless of link breakages. Mesh-based protocols thus achieve higher robustness against node mobility.

Core-assisted mesh protocol. CAMP¹ uses a shared mesh structure to support multicast routing in dynamic ad hoc networks. This structure ensures that the mesh includes the *reverse shortest paths*, the shortest paths from all receivers to the source.

Figure 1 shows how the protocol forwards data packets from node *h* to the rest of the group. To prevent packet replication or looping in the mesh, each node maintains a cache to keep track of recently forwarded packets. Periodically, a receiver node reviews its packet cache to determine whether it is receiving data packets from those neighbors not on the reverse shortest path to the source. When such situations arise, the node sends a heartbeat message to its successor in its reverse shortest path to the source. When the successor is not a mesh member, the heartbeat message triggers a *push join* message, which includes all nodes along any reverse shortest path in the mesh.

CAMP uses cores to limit the control traffic needed to create multicast meshes. Unlike the core-based tree protocol, CAMP does not require that all traffic flow through the core nodes. CAMP uses a receiver-initiated method for routers to join a multicast group. If a node wishing to join such a group finds it has neighbors that belong to the group, it simply updates its multicast routing table and uses a standard update procedure to announce its membership. When none of its neighbors are mesh members, the node either sends a join request toward a core or attempts to reach a group member using an

expanding-ring search process. Any mesh member can respond to a join request with a join ACK, which propagates back to the request originator.

On-demand multicast routing protocol (ODMRP). Based on a sender-advised approach to building a mesh, ODMRP² uses the *forwarding group* concept, in which a set of nodes forwards multicast data along the shortest paths between any member pairs. In ODMRP, each source establishes and updates a group membership and a multicast mesh *on demand*. By flooding a member advertising packet, a source node starts building a forwarding mesh for the multicast group, collecting membership information at the same time.

When a node receives a nonduplicate message requesting admission to the multicast group, it stores the upstream node identity and rebroadcasts the packet. When this request message packet reaches a multicast receiver, the receiver creates or updates the source entry in the *member table*. The system then uses the member table to prepare periodic control packets and broadcasts them via the receiver node. The nodes relay the packets back toward the source along the reverse path that the member-advertising packet traverses. This process constructs or updates the routes from sources to receivers and builds a mesh of nodes, called the *forwarding group*. Multicast sources send the member advertising packet periodically to refresh the membership information and update the routes. A soft-state approach maintains the multicast group and the mesh.

State maintenance

The techniques for maintaining multicast protocol states can be classified as stateless, constrained, or unconstrained. Stateless multicasting protocol nodes do not maintain any state information. Constrained state protocols reduce the state maintenance overhead through abstraction via application-layer multicasting or by aggregation via hierarchical multicasting. In unconstrained state protocols, both group members and nonmembers must maintain the protocol states to support a multicast group.

Broadcasting can provide a building block for route discovery in on-demand ad hoc routing protocols.

For multicasting in manets, a wider spread of protocol states restrict robustness and scalability. Changing network states requires more updates and exchange of control messages. If the routing tree or mesh involves fast-moving nonmember or member nodes, the multicast session will be severely hampered for unconstrained-state protocols. However, zero-state and constrained-state protocols are usually less affected by host mobility.

Location-aided multicasting

In networks that can access the Global Positioning System (GPS), the network provides each node with location and mobility information. Multicast protocols can also use this information to improve protocol robustness and performance. With GPS support, ODMRP can adapt to node movements and can use location and mobility information to estimate route expiration time, while receivers select the path that will remain valid longest. Sources can reconstruct routes in anticipation of route breaks, thereby making the protocol more resilient to node mobility.

Martin Mauve and colleagues³ proposed a *position-based multicasting* (PBM) technique that does not require flooding to maintain a tree or mesh structure. In PBM, a multicast source node finds a set of neighboring, next-hop nodes and assigns each packet destination to one next-hop node. The next-hop nodes, in turn, repeat the process. Thus, no global distribution structure is necessary.

Researchers have proposed two forwarding techniques for PBM. In *greedy multicast forwarding*, the next hop is selected based on the position of the forwarding nodes, its neighbors, and the destination. The distance toward the destination node is reduced at each hop.

When the greedy forwarding approach fails, the system adopts a recovery process using *perimeter forwarding*, in which it forwards the packet by traversing the network boundary gaps until it can resume greedy forwarding.

Gossip-based multicasting

Some multicasting protocols use gossip as a form of probabilistically controlled flooding to solve several problems, including network news dissemination. The basic idea of applying gossip to multicasting involves having each member node periodically talk to a random subset of other members. After each round of talk, the gossipers can recover their missed multicast packets from each other. In

contrast to deterministic approaches, a probabilistic scheme will better survive a highly dynamic ad hoc network because it functions independently of network topology and its nondeterministic property matches the network's characteristics.

Anonymous gossip. Manet designers can apply the anonymous gossip⁴ multicast performance enhancement technique atop any tree-based or mesh-based protocol with minimal overhead. This technique does not require a group member to have any knowledge of the other group members.

An anonymous gossip multicast protocol proceeds in two phases. In the first phase, a protocol multicasts packets to the group. In the second phase, periodic anonymous gossip takes place in the background as each group member recovers any lost data packet from other members of the group that might have received it.

Route-driven gossip (RDG). The route-driven gossip protocol⁵ relies on a unicast protocol such as DSR to provide routing information for guiding the gossip process. Each node maintains two data structures for a multicast group: a *data buffer* that stores received data packets and a *view*, which lists all other group member nodes known to this node. The view at each node consists of two parts: the *active view*, which contains the IDs of known members to which at least one routing path is known, and the *passive view*, which contains the IDs of known members to which no routing path is currently available.

A node seeking to join a group floods the network with a Group-Request message. All members receiving the message update their active view. They also return a Group-Reply to the request initiator with a certain probability. The initiator updates its active view after receiving the Group-Reply message. Each member node periodically generates a gossip message and sends it to a set of other nodes randomly chosen from its active view. The message includes a selected subset of the data buffer and the sequence number of the most recent missing data packets. A group member receiving a gossip message will update its view of other group members and update its data buffer with newly received data.

BROADCASTING

Network-wide broadcasting, which attempts to deliver packets from a source node to all other nodes in the network, serves an important function in manets. Broadcasting often provides a building block for route discovery in on-demand ad hoc routing protocols. When designing broadcast protocols for ad hoc networks, developers seek to

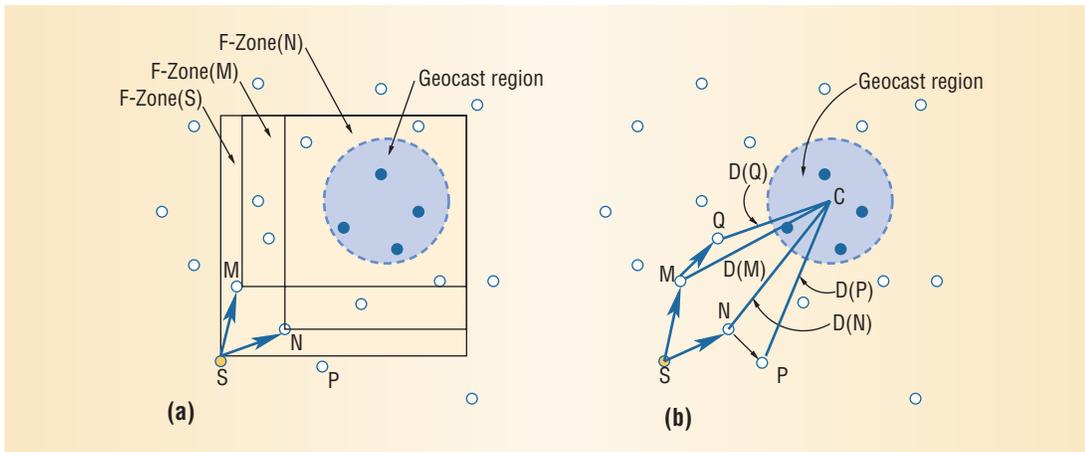


Figure 2. Location-based multicast forwarding schemes. (a) The first LBM scheme defines a rectangular forwarding zone, one corner of which lies at the source node and spans the geocast region. (b) The second LBM scheme uses a distance-based heuristic in which source node S defines the center point C of the geocast region in the geocast packets. Each intermediate node decides whether to forward a geocast packet by comparing its distance to the packet sender's distance.

reduce the overhead—such as collision and retransmission or redundant retransmission—while reaching all the network's nodes.

Although a wireless signal broadcast causes more contention and collisions in the shared wireless channel, it also allows a single transmission to reach multiple neighboring nodes. One comparison of existing techniques categorizes manet broadcast protocols into four types: simple flooding, probability-based, area-based, and neighbor-knowledge-based.⁶

In a more recent work,⁷ researchers proposed a general framework for self-pruning-based broadcast-redundancy-reduction techniques in ad hoc networks. Upon receiving a packet, intermediate nodes use the two proposed neighborhood coverage conditions to determine whether or not they should rebroadcast it. These coverage conditions depend on neighbor connectivity and the history of visited nodes. Since global network information is costly, the manet can use a distributed and local pruning process to select the forwarding node set based on local information such as the k -hop neighbor. Researchers have used this framework to propose new algorithms that combine features of previous work and show better performance.

GEOCASTING AND ANYCASTING

Applications that need to deliver messages of interest to every node in a specific geographical area can adopt *geocasting*, which is either flooding- or route-based. Each node's position with regard to the specified geocast region implicitly defines group membership. Each node is required to know its own physical location, which it can identify using the Global Positioning System. This does not require any explicit join and leave procedures. The group members tend to be clustered both geographically and topologically.

The IPv6 specification includes *anycast*, a similar Internet-based network service. Several servers, which jointly support a particular service, receive an anycast address. When a host sends its packets to this address, the network delivers the packets to at

least one and preferably only one of the servers in the anycast group. Although little work has been proposed for using anycasting in manets, researchers have used it in other applications, especially in battlefield or disaster-recovery communications.

Flooding-based geocasting

Flooding is the simplest way to deliver a message to all nodes in the network. Although expensive and inefficient, a simple flooding algorithm achieves the geocasting goal. Some flooding-based geocast protocols use the *forwarding zone technique* to constrain the flooding scope. A forwarding zone is a geographic area that extends from the source node to cover the geocast region. The source node defines a forwarding zone in the header of a geocast data packet. Upon receiving a geocast packet, other nodes will forward it only if its location is within the forwarding zone.

A geocast protocol's accuracy is defined as the probability that the transmission delivers a geocast packet to each geocast group member. Enlarging the forwarding zone can increase the accuracy. Given that the protocol overhead increases dramatically with an increase in the forwarding zone's size, a geocast protocol must achieve a workable tradeoff between the two factors.

Young-Bae Ko and Nitin H. Vaidya⁸ have proposed two flooding-based geocast protocols, both termed *location-based multicast* (LBM) schemes. Figure 2 shows these two LBM forwarding schemes. As Figure 2a shows, the first scheme defines a rectangular forwarding zone. One corner of the zone is at the source node and extends across the full geocast region. An adaptive-forwarding-zone technique ensures that each intermediate forwarding node redefines the forwarding zone by its location relative to the geocast region.

The second scheme uses a distance-based heuristic. As Figure 2b shows, instead of a forwarding zone, the source node S defines the center point, C, of the geocast region in the geocast packets. Each intermediate node decides whether to forward a geocast packet by comparing its distance to that of

In manets, an anycast protocol can simplify the access management and building process of a network distributed service.

the packet's sender. Thus, nodes M, N, and Q will forward the packet. However, node P decides not to forward the packet because the node also receives a geocast packet from node N, whose distance to center point C is shorter.

Route-based geocasting

Route-based geocast protocols use a two-step method for packet delivery. First, the protocol performs an anycast that delivers a geocast packet to any node within the geocast region. Thus, the source node builds a route to one or a few selected nodes in the intended region. Upon receiving a geocast packet, the selected nodes use a localized flooding method to further deliver the packet to all reachable nodes within the geocast region.

The GeoTORA⁹ protocol extends the unicast *temporally ordered routing algorithm* (TORA) for geocast routing. TORA potentially builds multiple routes from any source to a desired destination. The routing procedure assigns a height value to each network node, then uses their heights to determine the logical direction of a link between two nodes, working always from the higher to lower node. Thus, the destination node always has zero height, which the routing procedure uses to derive a destination-oriented *directed acyclic graph* (DAG). Any node that intends to send or forward a packet to the given destination simply follows the logical direction of the adjacent links.

GeoTORA adopts a similar method, building and maintaining a DAG for each geocast group. All nodes that belong to the geocast region have a zero-height link between a pair of nodes. If both end nodes have zero height, the system does not assign a direction. If a node wants to geocast to a region, it forwards the packet to a single node in that region—which, in turn, floods the packet within the region to reach all members.

Anycasting

Anycasting is defined as a point-to-point flow of packets between a single source and the “nearest” destination server identified by an anycast address.

In manets, an anycast protocol can simplify the access management and building process of a network distributed service. The protocol provides a route to the nearest server for a distributed service. It also maintains this route in the face of node mobility and can switch the connection to another server if needed. Rather than designing a completely separate anycasting protocol, extensions of

several different classes of unicast routing protocols—such as link state and distance vector—can efficiently construct and maintain anycast routes.

COMMON ISSUES IN GROUP COMMUNICATIONS PROTOCOLS

In addition to performance, some common issues that researchers have considered while designing most group communications protocols include energy conservation, reliability, security, and QoS support. However, the techniques for achieving these goals can differ significantly.

Reliability

Given the dynamic nature of manets, reliable group communications presents a challenging task. When node mobility is high, flooding becomes a viable approach for reliable group communications. We assume that mobility is not so high that flooding, or even its more persistent variations, becomes the only choice for reliable multicast and broadcast. Given node mobility and network dynamics, more efficient and flexible alternatives are available for reliable group communications in manets.

A broadcast protocol based on a clustering technique assumes that an underlying clustering protocol manages construction of a clustered architecture that covers the entire population of network hosts.¹⁰ The clustering protocol distributes the broadcast packets to form a forwarding tree consisting of cluster head nodes. This approach achieves reliability at the cost of maintaining the cluster structure proactively, even in the absence of traffic, and by using acknowledgments that travel backward along the path to the source node. Its efficiency also relies on the accuracy of the forwarding tree and underlying cluster structure, which present a challenging task when nodes move swiftly.

Obtaining a hard guarantee of manet reliability becomes extremely difficult when network size and mobility increase. The RDG protocol⁵ adopted a practical probabilistic specification that achieves high reliability without relying on any inherent multicast primitive. In RDG, each node has only a random partial view of the group, which results from the randomness of the routing information that any given node has.

RDG uses a pure gossip scheme in the sense that it gossips uniformly about multicast packets, negative acknowledgments, and membership information. A *gossiper push* mainly propels the spread of information, with each group member forwarding a multicast packet to a random subset of the group. This technique is complemented by a *gossiper pull* in which

the multicast packets piggyback any negative acknowledgments that the forwarding group members may generate. Given its nondeterministic characteristics, the notion of probabilistic reliability seems quite fitting in the dynamic manet environment.

A few efforts have focused on developing *medium-access-control* support for reliable group communications in manets. A new wireless ad hoc MAC protocol proposes a *broadcast medium window*.¹¹ BMW strives to ensure the reliable round-robin transmission of each packet to its neighbor. The protocol borrows some concepts from IEEE 802.11 and attempts to achieve reliable broadcast support at the MAC layer when the traffic load is manageable. If reliable transmission is counterproductive, BMW reverts to the unreliable delivery of IEEE 802.11. BMW's round-robin approach does not take advantage of the wireless signal's broadcast nature, so it can incur significant overhead by unicasting packets to each neighbor.

Energy efficiency

Since a limited battery source typically drives nodes in manets, designing energy-conserving protocols becomes essential. Even when energy is not a stringent constraint, reducing power consumption can result in less interference and better throughput.

Researchers can use various techniques to build power-aware and energy-efficient broadcast and multicast infrastructures in manets. Wireless transmission provides the greatest contributor to energy consumption in ad hoc networks, so reducing the number of nodes that participate in transmissions can reduce the total energy for a broadcast and multicast process. Many protocols thus strive to minimize the forwarding node set.

Several proposed techniques for energy-efficient broadcast and multicast share a common feature: combining a minimum or reduced forward-node set with power-level selection. The *broadcast incremental power* (BIP) protocol¹² adds new nodes to the multicast tree one at a time, starting from the source node. BIP bases its decision to add a specific node at each step on which node it can add with a minimum of additional transmission energy. A leaf or parent node with increased transmission power also can reach this new node. A greedy heuristic, BIP requires global network information, but it might not generate the minimum-cost tree. Another proposed localized algorithm requires only neighborhood information and attempts to take advantage of wireless transmission's broadcast nature.¹³

Energy consumption from retransmission at the

data-link layer when computing the minimum-energy-cost tree should also be considered when designing protocols. Although many efforts have been made to design energy-efficient broadcast and multicast protocols, issues such as how to address energy efficiency in highly mobile manets and how to factor in traffic conditions when using contention-based MAC protocols still present wide-open challenges.

The reception and idle-listening process provides another power-consumption source.

To reduce power consumption during idle listening, the power-aware PAMAS MAC protocol¹⁴ selectively turns off some network nodes for certain durations. PAMAS has a separate signaling channel for manets. PAMAS overhears exchanges of node request-to-send and clear-to-send messages and uses this information about traffic demand and neighboring node status to determine when a mobile node should sleep, for how long, and what to do if a neighboring destination node is asleep.

The *wakeup mechanism* also plays a critical role in designing power-efficient protocols. Existing wakeup mechanisms fall into three categories:¹⁵

- *On-demand wakeup* typically uses a wakeup radio to awaken a neighboring node.
- *Scheduled rendezvous* requires that sleeping nodes wake up at the same time periodically to meet the potential demand for intercommunication. This approach is unsuitable for a multihop environment because it requires time synchronization among all nodes.
- *Asynchronous wakeup mechanisms* do not require time synchronization among different nodes. The sleep and wakeup schedules are designed so that any two neighboring nodes will have overlapped active time within a specified number of cycles.

Quality of service

QoS is usually defined as a set of service requirements that the network must meet while transporting a packet stream from a source to its destination. The network is expected to guarantee the performance of a set of measurable prespecified service attributes such as delay, bandwidth, probability of packet loss, and delay variance. Two other QoS attributes, power consumption and service coverage area, are more specific to manets.

With the increase in QoS needs in evolving applications, supporting QoS-aware group communi-

Since a limited battery source typically drives nodes in manets, designing energy-conserving protocols becomes essential.

With advances in support for group communications, the use of multimedia objects will proliferate in manet application domains.

cations in manets is also desirable. Resource limitations and variability further add to the need for QoS in such networks.¹⁶ However, the characteristics of these networks make QoS support a complex process, thus QoS-aware group communications remains an open problem.

A proposed QoS-aware core migration protocol for the multicast protocols uses a group-shared multicast tree.¹⁷ This protocol seeks to construct a tree in which the leaves achieve the multicast application's desired qualities. To reduce the communications cost, the protocol conducts the core selection algorithm only on the current core node. If, for example, delay is the standard QoS metric, the core records the history of delays to group members in terms of the relative time difference between sending the packets to the core and receiving the corresponding acknowledgments from the respective subtree branches. If the averaged delay exceeds the QoS requirement by a given threshold, the core selects a better core candidate from the members close by. Thus, core migration occurs incrementally, which is more suitable to manet dynamics.

Security

Because security is an essential manet requirement, researchers have proposed secure routing protocols for unicast routing in ad hoc networks. Ariadne¹⁸ is a secure on-demand unicast routing protocol that prevents attackers or compromised nodes from tampering with uncompromised nodes. It uses symmetric cryptography and an efficient broadcast authentication scheme to prevent denial of service attacks.

Group communications amplify security concerns because they involve more nodes. However, research in this area is just beginning. Ad hoc networks have created additional challenges for implementing security services for wireless communications. The wireless broadcast media is more prone to both passive and active attacks. MAC layer solutions to group-key management and source-authentication proposals for wireline networks must be modified and enhanced for use in a wireless environment. Compared with other wireless communications such as cellular networks, ad hoc networks require even more sophisticated, efficient, and lightweight security mechanisms to achieve the same goals.

Once again, the dynamic characteristics of manets cause these extra challenges. First, ad hoc

networks lack trusted centralized infrastructure, which previous security proposals for wireline networks often required. Threshold-based and quorum-based approaches have been investigated to address this problem.

Second, the wireless links between nodes in a manet form and dissolve unpredictably, resulting in ephemeral relationships between nodes. These relationships make it more difficult to build trust based on direct reciprocity.

Third, proposed ad hoc group communication schemes differ markedly from those proposed for wireline networks. In some applications, especially in hostile environments such as battlefield communications, attackers can capture and compromise individual mobile nodes, posing a severe threat to the entire ad hoc network.

Finally, given the stringent nodal budgets in many applications, any solutions proposed for manets must view overhead as a key concern. These applications need strong security mechanisms, yet the solutions must be lightweight in terms of message overhead and computational cost.

For most mobile ad hoc network applications, group communications will be as frequent as unicast communications. Considering the wireless medium's broadcast nature, manets require effective and efficient support for group communications. Although much work has been reported on unicast communications, efforts to improve group communications in manets have not kept pace. Important aspects that researchers must pursue more aggressively include efficient MAC layer support for multicasting and broadcasting and providing reliable and efficient transport layer support.

With advances in support for group communications, the use of multimedia objects, such as video, audio, or images from various sites will proliferate in manet application domains. Thus, developers must design and integrate QoS support into group communications protocols. The wireless medium's nature, as well as that of its potential applications, warrants the integration of security aspects in all group communications protocols. ■

Acknowledgment

This work was supported in part by the National Science Foundation under grants ANI-0240502 and ANI-0296034, and a generous gift from Hewlett-Packard Corporation.

References

1. J.J. Garcia-Luna-Aceves and E.L. Madruga, "The Core-Assisted Mesh Protocol," *IEEE J. Selected Areas in Comm.*, Aug, 1999, pp. 1380-1394.
2. S.J. Lee, W. Su, and M. Gerla, "On-Demand Multicast Routing Protocol in Multihop Wireless Mobile Networks," *Mobile Networks and Applications*, ACM/Kluwer, vol. 7, no. 6, 2002, pp. 441-453.
3. M. Mauve et al., "Position-Based Multicast Routing for Mobile Ad Hoc Networks," poster paper, MobiHoc 2003.
4. R. Chandra, V. Ramasubramanian, and K. Birman, "Anonymous Gossip: Improving Multicast Reliability in Mobile Ad-Hoc Networks," *Proc. IEEE Int'l Conf. Distributed Computing Systems*, IEEE CS Press, 2001, pp. 275-283.
5. J. Luo, P.T. Eugster, and J-P. Hubaux, "Route Driven Gossip: Probabilistic Reliable Multicast in Ad Hoc Networks," *Proc. IEEE Infocom 2003*, IEEE Press, 2003, pp. 2229-2239.
6. B. Williams and T. Camp, "Comparison of Broadcasting Techniques for Mobile Ad Hoc Networks," *Proc. ACM Int'l Symp. Mobile Ad Hoc Networking and Computing* (MobiCom 2002), ACM Press, 2002, pp. 194-205.
7. J. Wu and F. Dai, "Broadcasting in Ad Hoc Networks Based on Self-Pruning," *Proc. IEEE Infocom 2003*, IEEE Press, 2003, pp. 2240-2250.
8. Y-B. Ko and N.H. Vaidya, "Flooding-Based Geocasting Protocols for Mobile Ad Hoc Networks," *Proc. Mobile Networks and Applications*, Kluwer Academic, vol. 7, no. 6, 2002, pp. 471-480.
9. Y-B. Ko and N.H. Vaidya, "GeoTORA: A Protocol for Geocasting in Mobile Ad Hoc Networks," *Proc. 8th Int'l Conf. Network Protocols*, IEEE Press, 2000, pp. 240-250.
10. E. Pagani and G.P. Rossi, "Providing Reliable and Fault-Tolerant Broadcast Delivery in Mobile Ad-Hoc Networks," *Mobile Networks and Applications*, vol. 4, 1999, pp. 175-192.
11. K. Tang and M. Gerla, "MAC Reliable Broadcast in Ad Hoc Networks," *Proc. MilCom 2001*, IEEE Press, 2001, pp. 1008-1013.
12. J. Wieselthier, G. Nguyen, and A. Ephremides, "On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks," *Proc. IEEE Infocom 2000*, IEEE Press, 2000, pp. 585-594.
13. J. Cartigny, D. Simplot, and I. Stojmenovic, "Localized Minimum-Energy Broadcasting in Ad Hoc Networks," *Proc. IEEE Infocom 2003*, IEEE Press, 2003, pp. 2210-2217.
14. C.S. Raghavendra and S. Singh, "PAMAS: Power Aware Multi-Access Protocol with Signaling for Ad Hoc Networks," *ACM Computer Comm. Rev.*, July 1998, pp. 5-26.
15. R. Zheng, J.C. Hou, and L. Sha, "Asynchronous Wakeup for Ad Hoc Networks: Theory and Protocol Design," *Proc. 9th Ann. Int'l Conf. Mobile Computing and Networking* (MobiCom 2003), ACM Press, 2003, pp. 35-45.
16. P. Mohapatra, J. Li, and C. Gui, "QoS in Mobile Ad Hoc Networks," *IEEE Wireless Comms.*, June 2003, pp. 44-53.
17. M. Kochhal et al., "An Efficient Core Migration Protocol for QoS in Mobile Ad Hoc Networks," *Proc. IEEE Int'l Performance Computing and Comm. Conf.*, IEEE Press, 2002, pp. 387-391.
18. Y. Hu, A. Perrig, and D.B. Johnson, "Ariadne: A Secure On-Demand Routing Protocol for Ad Hoc Networks," *Proc. 8th Ann. Int'l Conf. Mobile Computing and Networking* (MobiCom 2002), ACM Press, 2002, pp. 12-23.

Prasant Mohapatra is a professor in the Department of Computer Science at the University of California, Davis. His research interests include wireless networks, sensor networks, and quality-of-service issues in Internet and Web servers. Mohapatra received a PhD in computer engineering from Pennsylvania State University. Contact him at prasant@cs.ucdavis.edu.

Chao Gui is a PhD candidate in the Department of Computer Science at the University of California, Davis. His research interests include wireless networking and mobile computing. Gui received an MS in computer science from the University of Central Florida. Contact him at guic@cs.ucdavis.edu.

Jian Li is a PhD candidate in the Department of Computer Science at the University of California, Davis. His research interests include mobile ad hoc and sensor networks. Li received an M.Eng. in pattern recognition and automation from Tsinghua University, Beijing, China. Contact him at lijian@cs.ucdavis.edu.

For further information on this or other computing topics, visit our Digital Library at <http://www.computer.org/publications/dlib>.