

Abstract

Unlike static networks, ad hoc networks have no spatial hierarchy and suffer from frequent link failures which prevent mobile hosts from using traditional routing schemes. Under these conditions, mobile hosts must find routes to destinations without the use of designated routers and also must dynamically adapt the routes to the current link conditions. This article proposes a distributed adaptive routing protocol for finding and maintaining stable routes based on signal strength and location stability in an ad hoc network and presents an architecture for its implementation.

Signal Stability-Based Adaptive Routing (SSA) for Ad Hoc Mobile Networks

ROHIT DUBE, CYNTHIA D. RAIS, KUANG-YEH WANG, AND SATISH K. TRIPATHI
UNIVERSITY OF MARYLAND

Mobility is becoming increasingly important for users of computing systems. Technology has made possible wireless devices and smaller, less expensive, and more powerful computers. As a result users gain flexibility and the ability to maintain connectivity to their primary computer while roaming through a large area. The number of users with portable laptops and personal communications devices is increasing rapidly. These users would like the ability to exchange information in any location. The necessary mobile computing support is being provided in some areas by installing base stations and access points. Mobile users can maintain their connectivity by accessing this infrastructure from home, from the office, or while on the road.

Such mobility support is not available in all locations where mobile communication is desired. Access points may not be set up due to high cost, low expected usage, or poor performance. This may happen during outdoor conferences or in emergency situations like natural disasters and military maneuvers in enemy territory. If mobile users want to communicate in the absence of a support structure, they must form an *ad hoc network*.

Ad hoc networks consist of mobile hosts in a network bereft of base stations and predesignated routers, and are characterized by a highly dynamic network topology. The network topology changes frequently due to host migration, signal interference, and power outages. The mobile hosts in such a network can communicate directly with neighboring hosts through the shared wireless media, but communication with non-neighboring hosts requires a distributed routing algorithm. Traditional routing protocols pass detailed topology information between hosts and are not effective in ad hoc networks due to the high rate of topology change. In a dynamic network, the topology information (in the routing tables) is soon out of date, and the propagation of the routes is too slow to be accurate. In addition, the changing topology exacerbates the problem of looping routes. Hence,

new routing algorithms need to be developed to support ad hoc networks.

Ad hoc networks require a highly adaptive routing scheme to deal with the frequent topology changes. In this article, we propose a routing protocol that utilizes the ad hoc network characteristics to select the most stable routes through the dynamic network. The proposed Signal Stability-Based Adaptive Routing (SSA) protocol is novel in its use of signal strength and stability of individual hosts as route selection criteria. Selecting the most stable links (i.e., those which exhibit the strongest signals for the maximum amount of time) leads to longer-lived routes and less route maintenance. In this protocol, a host initiates route discovery on demand: only when a route is needed to send data. The source broadcasts route-search packets which will propagate to the destination, allowing the destination to choose a route and return a route reply.

Since it is not feasible to implement a large ad hoc network for experimental purposes, we perform simulations to investigate the benefits and costs of using signal strength and location stability as the route selection criteria. The results show that the use of signal strength consistently decreases the route maintenance required by providing longer-lived routes. Our results also show that location stability should be used only in certain scenarios since misinformation about stability patterns is very costly and has a negative impact on routing and performance.

The next section of this article discusses background and related work. The third section introduces the basic SSA Protocol and gives an example of its use. The fourth section presents details of the protocol, and the fifth section discusses extensions to the protocol and interoperability with Mobile IP (Internet Protocol). The sixth section presents our simulation results. The final section presents conclusions and some directions for future work.

Background and Related Work

Traditionally, next-hop routing protocols have been classified into distance-vector and link-state approaches. Distance vector protocols, like RIP [1], are based on the distributed Bellman-Ford (DBF) algorithm and rely on main-

This work is supported in part by National Science Foundation grant CCR 9318933, IBM equipment grants, and support from Novell.

taining distance estimates to all nodes in the network. Each node monitors the cost of its outgoing links and periodically broadcasts, to its neighbors, its minimum distance estimate to each destination in the network. Although the storage requirements are minimal, the DBF algorithm may converge very slowly and may require a large number of messages to respond to a topology change [2]. In link-state protocols, like OSPF [3], each node maintains the shortest-path tree to all nodes in the network which it uses to make routing decisions. Updates are broadcast to maintain consistent views of the network at each node. The maintenance of a view of the entire network involves large storage and communication overhead. Both mechanisms rely on exchanging detailed information among the nodes in the network [4]. These traditional routing algorithms, which depend on hierarchical addressing schemes and pass detailed routing information between designated routers, are not effective in ad hoc networks due to the high rate of topology change [5-7].

Routing approaches for ad hoc networks assume a rate of change that is not high enough to make flooding the only alternative and not low enough to make the use of traditional routing algorithms effective. One type of ad hoc routing method seeks to modify existing DBF routing algorithms for use in a dynamic topology. The Destination Sequenced Distance Vector (DSDV) algorithm [8] modifies the DBF algorithm to prevent looping by including sequence numbers to order the routing information. It also delays the advertisement of unstable routes to reduce fluctuations of the routing tables and the number of rebroadcasts of the same sequence route entries. The Wireless Routing Protocol (WRP) [9] uses a path-finding algorithm that includes second-to-last hop to a destination to identify a route and prevent temporary looping. The approach presented in [10] finds and maintains clusters in the ad hoc network. The boundary nodes connect clusters and perform routing using a traditional distance vector routing protocol. The cost of this method is maintaining the cluster membership when the network topology changes. For each of these methods, each node, or each boundary node, maintains routes to all destinations in the network. Information is propagated through the network to achieve this goal.

Another type of ad hoc routing algorithm uses an on-demand route discovery philosophy. Routes to a destination are sought only if the node has data to send to that destination. The Dynamic Source Routing (DSR) proposed in [11] uses broadcasts to propagate the route search, and each node that forwards the route search adds its address to the source list. When the request reaches the destination, a complete route is listed in the packet. Routes discovered in this way or by hearing neighboring hosts' communications are recorded in the cache. Data packets are forwarded by placing the source route (hop list) in the packet. The Lightweight Mobile Routing (LMR) approach proposed in [12] floods the network with a query broadcast when the route is desired. The nodes receiving this query broadcast either broadcast a reply with the requested route, if this node has a route to the requested destination; otherwise, they simply forward the broadcast query. This approach uses link status to create routes and prevent looping. Route erasure is required when the topology changes. The Associativity Based Routing

Method	Distance vector	On-demand	Location stability	Signal strength	Routing overhead	Pkt processing overhead
DSDV	ρ				High	Low
WRP	ρ				High	Low
Cluster	ρ				High	Low
DSR		ρ			Low	High
LMR		ρ			Low	Low
ABR		ρ	ρ		Low	Low
SSA		ρ	ρ	ρ	Low	Low

■ Table 1. Ad hoc routing methods.

(ABR) approach [13] also uses broadcast queries to find desired routes, but the optimal route is selected by the destination based on the stability of the route and shortest path. The main criterion is the intermediate node stability, which is based on the idea that nodes which have been stationary for a threshold period are less likely to move. This method also uses route erasure and maintenance when topology changes cause a route failure.

The main difference between the approaches discussed above and our approach is our utilization of the information available at the link level to choose routes. The signal quality of the channel is used to determine whether portions of the topology are stable or fluctuating at any given time. Like the second type of algorithm, routes are determined only on demand; however, we do not limit the rate of change of the topology or suggest that all parts of the topology are equally stable. We select routes through the most stable areas of the network, using an adaptive algorithm to ensure successful data transmission in a highly dynamic topology. We compare the various protocols' routing types, criteria for selecting routes, and overhead in Table 1. Routing overhead refers to the cost of route setup and maintenance packets. Packet processing overhead refers to the cost at each intermediate node of routing data packets and the per-packet header overhead.

In addition to correctness, a good routing protocol for ad hoc networks should have a low communication overhead, by minimizing route setup and maintenance messages. The protocol should also be simple and efficient, since it must quickly converge to select a route before network changes make the route invalid. The routing protocol must be distributed since no centralized host is available and to make the protocol scalable. It must be loop-free and should have minimal memory overhead. Finally, it should take advantage of the technology available and utilize information about the current state of the network.

Signal Stability-Based Adaptive Routing (SSA) Protocol

The SSA protocol performs on-demand route discovery by selecting longer-lived routes based on signal strength and location stability. The signal strength criterion allows the protocol to differentiate between strong and weak channels. Each channel is characterized as strong or weak by the average signal strength at which packets are exchanged between the hosts at either end of the channel. The location stability criterion biases the protocol toward choosing a channel which has existed for a longer period of time. Together, these two concepts form the signal stability criterion that chooses strong channels which have been in existence for a time greater than some threshold.

Protocol Overview

A source initiates a route discovery request when it has data to send to a destination that is not in the routing table. The route search is broadcast to all neighboring hosts. These hosts propagate the broadcast if it is received over a strong channel and the request has not been propagated previously (to avoid looping). The route search packet stores the address of each intermediate host in the route taken. The destination chooses the route recorded in the first arriving request, since this route is probably shorter and less congested than routes for slower-arriving requests. The destination returns the route reply along the selected route, and each intermediate node includes the new next-hop, destination pairs in its routing table.

Functionally, the SSA protocol consists of two protocols, the Forwarding Protocol (FP) and the Dynamic Routing Protocol (DRP), which utilize the extended device driver interface. This interface is responsible for making available to the routing protocols the signal strength information from the device. DRP maintains the routing table by interacting with the DRP on other hosts. FP performs the actual routing table lookup to forward a packet to the next hop.

Protocol Modules

The Forwarding Protocol (FP) and Dynamic Routing Protocol (DRP) work together to route packets in the ad hoc network. The extended device driver interface enables communication between these routing protocols and the link layer for the sending and receiving of packets and the receiving wireless link quality information. Two tables are maintained to enable SSA routing: the Signal Stability Table (Table 2) and the Routing Table (Table 3).

Each host sends out a link layer beacon¹ to its neighbors once every time quantum, denoted by a *click*. Every host receiving this beacon records the *signal strength* at which the beacon was received in the Signal Stability Table (SST). Each host also classifies its neighbors as strongly connected (*SC*, hence belonging to the *SC set*) if the host has been receiving strong beacons from the neighbor for the past few clicks. The neighbor is otherwise classified as weakly connected (*WC*, hence belonging to the *WC set*). A host marked as SC in the SST also has an entry in the Routing Table (RT) which stores destination and next-hop pairs for each known route. The SST also has a column, *Last*, to indicate that a beacon was received from a host within the last click and a column, *Clicks*, to record how long beacons have been continuously received with a strong signal from each neighboring host.

The availability and processing of signal strength information is made possible by the extended device driver interface, which provides the DRP with the average signal strength at which a packet was received and the address of the immediate sender. The DRP uses the extended interface to maintain the statistics in the SST. It then uses the SST to maintain routes to neighboring hosts in the RT and non-neighbor routes via information provided by route search, route reply, error, and erase messages.

The FP functions by looking up the destination in the RT

¹ Beacons are "I am alive messages" exchanged between wireless devices at regular intervals to maintain connectivity. SSA does not add overhead by defining any new beacons.

Host	Signal strength	Last	Clicks	Set
Y				
Z				

■ Table 2. The signal stability table (SST).

Destination	Next hop
Y	
Z	

■ Table 3. The routing table (RT).

and forwarding the packet on the next hop for the destination. When there is no entry for the destination in the RT, the FP initiates a *route search* to find a route to this destination. The route search message has a hop list which records the path taken by the message. Each intermediate DRP uses this list to prevent loops and adds its own address to the hop list. Although the destination DRP may receive multiple copies of a route search message, it simply selects


the route contained in the first arriving message and tunnels a *route reply* message on the reverse path to the source. The DRP at each intermediate host installs the appropriate next-hop entry for the destination and the source in its RT. When a route reply message is received, the DRP at the source updates the RT, and the FP routes the data via the next hop found in the RT.

A route may become unavailable due to migration of the hosts along the route. When a host moves out of range of its neighbors or shuts down, the neighbors will recognize that the host is unreachable since they no longer receive beacons from that host. The DRP will modify the SST and RT to reflect the changes. Any routes that have this unreachable host as the next hop will become invalid. When the host receives a packet to forward along an invalid route, FP will determine the lack of a route and will notify the source via an error message. The source FP will initiate a new route discovery to find an available route, and it will send a message to erase the invalid route.

An Example

Consider the example ad hoc network shown in Fig. 1. A solid edge indicates that the vertices that make up the edge are in each others' SC sets. The dashed edges indicate that the corresponding vertices are in the WC set.

In part A of Fig. 1, host A has data to send to host F and therefore wants to find a route to destination F. The protocol starts by sending out a broadcast route-search message seeking destination F. Route-search packets that arrive at any host over WC links are dropped, such as the packet sent from A to C. The route-search packet sent to B over the SC link will be broadcast by B to its neighbors. A will drop the packet, having already forwarded it, D and C will forward the packet and mark that they have seen it. Note that C has not already marked the packet as seen since it dropped the packet over the WC link from A without processing it. Packets will arrive at the destination F via two paths: A, B, D, F and A, B, C, E, F. Assuming that the links have similar latencies and traffic levels, the route-search will arrive over the shorter path first: A, B, D, F. The destination, F, will select this route and return the route-reply along the reverse path, F, D, B, A. Each intermediate host along the path will install entries in its RT. Once the route is installed in the RT of the source, A



■ Figure 1. An example network.

forwards the data packets via the next hop to destination F, which is host B.

Assume now that the link B-D disappears, as shown in part B of Fig. 1. When B realizes that host D is unreachable, it sends an error message to A. A then initiates another route search and sends an erase message to erase the invalid route. F again selects the first arriving route search, which will be the route A, B, C, E, F.

If, in the future, the link between B and C disappears, a strong route will no longer exist to the destination. In this case, no route-search packets will be propagated to the destination, since packets that arrive over the WC links are not forwarded. When A does not receive a reply after some timeout period, it must decide whether it wants to find any route or wait and try to find a strong route at a later time. If it wants any route, it will send a route-search message specifying *any* route, and will find the route A, C, E, F, which has one WC link.

SSA Protocol Details

We present the architecture, packet format, and protocol details in this section to provide a clearer understanding of the protocol as well as to illustrate the architecture that could be used to implement such a protocol. For clarity, the algorithm discussed below presents a synchronous processing scenario, but an implementation of the algorithm could also be asynchronous.

Since ad hoc networks exhibit no spatial hierarchy, we assume a flat addressing scheme. For simplicity of exposition we use the medium access control (MAC) addresses of the wireless device as the address of a node.

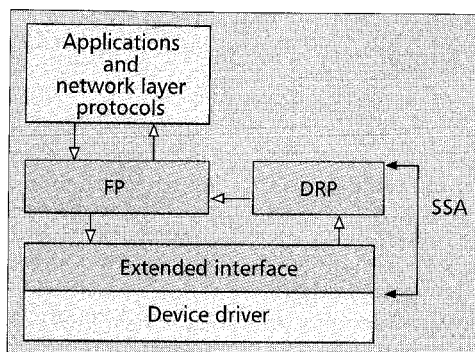
Protocol Stack

The FP and DRP are located between the network layer and the link layer, as shown in Fig. 2. This makes the SSA routing protocol interoperable with Mobile IP, as will be discussed later. All incoming packets pass from the extended device driver interface to DRP. The DRP updates the SST and the RT and relays the appropriate types of packets to the FP. The FP then either passes a packet up the stack or forwards the packet through the wireless device driver on to the next hop. All transmissions go out via FP, and all receives come in through DRP. This division simplifies the protocol by exporting a single interface for outgoing packets and also separating out the filter for incoming packets.

Packet Format

Figure 3 shows the packet format expected by the protocols. *SA* and *DA* refer to the source and destination addresses, respectively. *SEQ* is a sequence number assigned by the source, which is useful for route searches. *TTL* is a time-to-live field used to eliminate erroneous packets looping in the network. *TYPE* distinguishes between messages and is one of the following: UNICASTDATA, FLOODDATA, ROUTESEARCH, ROUTEREPLY, ERROR, or ERASE. *PREF* allows a host initiating a route-search to specify the quality of the route desired. This field can be STRONGLINKONLY or NOPREFERENCE. *LEN* is the length of the entire packet, and *CRC* is the checksum.

The rest of the packet contains either data (for a packet of type UNICASTDATA or FLOODDATA); the recorded route hop list (for a packet of type ROUTESEARCH or ROUTEREPLY); or the destination address of a stale route (for a packet of type ERROR or ERASE).




■ Figure 2. The protocol stack at each mobile host.

SSA Broadcast and Flooding

It is important to distinguish between MAC-level broadcasting and SSA broadcasting. Every SSA packet is encapsulated in a MAC frame before transmission, and every MAC frame is decapsulated to an SSA packet on reception. If the MAC address of the encapsulated frame is the broadcast address, then all hosts on that shared wireless media which receive the frame (immediate neighbors) will pass the packet up to the DRP for processing.

On the other hand, an SSA broadcast has its DA equal to the broadcast address. This type of broadcast packet will be delivered to all hosts within this ad hoc network and passed up to the network layer on each receiving host as well as forwarded. To achieve this, the MAC address of this packet is also the broadcast address.

A flooded packet, of TYPE FLOODDATA, must be forwarded by any host that receives it and is not the SSA destination. Such a packet is forwarded through a MAC-level broadcast even if a route to the destination is unknown. The packet is not passed up to the network layer until it reaches the SSA destination.



■ Figure 3. SSA packet format.


In summary, an SSA broadcast reaches the entire ad hoc network; flooding tries to reach a specific destination by propagating through the entire network; and a MAC-level broadcast sends the packet to the SSA at all immediate neighbors.

The FP

The FP accepts packets from the DRP and higher-layer protocols. If the destination address of the packet matches that of the host, the packet is pushed up the protocol stack. FP forwards all other packets through the device driver. Broadcasts are sent out without checking the RT, and the FP performs a routing table lookup for unicast packets. If no entry is found for the unicast destination, a route needs to be found. If this host is the source of the packet, a route search is initiated. When a route is found, the DRP will receive the route reply and install the route in the RT. The FP then forwards the original packet to the next hop. Alternatively, if this host is not the source of the packet, a packet of type error is sent back to the source, which will send a message of type ERASE to tear down the old route and initiate a new route search. In this case, the original packet will be dropped by the intermediate host.

Table 4 shows the details of the FP packet processing. `initiateRouteSearch()` produces and broadcasts a route search packet with an empty data portion, a unique sequence number in *SEQ*, and a *TYPE* field of ROUTESEARCH. The source can choose an appropriate *PREF* value depending on the needs of the upper-layer application or protocol. We choose STRONGLINKONLY for the first try and NOPREFERENCE for any ensuing request retry after a timeout. If, after several attempts, no satisfactory route is found, the data packet is broadcast and reaches the destination via flooding. An alternate approach to flooding would be to report an exception to the application which generated the packet.

The routine `sendRouteError()` initiates a route error



■ Figure 4. Establishment of a route.

packet, with DA equal to srcAddr and TYPE equal to ERROR. It sends the error packet to the original source for which forwarding was being attempted. The route maintenance packets (of TYPE ERROR or ERASE) are unicast with best-effort delivery. This means that any hosts which do not have an entry for the forwarding destination will drop the packet without further action. The subsection on route maintenance explains this further.

All outgoing SSA packets are encapsulated with the current MAC address and the next-hop MAC address, or with the MAC broadcast address if every neighboring host should receive the packet. In the pseudo-code, forward(pkt) is defined to be a unicast to a particular host's MAC address; broadcast(pkt) sends the packet to all reachable hosts by using the MAC broadcast address at the destination. Unicast data is processed only by the host with the matching address, while broadcasts are processed by every host that receives the packet. Packets with MAC broadcast addresses may have unicast higher-layer addresses, and the FP forwards the packet or passes it up to the higher layers depending on the packet type and destination address.

```

/* input : packet pkt */
destAddr = pkt->header.DA;
if (destAddr == myAddr || isBroadcast (destAddr) == YES)
    /* pkt for this host */
    /* pass packet to application */
    enQ (pkt);
if (destAddr != myAddr) { /* pkt not for this host */
    if (isBroadcast (destAddr) == YES || pkt->header.type == ROUTESEARCH
        || pkt->header.type == FLOODDATA)
        /* pkt forwarded by DRP for broadcast */
        broadcast (pkt);
    else {
        if (isInRt (destAddr) == YES) {
            /* if the destAddr is in the RT, forward the packet */
            next = nextHop (destAddr); /* get the next hop */
            forward (next, pkt);
        }
        else {
            if (pkt->header.type != ERASE && pkt->header.type != ERROR) {
                if (srcAddr == myAddr) {
                    /* initiate route request. At the completion of this call
                     there will be an entry for the destAddr in the RT */
                    next = nextHop (destAddr); /* get the next hop */
                    if (next != NULL) forward (next, pkt);
                }
                else {
                    pkt->header.type = FLOODDATA;
                    broadcast(pkt);
                }
            }
            else {
                /* route fails; inform the source and drop packet */
                sendRouteError(srcAddr);
                discardPkt(pkt);
                return;
            }
        }
    }
}
}
}
}

```

■ Table 4. Pseudo-code for FP.

The DRP

The DRP is more complex than the FP since it processes incoming packets and maintains the RT and SST. On receiving a packet from the device driver, DRP deciphers the packet type, updates the tables, modifies some of the header fields, and then passes it to FP. The DRP pseudo-code is given in Table 5.

In addition to the packet, the device driver passes the signal strength (sig) at which the packet is received and the address of the host (sender) that transmitted this packet (not the original source). The updateTables() function then updates the signal strength field in the SST according to the following formula:

$$SS_{cumulative} = \alpha \times SS_{cumulative} + (1 - \alpha) \times SS$$

$SS_{cumulative}$ is the quantity recorded in the SST, and SS is the value of average signal strength for the packet supplied by the device driver. α is an experimentally determined constant. updateTables() also marks the Last field in the SST to indicate that a beacon was received from the sending host during the current time click.

When no packet is received during a certain period of time, the device driver may still provide the DRP signal information obtained through its beacon.

Periodically, an asynchronous process runs through the SST comparing $SS_{cumulative}$ to an experimentally determined quantity $SS_{threshold}$. The calculation shown in Table 6 is then carried out to classify hosts into the SC or WC set. A mobile host that exhibits strong signal strength for $clicks_{threshold}$ consecutive clicks is included in the SC set and added to the RT (with itself as the next hop). $clicks_{threshold}$ is another experimentally determined quantity.

The DRP processes each packet depending on its packet type. For route-search packets, isStale() determines whether the packet should be dropped. It is dropped if the packet has previously been seen by this host or if the TTL has expired. The host records route-search packets that it has seen in a table of source address (SA), sequence number (SEQ) pairs. The routine seenRequest() records the pair when the host processes the route-search packet for the first time.

constructRouteSearchForward() modifies the route-search packet by adding the address of the resident host to the hop list to construct a new route-search packet for forwarding.

When the intended destination receives a route-search packet, a route reply is produced by constructRouteReply() and sent back to the source along the reverse path of the route contained in the route-search packet. The next hop for route reply is then installed in the RT and the reply is forwarded. Every intermediate node echoes this procedure (installHop() and constructRouteReplyForward()). installHop() installs all possible routes that are implied by the route reply. For example, Fig. 4 shows a route found from A to G. When host D receives a route reply from G, it

installs routes to A and B (with a next-hop of C) and installs routes to F and G (with a next hop of E). D already has routes for its neighboring hosts, C and E, from beaconing information. (We assume that wireless channels are symmetric; i.e., if the link is SC in one direction, it is also SC in the reverse direction. In cases where the forward link is SC and the reverse link is WC, we would have a less stable route than expected.)

Once the route installation is completed, data packets starting from the source are forwarded along the next hops installed in the RTs of the intermediate nodes. In cases where all hops are along an SC-set path, the packets would be routed quickly and efficiently. If a host cannot forward a packet due to link failure, FP sends a route-error packet back to the source. When the source receives such a route-error packet, it sends a route-erasure packet as constructed by the function `constructRouteErase()` with the DA equal to the destination of the stale route. The stale route is deleted from its RT (`deleteHop()`). Any intermediate host receiving this route-erasure packet forwards it to the next hop and deletes the stale route. If the host is unable to forward the packet, the packet is simply dropped.

Extended Interface

The extended device driver interface provides the `updateTable()` call by which the higher-layer protocols update the SST. The interface allows changes to the time period between beacons. It also allows $SS_{\text{threshold}}$ and $click_{\text{threshold}}$ to be controlled. The $SS_{\text{threshold}}$ determines the extent of the host's coverage area within which the neighbor nodes have strong signals. The $click_{\text{threshold}}$ regulates the impact of location stability considerations on the protocol. It should be determined based on known mobility patterns of hosts in the ad hoc network. In the case where it is set to one, the protocol's routing decisions are based solely on signal strength.

Route Maintenance

Route maintenance is triggered when a host has data to send over a failed link. Intermediate nodes send an error message to the source when such a failure occurs. The source host sends a route-search packet to find a new route and sends an erase message to remove the old route. The erase message should reach the intermediate host which discovered the failed next-hop. Error and erase messages are unicast with best-effort delivery and are dropped if a host is unable to forward the packet. This prevents a cycle of error and erase messages from wasting network resources.

In cases of multiple failures, some routing messages may not reach their destinations. This may result in the existence of stale routes, but it will not cause any routing errors or loops. The stale routes will be discovered and erased by the next data packet that tries to use the invalid route.

```

/* input : signal_info sig, packet pkt */
updateTables(sender, sig);
/* link layer has removed the lower level header from pkt */
if (pkt != NULL) {
    if (isStale(pkt) == YES) {
        discardPkt (pkt);
        return;
    }
    switch (pkt->header.type) {
    case (ROUTESEARCH) :
        if (isFromWC(pkt) == YES &&
            pkt->header.pref == STRONGLINKONLY) {
            discardPkt(pkt);
            return;
        }
        if (isSearchForMe(pkt->header) == YES) {
            outPkt = constructRouteReply (pkt);
            if (outPkt == NULL) break; /* reply not yet constructed */
            else installHop(outPkt); /* update RT */
        }
        else {
            outPkt = constructRouteSearchForward (pkt);
        }
        SeenRequest(pkt->header.SA, pkt->header.seq);
        FP(outPkt);
        break;
    case (ROUTEREPLY) :
        installHop(pkt); /* update RT */
        outPkt = constructRouteReplyForward(pkt);
        FP (outPkt);
        break;
    case (UNICASTDATA) :
    case (FLOODDATA) :
        FP(pkt);
        break;
    case (ERROR) :
        if (pkt->header.DA == myAddr) {
            outPkt = constructRouteErase(pkt);
            FP(outPkt);
            deleteHop(pkt);
        }
        else {
            FP(pkt);
        }
        break;
    case (ERASE) :
        FP(pkt);
        deleteHop(pkt); /* update RT */
    }
}

```


■ Table 5. Pseudo-code for DRP.

```

/* clicks assumed to be initialized to 0 */
if (last not marked) {
    delete entry from SST and RT;
    return;
}
unmark last;
if (SS_cumulative >= SS_threshold) clicks++;
else clicks = 0;
if (clicks == clicks_threshold) {
    set = SC;
    add host to RT;
}
/* if clicks > clicks_threshold, then host is already in the RT */
else set = WC;

```

■ Table 6. Pseudo-code updating SST.



■ **Figure 5.** Failure of a route-erase packet.

If a link failure prevents the route reply from reaching the source, the source will time out and retry the route search. Meanwhile, the intermediate hosts between the failed link and the destination will have incorrect routes to the source. If any of these hosts use these routes, error and erase messages will be generated to correct the routes.

In cases where multiple link failures occur nearly simultaneously, erase and error messages may not reach their destinations. If an error packet cannot be delivered to the source, it must be due to another link failure which will also trigger an error packet closer to the source. This second error packet will inform the source of the second failure so that the source can take appropriate action. If a data packet is en route between two links which fail, as shown in Fig. 5, the resulting error message from E will not reach the source A, since it will be dropped due to the failed link at C. When A sends a data packet again, B will send the error message to inform A of the route failure. The erase message from A will erase the route at B, but the stale routes to G at C and D will only be erased if either C or D tries to send data to G.

If the second link failure occurs after the error message arrives at A, the erase message from A will be dropped at B. This creates the same situation, where C and D have an invalid route to G which will be erased when either C or D sends a packet toward G.

Since intermediate hosts relay route-search packets even if they already know a route to the desired destination, the algorithm works correctly in the presence of stale routes. The source is informed of the error and initiates a new route search. SSA uses route erasure only to avoid wasting resources by forwarding data packets over routes known to be stale. The cost of this simple error and erase method is a few stale routes. Since multiple failure cases are rare, the best-effort unicast of error and erase packets is an effective method of reducing excessive packet transmissions.

Other Issues

Enhanced SSA

The protocol presented in the previous sections provides a basic routing function. However, there are several enhancements to give the hosts more options to deal with varying needs. One is the route quality option (the PREF field in the header), which the source sets when searching for a route. The host may prefer routes with more strong links but not want to exclude routes with any weak links. To accommodate this need, the hosts may implement another PREF option: STRONGPREFERRED. Any intermediate host receiving such a route search should forward it (by broadcast) unless its hop list contains a loop or its TTL has expired, even if the host

has seen this request before (through a different path). The destination host should wait for a period of time to allow several

route-search packets to arrive via different routes. The destination then selects the best one according to a certain criterion, such as shortest path with minimal weak links. The source host may choose this option to find a route if it is probable that no strong route exists.

Another optimization decreases route discovery latency and route-search propagation by allowing intermediate hosts to participate in route discovery. If an intermediate host receiving a route search already has a route to the destination, it may send a route reply immediately back to the source to decrease the latency, in addition to forwarding the route search to the destination. If this route is nonoptimal or stale, it will be overwritten later by a route reply from the destination. The source may decide whether it wants to risk losing a few packets to start transmitting data sooner.

Interoperability with Mobile IP

Assuming that some hosts in the ad hoc network are base stations with both wired and wireless connectivity, SSA can be easily integrated into the global Internet through Mobile IP [14]. These base stations serve as the home and foreign agents for Mobile IP.

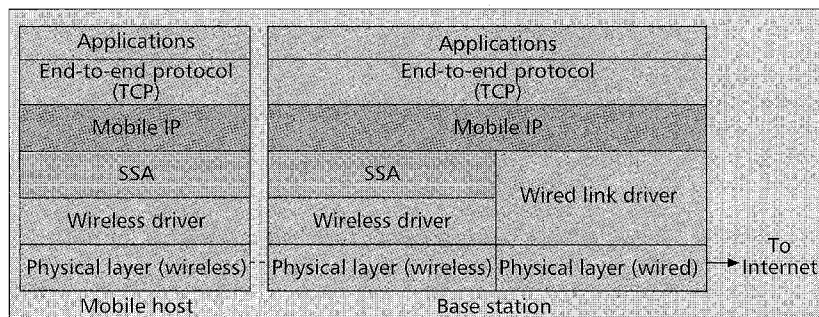
Conceptually, Mobile IP sits on top of SSA (Fig. 6). As far as SSA is concerned, Mobile IP is just another higher-layer protocol. On the other hand, Mobile IP treats SSA as a link-layer protocol, and from the viewpoint of Mobile IP, a base station can "directly" reach any host that resides in the same ad hoc network. SSA will take the responsibility of delivering packets within the ad hoc network. The base station will take care of any encapsulation or decapsulation required by Mobile IP.

Mobile IP on a base station may broadcast its agent advertisement, which should be heard by any potential client, in this case by any host within the ad hoc network. In order to achieve this, all hosts in this ad hoc network would be required to rebroadcast the agent advertisement, which takes a lot of resources. We suggest that a base station broadcast agent advertisements only rarely, or not at all, to conserve the limited bandwidth. If access to the wired network is not always available in the ad hoc network, occasional advertisements may be worthwhile; otherwise, a mobile host should use agent solicitations when it desires the service of a base station.

When sending an agent solicitation, the SSA on the mobile host should send a packet addressed to the SSA broadcast address. Non-base-station hosts will simply rebroadcast the packet, while any reachable base stations will send out a reply if willing to serve as an agent. The base station's reply would trigger SSA to find a route to the base station.

After this handshake, the mobile host can continue with the usual registration process required by Mobile IP and start sending and receiving datagrams to and from the Internet through the base station.

If a mobile host becomes separated from its agent, it may send out a new agent solicitation to find another reachable base station. Alternatively, it may occasionally send agent solicitations to find other base stations while maintaining a connection with the current one. Having this list of base stations would decrease the latency of a base station switch in case the first base station becomes unreachable.



■ **Figure 6.** Integration of Mobile IP and SSA.

Evaluation

We performed simulations to evaluate the benefits and costs of the SSA routing approach. The simulations quantify the length and longevity of the routes determined by SSA under various node densities and mobility rates. They also determine the relative efficacy of using signal strength and location stability as selection criteria for routing.

We studied a large range of cases by varying many of the input parameters, including area, node density, transmission range, rate of topology change, pattern of individual host mobility, session length, and the routing algorithm criteria for route selection. We measured and compared the number of route reconstructions required, the average route hop length, the percentage of strong routes available, and the transmission cost of the SSA routing algorithm. We saw improved reconstruction costs for most sets of parameters. For brevity, only a representative set of simulations are presented in the following sections.


Simulation Setup

Our simulation is based on the environment of a closed 1500 x 1500 unit area in which there are a number of randomly distributed mobile hosts. A signal is considered strong if it comes from a host *strong-radius* units or less away, and two hosts separated by more than *weak-radius* units are considered disconnected. Note that in a real environment, these quantities would be dependent on and controlled by the $SS_{\text{threshold}}$ set through the device driver interface and the capability of the wireless device, respectively. In our simulation, we assume that the signal strength depends solely on the distance between the sending and receiving hosts. If a signal is weaker than that from a host *weak-radius* units away, it is considered noise and dropped at the physical layer. $SS_{\text{threshold}}$ is assumed to be equal to the strength of a signal from a host *strong-radius* units away. Since $SS_{\text{threshold}}$ is experimentally determined, we ran simulations with *strong-radius* 200 and 300 units. The *weak-radius* is kept constant at 400 units because it is a physical quantity dependent on the wireless device.

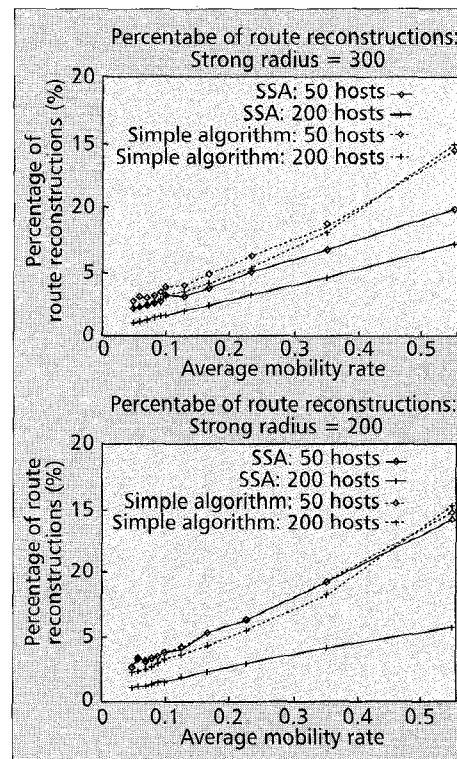
20 percent of the hosts are stationary during the simulation. The other 80 percent move for a number of

clicks and then stay for some number of clicks and then move again, continuing the cycle. The lengths of the moving periods are normally distributed with average 10 clicks and standard deviation 1 click. There are two classes of staying period lengths. The short-stay class is normally distributed with average 3 clicks and standard deviation 1 click, and the long-stay class with average 150 and standard deviation 10. Initially, every nonstationary host is assigned a probability that determines whether it falls into the short-stay or long-stay class each time it enters a staying period. The initial staying probability is chosen from a normal distribution with standard deviation 0.05 and a mean that ranges from 0 to 1 with step 0.1. A simulation is run for each step of the staying probability to obtain data for a range of mobility rates.

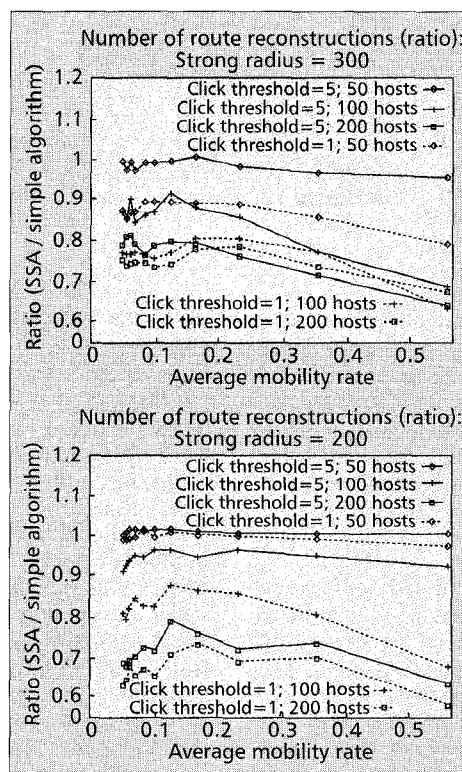
If a host moves during a certain click, it moves 20 units of distance in a randomly chosen direction. If it did not move during the previous click, the direction is chosen from a uniformly distributed random number. Otherwise, the direction is chosen from a normally distributed random number with average equal to the previous direction and standard deviation of 10 degrees. Hence, a host is likely to continue to move in the same general direction as the previous movement. If a host hits the boundary of the area, it will reflect off the boundary at the same angle so that the total moving distance during this time click is still 20 units.



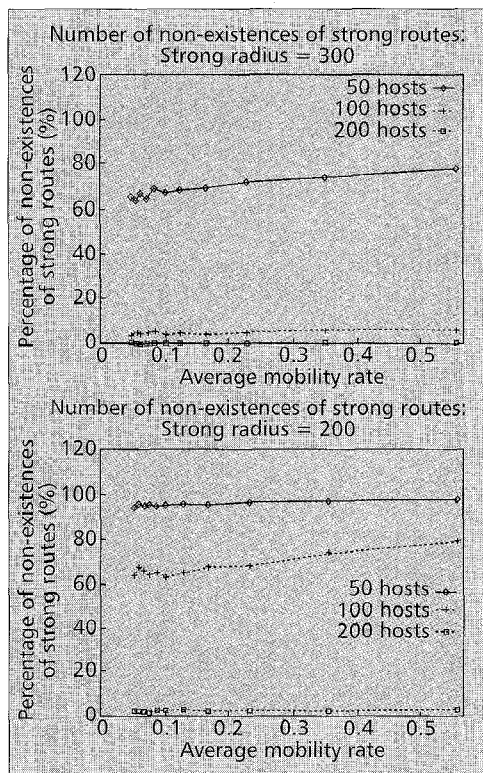
■ Figure 7. Number of route reconstructions: stability considered.



■ Figure 8. Number of route reconstructions: stability not considered.



■ Figure 9. Improvement of SSA over the simple algorithm



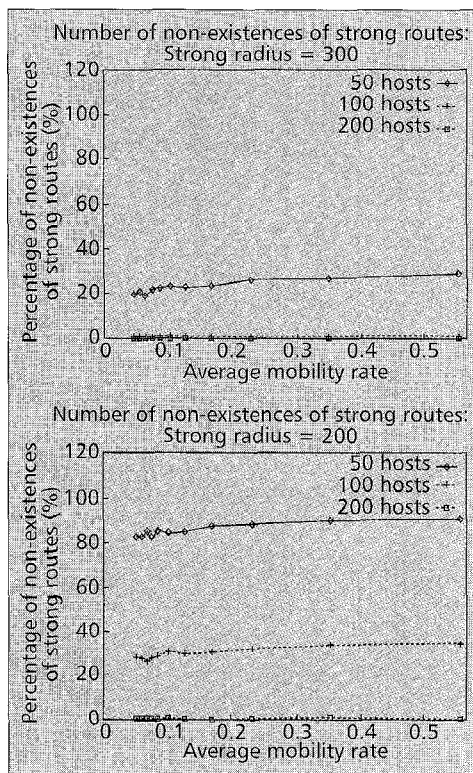
■ **Figure 10.** Probability of no strong routes: stability considered.

We run simulations for networks of sizes 50, 100, and 200 hosts with $click_{threshold}$ of 1 and 5. $click_{threshold}$ determines the threshold above which routes are considered stable. Note that $click_{threshold}$ equal to 1 means that location stability is not considered. When considering location stability, this threshold should be slightly greater than the mean of the short-stay period so that only hosts which have a long-stay period are considered stable. A $click_{threshold}$ of 5 is chosen since the mean of the short-stay period is 3.


During each run, we randomly generate the initial positions of the hosts and let them move and exchange beacon signals for 10 time clicks to establish an initial state. Then we randomly choose a source and a destination and run our algorithm. After each time click, we send a data packet through the established route to trigger route maintenance actions (route erasure and rediscovery) if any of the links failed. After each session (of length 300 clicks), we observe the average number of hops in a route and the number of required route reconstructions. These quantities are averaged over several hundred runs for each combination of the input parameters.

Results

We compare the results from SSA (with and without location stability) with those from an imaginary routing algorithm in which a shortest path is chosen, regardless of the strength of its links. We call the later approach the *simple algorithm*. The performance parameters are plotted against the average mobility rate, which is the average of all the host mobility rates. The mobility rate of a host is the number of clicks during which the host moves divided by the total number of clicks in a session. Clearly, this average mobility rate is inversely dependent on the average initial staying probability. As the



■ **Figure 11.** Probability of no strong routes: stability not considered.



■ **Figure 12.** The buffer zone effect.

mobility rate increases, the number of route reconstructions consistently increases, as shown in Figs. 7 to 9. The number of reconstructions also increases as the number of hosts decreases, due to the increasing sparseness of the topology. In the following graphs, we compare the simple algorithm to the SSA algorithm for the same number of hosts.

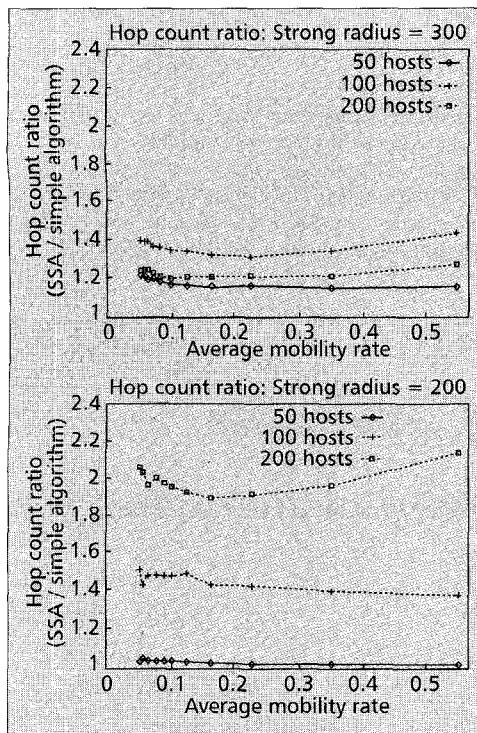
Figures 7 and 8 show that the fraction of sessions requiring route reconstructions is consistently lower for SSA both with and without stability as compared to the simple algorithm. However, SSA with location stability performs worse than that without location stability. At first glance, this is somewhat surprising. However, Figs. 10 and 11 reveal that taking stability into account increases the probability of nonexistence of strong routes considerably. This is because

location stability introduces a much stronger criterion for a link to be SC. If we are unable to find a strong route, route discovery takes longer, and the route is likely to fail sooner. The advantage of SSA arises from the *buffer zone* effect. If an SC link is chosen as part of a route, it will have to become WC before breaking, and this tends to give the individual links, and therefore the entire route, a longer life. As shown in Fig. 12, the buffer zone allows mobile hosts to roam within a certain vicinity of each other without triggering a route reconstruction.


Comparing the SSA curves for 50 hosts in Fig. 8, we see that the performance with the larger strong radius is superior to that of the smaller one. This is somewhat contrary to the buffer zone effect just described since the smaller radius allows the mobile hosts to travel a longer distance in the weak region before the link breaks. Figure 11 offers insight: the percentage of nonexistence of strong routes is more than 80–90 for a *strong-radius* of 200, whereas it is only about 20–30 for a *strong-radius* of 300. The decreased number of strong routes more than offsets any gains due to the increase in the buffer zone.

Clearly, SSA performs well when there are an adequate number of strong routes. This, in turn, depends on the node density (the number of hosts in our environment), the strong radius, the mobility rate, and the criteria defining a strong link. Many combinations of these parameters result in a configuration where SSA drops the number of route reconstructions required. SSA reduces the route reconstructions needed by up to 60 and never performs worse than the simple algorithm. A careful comparison of the SSA curves with and without location stability considerations indicates that in most cases not taking stability into account results in better performance.

Since SSA prefers routes with strong links which are likely



■ Figure 13. Hop counts: stability considered.



■ Figure 14. Hop counts: stability not considered.

to be between two hosts close to each other, we tend to get routes with more hops than with the simple algorithm. On the other hand, strong links are less vulnerable to interference and hence result in less packet loss and corruption. As a rough weight, we count each weak link as 1.25 hops to reflect its vulnerability. Figures 13 and 14 show the hop count ratio between SSA and the simple algorithm. Through the mobility rate range, the hop count ratio usually ranges between 1.1 and 1.5. The drop in the percentage of route reconstructions is more than enough to justify this increase in the average hop count.

Conclusion and Future Work

The SSA protocol proposed in this article focuses on obtaining the most stable routes through an ad hoc network. This approach seeks to maximize the duration of the discovered routes. Our simulations have shown savings of up to 60 in the number of route reconstructions as a result of using signal strength to select routes. Using location stability, on the other hand, is shown to be very sensitive to the particular configuration of the ad hoc network being considered. Since a general ad hoc network is likely to have unpredictable and variable mobility patterns, we propose the adoption of signal strength as a criterion for routing, with configurable parameters to take location stability into account where applicable. We plan to do further simulations using a packet-level simulator to determine the costs and benefits of this approach more accurately and to study the effect of this approach on various transport protocols, including Transmission Control Protocol (TCP).

Although it is intuitively clear that the algorithm is loop-free and converges, we plan to prove these properties using theoretical constructs. We also plan to analyze convergence time and routing overhead for this algorithm.

Acknowledgments

We would like to thank Sambit Sahu (University of Massachusetts, Amherst) and anonymous reviewers for their comments.

References

- [1] G. Malkin, "RIP Version 2 Carrying Additional Information," RFC 1723, 1994, <ftp://ds.internic.net/rfc/rfc1723.txt>.
- [2] D. Bertsekas and R. Gallager, *Data Networks*, Englewood Cliffs, NJ: Prentice-Hall, 1992.
- [3] J. Moy, "OSPF Version 2," RFC 1583, 1994, <ftp://ds.internic.net/rfc/rfc1583.txt>.
- [4] W. R. Stevens, *TCP/IP Illustrated*, vol. 1, Reading, MA: Addison-Wesley, 1994.
- [5] B. M. Leiner, D. L. Nielson, and F. A. Tobagi, "Issues in Packet Radio Network Design," *Proc. IEEE*, vol. 75, no. 1, Jan. 1987.
- [6] J. Rubin and J. D. Tarnow, "The DARPA Packet Radio Network Protocols," *Proc. IEEE*, vol. 75, no. 1, Jan. 1987.
- [7] S. Corson, J. Macker, and S. Batzell, "Architectural Considerations for Mobile Mesh Networking," Internet Draft RFC Version 2, work in progress, May 1996, <http://tonnant.itd.nrl.navy.mil/mmnet/mmnetRFC.txt>.
- [8] C. E. Perkins and P. Bhagwat, "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers," *SIGCOMM Conf. Proc.*, 1994.
- [9] S. Murthy and J. J. Garcia-Luna-Aceves, "A Routing Protocol for Packet Radio Networks," *MOBICOM Conf. Proc.*, Nov. 1995.
- [10] P. Krishna *et al.*, "A Cluster-Based Approach for Routing in Ad-Hoc Networks," *USENIX Tech. Conf. Proc.*, Jan. 1996.
- [11] D. B. Johnson and D. A. Maltz, "Protocols for Adaptive Wireless and Mobile Networking," *IEEE Pers. Commun.*, vol. 3, no. 1, Feb. 1996.
- [12] S. Corson and A. Ephremides, "A Distributed Routing Algorithm for Mobile Wireless Networks," *ACM/Baltzer Wireless Networks J.*, vol. 1, no. 1, 1995.
- [13] C.-K. Toh, *Wireless ATM & Ad-Hoc Networks*, Kluwer, Nov. 1996.
- [14] C. E. Perkins, ed., "IP Mobility Support," IETF, Oct. 1996, RFC 2002.

Biographies

ROHIT DUBE received the M.S. degree in computer science from the University of Maryland, College Park, in 1996 and the B.Tech degree in computer science and engineering from the Indian Institute of Technology, Bombay, in 1994. His research interests include wireless networks, mobile computing architectures, routing protocols, and performance evaluation of transport protocols and network applications. When Rohit is not hacking kernels or tweaking bits, he likes to photograph, travel, and hike.

CYNTHIA D. RAIS received an M.S. degree in computer science in 1995 from the University of Maryland at College Park. She received a B.S. degree in computer science and physics from the College of William and Mary in 1993. Her research interests include mobile and wireless networks, ad hoc mobile network routing, and performance evaluation.

KUANG-YEH WANG received a B.S. degree in mathematics from the National Taiwan University, Taipei, Taiwan, in 1990 and an M.A. degree in mathematics from Brandeis University, Waltham, Massachusetts, in 1994. He is currently in the Ph.D. program at the University of Maryland, College Park. His research interests include wireless networks, mobile computing, routing protocols, and performance evaluation.

SATISH K. TRIPATHI is a professor in the Department of Computer Science at the University of Maryland, College Park. He served as the department chair during 1988–1995. In March 1997, Dr. Tripathi will join the University of California at Riverside as the Dean of Engineering and the Johnson Professor of Engineering. For the last 20 years Dr. Tripathi has been actively involved in research related to performance evaluation, networks, real-time systems, and fault tolerance. In the networking area his current projects are on mobile computing, ATM networks, and operating systems support for multimedia information. He serves on the editorial boards of *Theoretical Computer Science*, *ACM/Springer Multimedia Systems*, *IEEE/ACM Transactions on Networking*, and *IEEE Transactions on Computers*. He has edited books in the areas of performance evaluation and parallel computing.