

Chapter 10

Multipath Routing in Ad Hoc Networks

An-I Andy Wang* Geoffrey H. Kuenning[‡] Peter Reiher*

**University of California, Los Angeles*

[‡]*Harvey Mudd College*

Key words: multipath routing, mobile wireless networks, ad hoc networks

Abstract: In a mobile wireless network, multipath routing provides an effective way to recover from frequent network failures, balance load and energy resources among network nodes, and allow more secure and resilient data transmission. This chapter examines various design approaches to building multipath routing protocols, including handling different failure models, constructing routes of diverse characteristics, building multiple routes with constraints on accessing global knowledge, locating nodes with various approaches, and making decisions on how to recover from failures. We illustrate these approaches through a number of multipath routing protocols. Finally, we conclude this chapter with the most recent advances in multipath routing, along with future challenges.

1. INTRODUCTION

In an ideal network, a source always knows how to reach the destination, and the network connection is always reliable. In a wireless mobile network, or an *ad hoc network*, a source needs to update the location of the mobile destination and intermediate nodes constantly, and network connections may break frequently due to the changing network topologies and unreliable wireless connectivity.

Routing is a major challenge in this wireless mobile environment. Mobility renders standard Internet routing methods (e.g., single-path shortest

route) inappropriate. Typically, ad hoc networks operate on wireless links with limited bandwidth and transmission range, and the nodes constituting the network often operate off batteries, placing a further premium on efficient operations.

One implication of mobility is that frequent location updates make the conventional approach of maintaining source-destination distance tables impractical, since refreshing table entries requires a high rate of updates. In addition, mobility causes many table entries to become out of date before ever being used. Newer on-demand approaches maintain table entries only when a communication session is initiated, thereby reducing the table maintenance overheads. However, such approaches are still inefficient, since mobility can frequently break those on-demand routes and trigger expensive repair mechanisms to reestablish table entries.

Clearly, handling mobility demands protocols that have higher resiliency in the face of rapidly changing network topologies. Such resiliency can be achieved by using multipath routing solutions, which create several redundant routes for a source-destination pair. If one route fails, a backup route will still be available. Combined with on-demand approaches, multipath routing can handle mobility efficiently by tracking intermediate nodes and destinations only when necessary. The combined approach offers a greatly reduced route recovery time when a main route fails.

Multipath routing also offers other advantages. In a conventional network infrastructure, classical multipath routing allows load balancing among multiple routes, reducing network traffic congestion and improving the overall quality of service (*QoS*). Transmitting data through multiple paths in parallel also permits aggregation of network bandwidth. Higher resiliency can be achieved by transmitting data either redundantly or with error-correcting information through separate routes simultaneously.

In the context of ad hoc networking, all the classical applications of multipath routing still apply, but ad hoc multipath routing provides additional benefits. First, in a mobile environment, a pre-established route is likely to break often, and reducing the failure recovery time by having standby alternative routes can significantly affect the *QoS* perceived by end users. Alternating paths to transmit information can also spread the energy use among network nodes and prolong the battery life for the ad hoc network as a whole. In addition, transmitting encrypted data across multiple routes can significantly reduce the likelihood of man-in-the-middle, replay, and eavesdropping attacks. This property is especially important in mobile environments, since wireless communication is inherently more vulnerable to security failures.

This chapter describes the problems involved in constructing multiple paths in wireless mobile networks and surveys the approaches being taken to overcome those problems.

2. DESIGN SPACE FOR AD HOC MULTIPATH ROUTING PROTOCOLS

The typical problems encountered in designing single-path protocols for a mobile wireless environment also arise in designing multipath protocols. Mobility makes it difficult or impossible to maintain a global view of the network. Mobility also has implications in terms of caching policies. If cached information at intermediate nodes is frequently out of date, caching can degrade routing performance because detecting inaccurate routing information is not instantaneous. In addition to the dynamic topology, unreliable and range-limited wireless transmission makes resiliency a requirement rather than an enhancement in a routing solution. Since mobile transmitters are likely to be battery powered, routing protocols need to minimize the communication for coordinating network nodes.

At the protocol level, the design of multipath routing needs to consider failure models, characteristics of redundant routes, coordinating nodes to construct routes, mechanisms for locating mobile destination and intermediate forwarding nodes, and failure recovery.

2.1 Failure Models

Most multipath routing protocols are designed for *independent, isolated failures* in terms of network components. More precisely, each node and link in a route has a probability of failure p_f during some small interval T . The probability of a route failure is defined as the probability of at least one failed component in a route during the interval T . The isolated failure model is quite realistic, especially for hardware components. Redundant routes can handle this failure model gracefully.

Another form of failure is the *geographically localized and correlated failure*, which results in all nodes failing within a circle of a radius R . The choice of R is somewhat arbitrary and depends on the possible causes of failure. This model of failure reflects various environmental factors, such as poor weather conditions or natural disasters. Multipath routing protocols that form routes spanning a large geographical region are more likely to survive this type of failure. Protocols that construct paths that are adjacent to one another usually resort to reconstruction of multiple paths when faced with this type of failure.

2.2 Characteristics of Redundant Routes

The degree to which a multipath routing algorithm succeeds in building useful multiple paths depends not only on the design of the algorithm, but on why multiple paths are wanted. To illustrate various designs, *Figure 1* shows a sample ad hoc network. The dashed lines represent the available wireless links between nodes. The figure represents relative geographical positions, in addition to connectivity. The connectivity varies depending on physical distance, radio characteristics, and environmental conditions. In this figure, the source node *S* wants to send traffic over the network to *D* through multiple paths.




Figure 1. An example of an ad hoc wireless network

To achieve reliability, one possibility is to construct *node-disjoint routes*, where each route constitutes a different set of intermediate forwarding nodes (*Figure 2*). With k node-disjoint routes, a multipath scheme can tolerate at least $k - 1$ intermediate network component failures without disconnecting a source-destination pair. Node-disjoint routes can be geographically adjacent to the shortest path, if a goal is to minimize end-to-end delays and maximize network bandwidth aggregation at the same time. *Figure 2* shows two routes, S-1-2-3-D and S-4-5-D. In this case, multiple routes constructed include the shortest path (S-4-5-D).




Figure 2. Two node-disjoint routes, one of which is the shortest path

Routes can be also widely separated geographically, if the predominant goal is to reduce the probability of multiple routes being disrupted by a single regional, correlated failure. Figure 3 shows a different pair of routes, S-1-2-3-D and S-6-7-8-D, which will require a failure that affects a larger geographical region before both paths can be broken.




Figure 3. Two node-disjoint routes that are separated geographically

Another multipath approach is to form *link-disjoint routes*, in which links are not shared, but intermediate nodes can be shared when constructing multiple paths for a source-destination pair. (Note that node-disjoint routes are automatically link-disjoint.) Link-disjoint routes are not as resilient to geographically localized and correlated failures. However, detecting and repairing a single point of failure can be a more localized operation, because node-disjoint routes need to propagate network failures back to the source before an alternate can be deployed for recovery. In contrast, an alternative route in link-disjoint schemes can be set up by an adjacent node in the upstream direction. In an energy-constrained environment, constructing and maintaining link-disjoint routes can be more energy efficient than maintaining node-disjoint routes, because link-disjoint routes tend to be adjacent to the primary route, which is often the shortest route. On the other

hand, node-disjoint routes tend to cover a wider geographical region and spread battery consumption more evenly throughout the network.




Figure 4. Two link-disjoint routes (S-4-5-D and S-6-4-2-3), with node 4 shared

Figure 4 shows two link-disjoint routes, S-4-5-D and S-6-4-2-3-D. If the main route S-4-5-D fails due to node 5, the route recovery process can be handled at node 4 to reestablish the route S-4-2-3-D. If an additional link failure occurs at the link between node S and node 4, node S will reestablish the route S-6-4-2-3-D. On the other hand, if node 4 fails first, both routes can be destroyed in a single failure. In general, k link-disjoint routes do not provide as much resiliency as k node-disjoint routes. However, since the total number of link-disjoint routes in a given network almost always well exceeds the number of node-disjoint routes, a sufficient number of link-disjoint paths can achieve the resiliency of node-disjoint paths, assuming independent failures of network components.

A more relaxed form of the multipath routing scheme is to construct *partially disjoint routes*. The formation of partially disjoint routes is mostly for reliability purposes. The primary route is used for data transmission, while other partially disjoint routes are standby routes for failure recovery.




Figure 5. Partially disjoint routes (S-4-5-D and S-4-2-3-D), with link S-4 shared

Figure 5 shows a primary route (S-4-5-D) with one partially disjoint route (S-4-2-3-D). If node 5 fails, the primary route is recovered by switching to the partially disjoint alternative.

2.3 Construction of Multiple Paths

Multiple paths can be constructed by relying on global knowledge, incomplete global knowledge, or only local knowledge.

Routing protocols that rely on global knowledge allow a route-forming node to access the current status of all network nodes and links. At the time of formulating multiple paths, the node can produce multiple routes in a centralized manner. Obtaining such information requires an abundance of local node resources (unlikely for wireless mobile nodes), or the existence of a somewhat centralized database (which is difficult to scale). In any large-scale network, obtaining global knowledge about the entire network is very difficult and costly, even disregarding the additional constraints of wireless mobile environments. For that reason, wireless multipath routing methods avoid any reliance on global knowledge.

At the other extreme, relying on local knowledge means that a source constructs multiple paths in a distributed manner that is largely based on access to the current status of neighboring nodes and links. Relying solely on neighboring information may seem insufficient for coordinating nodes to form multiple paths. However, since wireless communication is broadcast-based, a node can potentially overhear information from neighboring nodes. In addition, certain communication can be implicit. Although routing protocols that rely on local knowledge can scale well due to their distributed nature, constructing multiple routes based on localized coordination is generally difficult because nodes need to determine locally whether global invariants are met (e.g., disjointness of routes, loop-free routes).

Thus, the majority of multipath schemes rely on incomplete global knowledge because of the ease of constructing and verifying centralized solutions. Incomplete global knowledge can be easily obtained by constrained flooding of the network. However, a few multipath schemes are distributed and rely on local knowledge gained through periodic exchange of information with neighboring nodes or by overhearing information through the promiscuous mode.

Without complete global information, one greedy approach to constructing multiple paths is to apply variants of the Dijkstra or Bellman-Ford pair-wise shortest-path algorithms iteratively. The partial network topology is represented as an undirected graph $G = (V, E)$ with wireless mobile nodes as vertices V , and network links as edges E . Each iteration on G will yield a pair-wise shortest path, which is added to the list of multiple

routes and removed from the original G before the next iteration. After k iterations, this approach will yield a greedy solution of the k shortest node-disjoint routes, which is useful when multiple routes are needed to maximize network bandwidth. If, after each iteration, the nodes in the shortest path are not removed from G , the same method will yield the k shortest link-disjoint routes. Multiple routes with different characteristics (e.g., minimal congestion) can be obtained by adjusting the lengths of the edges according to various constraints. For example, to avoid congestion, the length of edges can be increased for congested links.

Most well-known distributed or localized multipath routing protocols are inspired by biological and physical models, such as ants [Braginsky and Estrin 2001], water flow [Park and Corson 1997], and diffusion models [Estrin et al. 1999; Intanagonwiwat et al. 2000]. In the water flow method, for example, a source can send network traffic through multiple paths to a destination by properly defining heights of *terrains* at each intermediate forwarding node. Whenever the traffic is trapped within a section of terrain (an intermediate node), the terrain is modified so traffic can flow outwards again.

Distributed multipath schemes can be elegant and scale well, and each node can perform routing by keeping only local state. However, these schemes need to overcome the challenge of oscillations. For example, multiple nodes can independently trigger mechanisms for network detection, recovery, erasing routes, and creating routes, resulting in unstable network routes that change frequently due to both mobility and node behaviour. Although centralized multipath approaches are more intuitive to construct and verify, each route typically has a predefined set of forwarding nodes. In ad hoc networks, such predefined routes can be easily broken due to high mobility.

2.4 Location Discovery

Before forming multiple routes in a mobile environment, a source node needs to approximate the current state of the network so it can locate the destination and intermediate forwarding nodes. Location discovery methods for ad hoc multipath protocols are largely based on single-path protocols for such networks.

Proactive approaches actively maintain variants of routing tables for each source-destination pair. Commonly, these routing tables are organized in a hierarchy for scaling. The advantage of proactive approaches is that a source can immediately use local or nearby tables to construct multiple routes. However, mobility renders proactive approaches impractical, since keeping these distributed tables up to date requires high messaging overhead.

Constructing multiple routes on top of these distributed tables also means that a source-destination pair might have multiple table entries and thus higher storage overhead. Multiple entries are more likely to become inconsistent and lead to broken paths and routing loops. In addition, mobility can cause many table entries to become out of date before ever being used, and those unused entries waste precious storage resources and update efforts on resource-poor mobile nodes.

Reactive, or **on-demand approaches** flood the network right before a source initiates a communication session with a destination, so that only states of active routes are being maintained in the network. The flooding process also allows a source to update its view of the network to construct multiple routes. Since flooding mechanisms become prohibitively expensive as they scale up, most on-demand schemes impose constraints on flooding. For example, under certain conditions a node can decide to drop redundant route requests as opposed to forwarding them [Lee and Gerla 2001]. A source can also enlist the help of a geographic location service such as the global positioning system (**GPS**). With knowledge of the destination's prior location and mobility characteristics, a source can limit the flood to an area where the destination is likely to be located [Pei et al. 2000]. However, these constraints also imply that the source may obtain only a partial view of the network state, resulting in the building of multiple routes that are not as effective in achieving a particular set of goals.

A hybrid of the proactive and reactive approaches is also possible. For example, in a large ad hoc network, small groups of nodes can use proactive routing for neighbouring nodes. Each group can elect a head node to represent the group, so that reactive routing is performed among head nodes. This approach is based on two observations. First, the information from distant nodes is less likely to be correct. Therefore, an on-demand approach to obtain remote information is more appropriate. Also, nodes within the same vicinity are more likely to communicate among themselves, so maintaining the complete state at that scale may not differ from the cost of flooding initiated at the beginning of each communication session.

A reverse composition is also possible. The reasoning is that since head nodes experience relatively less mobility due to the distance between them, a proactive approach can work well. Lower-level nodes can use on-demand routing to accommodate frequent and local topological changes.

2.5 Route Recovery

Since one of the major motivations for having multiple routes is to reduce route recovery overhead, the route recovery process for multipath routing protocols is relatively simple.

The failure recovery approach largely depends on how redundant routes are used. If redundant routes are used as backup routes, failure recovery simply means that one of the backup routes will perform normal data delivery. On the other hand, if multiple routes are being used for achieving a certain QoS (such as bandwidth), the failure of one route will prompt the construction of a new redundant route, while the remaining routes can provide a graceful degradation of QoS. If multiple routes are used for load balancing and congestion control, the lack of response from one route may mean that a network is overly congested at the moment. Reestablishing a new redundant route may not be a desirable choice. Fortunately, as long as a source-destination pair has a well-defined priority for a given connection, the decision on the recovery action should be straightforward.

3. EXAMPLES OF AD HOC MULTIPATH ROUTING APPROACHES

Ad hoc multipath routing has generated considerable research interest over the past decade, and a number of algorithms have been developed to address the problem. We will summarize major classes of multipath approaches and consider a few examples of each. Since most ad hoc multipath routing protocols are extensions of single-path approaches, we will start with a brief review of single-path routing protocols.

3.1 Review of Ad Hoc Single-Path Routing Protocols

Many ad hoc multipath routing protocols are direct descendants of two popular single-path approaches: dynamic source routing (*DSR*) [Johnson and Maltz 1996] and ad hoc on-demand distance vector (*AODV*) [Perkins and Royer 1999]. We will briefly review each scheme.

Both DSR and AODV are on-demand approaches and establish routes as needed. Under DSR, a source locates a destination via flooding. Duplicate route-request messages are discarded at intermediate forwarding nodes. Once the destination is located, the destination will respond to the first request message and use the path recorded in the request packet to acknowledge back to the source. DSR uses the *source routing protocol*, in which the source precomputes the entire communication route, and the routing information is encoded in each packet header being transmitted.

One major advantage to using source routing is that intermediate nodes can perform *stateless forwarding*, in the sense that each intermediate forwarding node maintains no state regarding the routes being forwarded. Therefore, the overhead of forwarding is not as sensitive to the size of the

network. However, DSR does use caching to speed up the process of locating a destination. As mobility increases, caching contributes negatively because cache entries are often invalid. Stale routes, if used, may start polluting other caches [Li et al. 2000]. Also, in the case of network failures, new routing requests and associated flooding are required to recover the route.

AODV uses a table-driven approach instead. AODV uses the same on-demand flooding and route recovery mechanisms as DSR. However, each node maintains a routing table that lists the next hop for each reachable destination for each active route. A sequence number is associated with each entry to prevent routing loops. Periodic beaconing is required to keep those tables up to date.

3.2 Extensions of Ad Hoc Single-Path Routing Protocols

DSR and AODV have been modified in various ways to provide multiple routes. The key observation is that during the on-demand flooding phase, enough information can be gathered to form redundant routes without additional overhead. DSR and AODV can be extended to provide multiple routes by relaxing route-request broadcasting constraints during the flooding phase and aggregating network states at the source, the destination, or intermediate nodes. With additional knowledge of the network states, a node can make a more informed decision regarding disjoint route construction.

The *diversity injection* approach [Pearlman and Haas 1998] modifies DSR to compute multiple routes. The key observation is that the flooding process is typically constrained; therefore, replies to multiple route requests at the destination tend to produce routes that share many links. One potential fix is to relax the flooding constraint of the request messages to discover more routes, but the traffic produced by flooding is prohibitive. Since each route request contains a potentially different route back to a source, caching recent requests can build a library of routes back to different sources. When a destination replies toward the source, an intermediate node can inject diversity by probabilistically selecting a route from the library to the source.

Nasipuri and Das [1999] propose an on-demand multipath scheme that modifies the destination under DSR—causing it to reply to multiple route requests selectively. The goal of this approach is to construct partial disjoint routes, where alternate routes are connected from various nodes on the primary route to the destination. *Figure 6* shows an example. The primary route is S-4-5-D, and partially disjoint routes are S-4-2-3-D and S-6-7-8-D. If the primary route is broken at node 5, once node 4 detects the failure, it will alter packet headers to replace the primary route with the alternate route

S-4-2-3-D. This process continues until all routes break; then a fresh route discovery is initiated. Although the intent of the design is to construct many alternate routes along the primary route, the quality of routes constructed is largely a function of the thoroughness of flooding. Also, a failure point near the source will render many downstream alternate routes unavailable.




Figure 6. Multipath protocol proposed by [Nasipuri and Das 1999]

Split multipath routing (SMR) [Lee and Gerla 2001] is another on-demand, ad hoc multipath routing protocol using source routing. SMR is similar to DSR except that SMR tries to increase the probability of finding more disjoint routes during the route discovery phase by avoiding the use of cached routes and relaxing the constraint on forwarding duplicate route-request messages during the flooding phase. The choice of routes is based on the minimal overlapping of nodes and links among routes.

SMR-GPS [Prier et al. 2002] enlists the aid of GPS information to improve the disjointness of backup routes in SMR. SMR-GPS maximizes the minimal pair-wise distance between nodes within a route. SMR-GPS can outperform SMR in terms of surviving geographically localized and correlated failures.

Ad hoc on-demand multipath distance vector (*AOMDV*) [Marina and Das 2001] extends AODV by adding mechanisms to detect link-disjoint routes. A source can initially send different versions of a route request to each of its neighbor nodes. Based on the version stamp of the route request messages, a destination or an intermediate forwarding node can deduce the number of potential disjoint routes to the source. Based on the hop count of the route request, an intermediate node can decide whether to rebroadcast a certain version of the route request. Multiple routes are built incrementally during the forwarding process.

One challenge to modifying the existing AODV algorithm to support multipath routing is avoiding routing loops. The difficulty lies in the distributed storage of network states; maintaining a loop-free invariant is difficult because tables may become inconsistent. In addition, since each source-destination pair may contain multiple routes, a node may reach a

destination through different hop counts, which further complicates verifying the correctness of the algorithm. Marina and Das [2001] deal with this problem by ensuring that nodes receiving a route request only forward it if its hop count is lower than the hop count of any route request they have already received. This requirement ensures that a route request looping back on a node will not be forwarded again, though it may also prevent forwarding of non-looping routes. Other approaches rely on multiple entries in the routing table and version numbers on the route requests. These approaches must contend with difficult bookkeeping issues.

3.3 Other Ad Hoc Multipath Routing Protocols

The temporally ordered routing algorithm (*TORA*) [Park and Corson 1997] maintains a destination-oriented directed acyclic graph (DAG) to construct multiple routes. The use of a DAG assures that the algorithm is loop-free. TORA uses a height-based algorithm, with traffic flowing like water from the source to the destination through multiple paths. When traffic is trapped within the terrain, the terrain is modified so traffic can flow again. The use of this gravitational model enables TORA to compute multiple routes in a distributed fashion. Its localized computation allows TORA to scale and be responsive to changes in dynamic topologies. However, TORA may potentially encounter oscillations of multiple routes, especially when multiple sets of coordinating nodes are concurrently detecting partitions, erasing routes, and building new routes based on each other [Royer and Toh 1999]. Also, TORA needs to flood the network to erase invalid routes due to proliferation of states. In addition, the assumption of reliable, in-order delivery of routing control messages imposes high overheads [Broch et al. 1998]

Directed diffusion [Estrin et al. 1999; Intanagonwiwat et al. 2000] is designed in the context of sensor networks, where minimizing energy consumption is a top priority. Unlike conventional routing approaches, the diffusion model is a data-centric and application-specific approach to directing data from sources to destinations, or sinks. A sink may disseminate its interest in data with certain attributes. Nodes that have the data of interest or information on how to obtain the data will backtrack the trail of interest to the sinks. A group of sensor nodes can cluster, and nodes rotate roles to allow batteries to recharge. The motivation for this diffusion model is the use extensive caching to avoid end-to-end communication, thus prolonging the battery life of individual nodes and the life of the sensor network as a whole. However, the energy consumption is traded off against the storage needed to cache data, and the effectiveness of data duplication and caching is highly dependent on the mobility of sensor nodes.

4. RECENT ADVANCES IN AD HOC MULTIPATH ROUTING PROTOCOLS

Existing multipath approaches often involve storing additional state. Keeping distributed state consistent is usually complex. Therefore, recent advances try to move toward distributed schemes where decisions are made with local knowledge. Ideally, no coordination is required to build multiple disjoint paths in parallel, and each forwarding node performs stateless routing for better scaling. In addition, recent advances are also more energy-aware.

4.1 Braided Multipath Routing

Braided multipath routing [Ganesan et al. 2002] was developed in the context of sensor networks, stressing energy conservation. A shortest alternative path is created for each node in the primary path, resulting in braided paths (*Figure 7*). Location discovery is through low-rate dissemination of source and destination information throughout the network.




Figure 7. Braided multipath routing

Each node can use the promiscuous mode to overhear nearby routing information and form local detours around the nodes in the primary route. The total number of routes between the source and destination is proportional to the n^{th} Fibonacci number, where n is the number of nodes in the primary path. With a large number of alternative routes, the primary path under this approach can potentially sustain many independent failures.

Figure 7 shows an example of braided routing with one primary route (S-4-5-D) and two partially disjoint routes (S-4-2-3-D and S-6-7-5-D). If node 5 fails, the primary route will fall back to the alternate route S-4-2-3-D; if node 4 fails, the primary route will fall back to the alternate route S-6-7-5-D. However, if any of two neighboring nodes fail simultaneously, braided

routes can no longer rely on alternate routes for recovery. Therefore, braided multipath routing is not resilient to geographically localized and correlated failures.

4.2 Magnetic-Field-Based Multipath Routing

Magnetic-field-based multipath routing (MFR) [Nguyen et al. 2002] is an on-demand protocol that exploits the shape properties of magnetic force lines to build node-disjoint paths (*Figure 8*). For each communication pair, a source represents the positive pole, and a destination represents the negative pole. Multiple paths are formed by following or approximating a designated set of magnetic force lines going from the positive pole to the negative pole. By choosing field lines with different initial angles at the source node, MFR can control the distance between disjoint paths. An angle of 0 degrees represents the straightest and (most likely) the shortest path from the source to the destination. Destination discovery is based on flooding, and MFR assumes the assistance of GPS to identify the source and destination locations.

MFR is quite different from the foregoing approaches. Since knowing the position information of the source, the destination, and the node itself is sufficient to compute the direction of a magnetic force line, no explicit control messages are needed to coordinate the formation of multiple routes. Although each node makes local decisions to forward traffic, constructed paths are likely to be node-disjoint. In addition, each node can perform stateless forwarding without maintaining information for each route.

Figure 8 shows a pair of communicating nodes using MFR. The three routes shown are based on magnetic field lines with initial angles -60, 0, and 60 degrees. For independent failures, the multiple paths can serve as alternate routes. The node-disjoint routes formed under MFR can also tolerate geographically localized and correlated failures. However, since the disjoint routes can be significantly longer than the shortest route, energy consumption may be suboptimal. In terms of mobility, each route in MFR has no fixed set of nodes; therefore, the node membership for each route can change dynamically without breaking the multiple routes. For example, nodes 5 and 8 can exchange positions to form new routes S-4-8-D and S-6-7-5-D without affecting the multiple routes.




Figure 8. Magnetic-field-based multipath routing

4.3 Future Challenges

As we can see, no single routing approach can currently address all of the various requirements of ad hoc networks—resiliency, energy constraints, stability, scaling, and extreme density and mobility of nodes. As a result, there is not yet a consensus on the proper way to perform multipath routing in ad hoc mobile wireless networks. Substantial research remains on either finding a better alternative protocol than those already devised or making key improvements to an existing protocol.

Even within the range of existing protocols, some important issues are still inadequately addressed. For example, many ad hoc routing protocols face particular difficulties when the radio transmission characteristics of the environment are difficult. Urban areas with many nodes located indoors or distributed in tunnelled areas are some examples. In these cases, the routes actually available may be rather serpentine, and not all protocols are capable of finding even one of them, let alone several.

Few of the existing approaches have considered security, since getting basic services deployed for the mobile wireless environment is already challenging, and security mechanisms often impose high overhead. However, the security requirements of the ad hoc wireless environment are more challenging than those in a standard wired network, making security correspondingly more important. Attacks on wireless networks are becoming increasingly popular, based on the insecurity of commonly used protocols. The protocols proposed for ad hoc networks are not substantially more secure than wired protocols, particularly when compromised network nodes are participating. Whether there are special security issues related to building multiple paths, rather than a single path, remains to be seen. Certainly the design of the protocols should guarantee that in spite of the

actions of malicious participants, good paths will still be successfully constructed.

Advances in ad hoc multipath routing are moving at a rapid pace, in anticipation of the increasing need for such networks. A consensus regarding protocol design requirements will probably be reached in the near future. Much more research remains to be done in this area to ensure that the eventual choice of multipath routing algorithms for mobile ad hoc wireless networks is a wise one.

5. REFERENCES

- [Braginsky and Estrin 2001] Braginsky D, Estrin D. Rumor Routing Algorithm for Sensor Networks. *International Conference on Distributed Computing Systems (ICDCS-22)*, November 2001.
- [Broch et al. 1998] Broch J, Maltz D, Johnson D, Hu YC, Jetcheva J. A Performance Comparison of Multi-Hop Wireless Ad Hoc Network routing Protocols. *Proceedings of the IEEE/ACM MOBICOM*, 1998.
- [Estrin et al. 1999] Estrin D, Govindan R, Heidemann J, Kumar S. Next Century Challenges: Scalable Coordination in Sensor Networks. *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking*, 1999.
- [Ganesan et al. 2002] Ganesan D, Govindan R, Shenker S, Estrin D. Highly Resilient, energy Efficient, Multipath Routing in Wireless Sensor Networks. *Mobile Computing and Communications Review (MC2R)* 1(2), 2002.
- [Intanagonwiwat et al. 2000] Intanagonwiwat C, Govindan R, Estrin D. Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking*, August 2000.
- [Johnson and Maltz 1996] Johnson DB and Maltz DA. Dynamic Source Routing in Ad Hoc Wireless Networks. *Mobile Computing*, edited by Tomasz Imielinski and Hank Korth, Chapter 5, pp. 153-181, Kluwer Academic Publishers, 1996.
- [Ko and Vaidya 1998] Ko YB, Vaidya N. Location-Aided Routing (LAR) in Mobile Ad Hoc Networks. *Proceedings of ACM/IEEE MobiCom*, October 1998.
- [Lee and Gerla 2001] Lee SJ, Gerla M. Split Multipath Routing with Maximally Disjoint Paths in Ad Hoc Networks. *IEEE International Conference on Communications*, 2001.
- [Li et al. 2000] Li J, Jannotti J, De Couto DSJ, Karger D, Morris R. A Scalable Location Service for Geographic Ad Hoc Routing. *Proceedings of ACM MOBICOM*. August 2000.
- [Marina and Das 2001] Marina MK, Das SR. On-Demand Multipath Distance Vector Routing in Ad Hoc Networks. *Proceedings of IEEE International Conference on Network Protocols (ICNP)*, 2001.
- [Nasipuri and Das 1999] Nasipuri A, Das SR. On-Demand Multipath Routing for Mobile Ad Hoc Networks. *Proceedings of the IEEE International Conference on Computer Communication and Networks (ICCCN'99)*, October 1999.
- [Nguyen et al. 2002] Nguyen N, Wang AI, Reiher P, Kuenning GH. Magnetic Field Routing for Ad Hoc Networks. Submitted to the *First Workshop on Hot Topics in Networks (HotNets-I)*, October 2002.
- [Park and Corson 1997] Park VD, Corson MS. A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks. *Proceedings of IEEE INFOCOM*, April 1997.

- [Pearlman and Haas 1998] Improving the Performance of Query-Based Routing Protocols Through "Diversity Injection." *IEEE Wireless Communications and Networking Conference 1999 (WCNC'99)*, September 1999.
- [Pei et al. 2000] Pei G, Gerla M, Hong X, LANMAR: Landmark Routing for Large Scale Wireless Ad Hoc Networks with Group Mobility. *Proceedings of IEEE/ACM MobiHOC 2000*, pp. 11-18, August 2000.
- [Perkins and Royer 1999] Perkins CE, Royer EM. Ad Hoc On Demand Distance Vector Routing. *Proceedings of IEEE WMCSA'99*, pp. 90-100, February 1999.
- [Prier et al. 2002] Prier G, Schnaider M, Robinson M. SMR-GPS: Split Multipath Routing with GPS Based Route Selection. Unpublished manuscript, 2002.
- [Royer and Toh 1999] Royer EM, Toh CK. A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks. *IEEE Personal Communications*, April 1999.