A Routing Protocol for k-hop Networks†

William D. List
Department of Computer Science, and
Coordinated Science Laboratory
University of Illinois at Urbana–Champaign
Email: list@crhc.uiuc.edu

Nitin H. Vaidya
Department of Electrical & Computer Engineering, and
Coordinated Science Laboratory
University of Illinois at Urbana–Champaign
Email: nhv@crhc.uiuc.edu

***Abstract*— Recent years have witnessed the widespread deployment of IEEE 802.11 LANs in areas such as airports, campuses, and enterprises. These networks allow users to access network services and the Internet in remote locations and without the need for wires. The data rates for 802.11a, b and g far surpass that of wide-area cellular networks, however, the range of transmission of 802.11 is much less than that of cellular (250m versus 20km). Employing ad-hoc mode in 802.11 can extend traditional WLANs to multiple hops, thus increasing coverage and reducing the need for additional infrastructure. The amount of network extension (in terms of wireless hops) is limited by the density of the network (i.e., the availability of wireless devices that can serve as relays for other devices), and the scalability of stand-alone wireless ad-hoc networks. In this paper, we introduce a k-hop architecture and routing protocol utilizing a "beaconing" approach for route discovery and maintenance. We demonstrate through simulations the efficiency and reliability of our routing protocol in the presence of mobility and high node density.**

I. Introduction

Mobile Ad-Hoc Networks (MANETs) consist of a collection of mobile nodes that act in a distributed fashion without an established infrastructure. Each node in the network serves as a router for forwarding packets on behalf of other nodes, supporting multi-hop communications. MANETs are quickly becoming popular, with many potential applications already identified [1]. However, these wireless networks alone provide limited capacity, making widespread deployment with many users difficult [2], [3], [4]. Hence, MANETs are best-suited for operation within a limited range of hops, beneath the so-called *ad-hoc horizon* [5]. An attractive feature of MANETs is their ability to construct on-the-fly networks and adapt to changing conditions. This feature makes them a good candidate for the logical extension of fixed network infrastructures, such as WLANs. Extending the range of WLANs beyond a single hop has many benefits, including 1) increased coverage of the network, 2) more opportunities for path selection and in turn better signal quality, and 3) decreased deployment costs. Implementing an ad-hoc network on top of each wireless access point can benefit a large number of users that are just outside the reach of the WLAN. This concept is already being explored in other areas, for example to allow vehicles on the highway to send data back and forth and to the Internet via Dedicated Short-Range Communication (DSRC) radios.

† This work was supported in part by Motorola, Inc.

In this paper, we introduce our k-hop network architecture, comprising a fixed backbone network of gateways (GWs) that is extended by a k-hop MANET. Our main contribution is the presentation and analysis of our novel "beaconing" routing protocol. We describe and test the routing protocol in a variety of settings and draw conclusions on its design and performance that can aid in the development of future protocols.

II. Related Work

Since the advent of ad-hoc networks, researchers have explored ways to incorporate them with other network architectures to achieve various goals. For example, a considerable amount of literature exists regarding ways to integrate ad-hoc and cellular networks [6], [7], [8], [9], [10], [11], [12]. While the idea of combining 802.11 and wide-area cellular networks is interesting, ultimately it requires the need for two radios in each device, leading to a more expensive solution. Our architecture assumes that the only network devices that have multiple NICs are the gateways, keeping the complexity and cost of mobile devices low.

Similar work to our own first appeared early on in [13]. Super Mobile Hosts (Super-MHs) form a backbone network for Mini Mobile Hosts (Mini-MHs) to communicate across long distances. Routing between Super-MHs is performed using a separate channel than that of Mini-MHs. The route discovery scheme used by Mini-MHs attempts to find either (1) the destination directly, or (2) the nearest Super-MH.

Recent work done by Gerla *et al.* is similar in nature to our k-hop network model. The design methodology presented in [14] uses Landmark Ad Hoc Routing (LANMAR) [15] and a clustering technique to build hierarchical networks. An algorithm is employed to select a subset of mobile nodes to become Backbone Nodes (BNs), forming a backbone network. A primary difference between [14] and our k-hop architecture is that nodes in a k-hop network do not form clusters around the GWs. In addition, the GWs within a k-hop network are not dynamically chosen, and are assumed to be in fixed locations.

Another architecture that achieves similar same goals as ours is the Lightweight Underlay Network Ad-Hoc Routing (LUNAR) [5]. LUNAR implements a layer in-between the MAC layer and the IP layer to perform a variation of multi-hop ARP. The protocol is simple and is designed to work in small-hop environments. LUNAR uses a combination of re-activity and pro-activity for route discovery and maintenance.

A special *gatewaying node* handles passing packets back and forth between the Internet and the LUNAR network. Our architecture differs from LUNAR in the following ways — 1) We do not support multi-hop ARP, and use DHCP for address assignment rather than a randomized scheme (see [16]), 2) In LUNAR, broadcasts are performed using a tree; we use a more traditional broadcast scheme, and 3) We employ GW "beaconing" to identify MHs in the immediate area, LUNAR does not have any discovery process.

A new hybrid version of AODV called Hierarchical AODV (H-AODV) is presented in [17]. H-AODV performs RREQs at two separate levels, and improvements in path lengths are possible by relaying through BNs. H-AODV uses an additional field in the routing control packets to define subnets around each BN. Thus, all packets outside the local subnet must pass through the associated BN. This is not the case in a k-hop network, where nodes are able to locally route to anyone within the k-hop bound. A second difference between H-AODV and k-hop networks is that in our architecture nodes are able to utilize one *or more* GWs for routing whereas in H-AODV, each node is associated with only one BN at a time.

III. The k-hop Architecture

In this chapter, we describe in detail the fundamentals of our k-hop architecture.

A k-hop network comprises a fixed network of gateways (GWs) that support mobile hosts (MHs) via an ad-hoc network protocol. Each MH and GW is equipped with a wireless interface that allows it to communicate with other MHs and GWs over the wireless channel. Each GW possesses a second interface (wired or wireless) that allows it to talk to the backbone network. k in a k-hop network refers to the upper bound on the number of wireless hops in all ad-hoc connections. For example, in a 2-hop network paths between two MHs or a MH and GW are limited to 2 wireless hops. For a MH M, a GW or MH that is within the k bound of M is considered *local* to M. The remaining MHs in the network for M are *remote*, and must be accessed via one or more GWs. Thus, a k-hop network allows at most $2k$ wireless hops on any route.

Figure 1 shows a simple 2-hop network with 2 GWs (black circles) and 8 MHs (white circles). Dotted lines indicate connectivity between MHs and GWs. The backbone link between GW **G1** and **G2** is not shown. The following paths in the network are valid. $\mathbf{A} \Rightarrow \mathbf{D}$[1] could use $\mathbf{A} \rightarrow \mathbf{G1} \rightarrow \mathbf{B} \rightarrow \mathbf{D} \mid 3$[2] or $\mathbf{A} \rightarrow \mathbf{C} \rightarrow \mathbf{G2} \rightarrow \mathbf{E} \rightarrow \mathbf{D} \mid 4$. $\mathbf{E} \Rightarrow \mathbf{C}$ could use $\mathbf{E} \rightarrow \mathbf{G2} \rightarrow \mathbf{C} \mid 2$ or even $\mathbf{E} \rightarrow \mathbf{B} \rightarrow \mathbf{G1} \rightarrow \mathbf{A} \rightarrow \mathbf{C} \mid 4$ if necessary. Third, $\mathbf{G} \Rightarrow \mathbf{H}$ is limited to the path $\mathbf{G} \rightarrow \mathbf{B} \rightarrow \mathbf{G1} \rightarrow \mathbf{G2} \rightarrow \mathbf{C} \rightarrow \mathbf{H} \mid 4$. Notice that MH **F** is not able to reach MH **A** ($\mathbf{F} \Rightarrow \mathbf{A}$), since its shortest path to any GW is 3 hops which is above the k bound.

MHs take advantage of all nearby GWs to maintain connectivity. Using a unique "beaconing scheme", our routing protocol is designed to pro-actively ensure that a valid path to at least one GW is always known (if such a path exists).

[1] $\mathbf{S} \Rightarrow \mathbf{D}$ is read as "Source **S** connecting to destination **D**."
[2] $\mathbf{S} \rightarrow \mathbf{A} \mid 1$ stands for "**S** forwarding to **A** over 1 wireless hop."

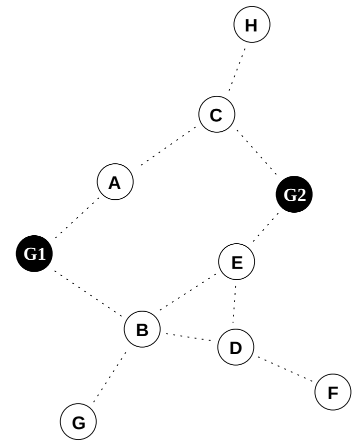

Fig. 1. A simple 2-hop network.

IV. Routing in k-hop Networks

The architecture of Section III calls for a routing protocol that is able to exploit the provided infrastructure when making routing decisions. In effect, the routing protocol should be "topology aware". In this section we describe the operation of our k-hop Routing Protocol (KRP), a protocol that extends upon the popular Ad Hoc On-demand Distance Vector (AODV) [18] protocol.

A. The Bulletin Board

Gateways communicate reachability of MHs in their local area via a shared *Bulletin Board* (BB). The BB consists of a set of destinations along with a list of GWs that are able to reach that destination. Also listed per GW is a hop count indicating the minimum number of hops that exist between the destination and the GW. GWs are individually responsible for updating the BB immediately after learning new routing information. For a discussion on the implementation details of the BB, see [16].

B. The Gateway List

When a MH receives any packet containing a path pertaining to a GW, it takes this information and creates or updates the corresponding entry in its own *Gateway List* (GL). The GL serves as a repository for selecting a default path when no known route to a destination exists. Each entry in the GL contains just two fields — the gateway address and the minimum number of hops needed to reach the GW. The list is kept in sorted order so as to always prefer GWs that are closer to the MH. The GL is closely tied to a MH's routing table in such a way that all additional information about a GW listed in the GL (such as the next hop, the time remaining till the entry expires, etc.) is only stored in the routing table but can be quickly referenced from the GL by hashing on the address of the GW.

C. Path Acquisition

KRP acquires routes by processing the information stored in received control packets. These packets often contain a `<source, sequence number>` pair that identifies the original source of the packet and the most recent *sequence number* for that source. Behavior for maintaining sequence numbers for each destination borrows from AODV [18]. Each KRP packet also contains a `hop count` field denoting the number of hops to the source, and is recorded in the routing entry when necessary. This field is incremented by one when sent or forwarded by a MH. GWs, on the other hand, either set (when sending) or reset (in the case of forwarding) this value to 1. A simple extension to KRP is possible to allow for storage of multiple paths per destination (see [16]), however, the benefit of this added feature is not presented in this paper.

D. Beaconing

KRP possesses a proactive *beaconing* mechanism used by GWs in order to track MH locations and refresh ongoing connections. The scope of each beacon is limited to k wireless hops. BEACON packets contain the beacon interval (`BEACON_INTERVAL`) of the GW. A MH who does not hear from the GW after (`BEACON_LOSS` $\times$ `BEACON_INTERVAL`) seconds deletes the corresponding GL entry. Finally, the most recent sequence number of the GW is incremented and placed in the BEACON packet. BEACON packets are broadcast using a counter-based scheme (see [16] for details).

Upon receiving a BEACON packet, a MH follows the scheme in Sections IV-B and IV-C. A MH will respond to BEACON packets if (1) the GW is new to the MH or the GW's entry has expired, or (2) the MH is participating in a remote connection. For both cases, a reply is issued by the MH to alert the GW of its presence. Non-duplicate beacons are scheduled for retransmission after a short random delay (maximum delay of 5 ms) or dropped if the remaining TTL is 0. Before retransmission, the hop count field of the packet is incremented to account for the current hop.

Replies from a MH come in the form of an AODV-style RREP. The RREP specifies how long the GW should keep the route valid, and is unicast back to the GW using the most recently acquired path. The receiving GW handles the RREP by performing the actions of Sections IV-A and IV-C.

E. Route Discovery

Route discovery in KRP is a multiphase process, starting with a source-initiated RREQ flood for the destination that is limited to k hops. If the destination is within this scope, it responds with a RREP along multiple paths as described in Section IV-C. A connection established in this fashion is considered *DIRECT*. If the initial RREQ attempt fails, the data packet is forwarded to the nearest GW as described in Section IV-B, and the MH switches to *REMOTE* state. It is possible that a MH does not have a valid path to any GWs, in which case it issues a RREQ with the destination address set to the special `ALL_GWS` identifier. A GW receiving a RREQ with this address replies once with its own information via a RREP. The number of times a MH will attempt this step is based on the parameter `DISCOVERY_ATTEMPTS`.

Route discovery is initiated by a GW when receiving a data packet that corresponds to a non-existent or expired BB route. If the GW is listed as the sole next-hop for the destination but the entry is marked as expired, the GW buffers the packet and performs route discovery for the destination. Otherwise, a network-wide *page* consisting of each GW issuing a RREQ for the destination with a TTL of k is executed. Paging is an expensive mechanism, and hence the rest of the protocol is designed to reduce the amount of paging. The number of times the network will page for a destination is based on the `PAGE_ATTEMPTS` parameter.

F. Route Recovery

In KRP, route recovery is performed only by GWs and connection endpoints. This process is triggered when all of the paths to a particular MH or GW are rendered invalid. We first describe the steps taken by a GW to deal with such an event, followed by MH handling.

When a GW receives notification of a broken path that involves one or more routing entries, it attempts to re-route any pending data packets by handing them off to another GW listed in the BB for the destination. If no alternative GW is available, the GW sets the *repair* flag in the BB entry for the destination, and adds the destination to a list of MHs that require route discovery. The resulting compiled list of destinations are copied into a RREQ packet, along with their most recent sequence numbers. The *multi-destination* RREQ is then broadcast locally by the GW. When the RREQ reaches one of the intended destinations, it responds normally with a RREP. The GW keeps track of those destinations that fail to respond, and follows up with an individual RREQ for the destination, up to `REPAIR_ATTEMPTS` times. After local recovery if a GW still cannot reach a destination, it activates the paging mechanism.

A RERR that eliminates the last path to a MH in a direct connection triggers a RREQ for the destination, following the rules of Section IV-E.

V. SIMULATION

The purpose of this section is to describe the set of experiments that we used to test the performance of KRP. For brevity, many of our results are not included here, but the interested reader can refer to [16]. In this paper, we chose to concentrate on highlighting the behavior of our beaconing scheme.

A. Simulation Environment

Simulations were completed using *ns2*. All runs were performed with version 2.1b9a of the simulator. Common simulation parameters can be found in Table I.

For all experiments, we chose 100 MHs to roam around a rectangular area of 2400 m $\times$ 600 m at four different maximum speeds of 1, 5, 10, and 20 m/s. GWs are placed 600 m apart and 300 m from the top and bottom of the simulation area. To cover a wide range of mobility and traffic patterns,

TABLE I
SIMULATION PARAMETERS.

Transmission Range	250 m
Simulation Duration	900 s
Mobility Model	Random way-point [19]
Pause Time	0 s
Medium Access Protocol	IEEE 802.11 DCF
Radio Propagation Model	Two-ray Ground
Link Bit-rate	2 Mb/s

for each point of data is a result of 36 simulations consisting of 6 different mobility patterns and 6 traffic files. The traffic is comprised of Constant Bit Rate (CBR) communications between MH sources and sinks. Each CBR source sends 64 byte packets with 200 ms spacing or an equivalent rate of roughly 2.5 kb/s. This traffic load is similar to that used in previous tests of popular routing protocols [19]. A 95% confidence interval was computed for each data point, and is indicated in each graph by vertical bars.

B. Metrics

The overall goal of our experiments is to measure and compare the ability of the proposed routing protocol to react to a changing network topology. To accomplish this goal, we subject it to a series of different scenarios representing a range of conditions. Its performance is marked by the following metrics — (1) Packet Delivery Ratio (PDR), (2) End-to-End Path Latency (E^2PL), and (3) Routing Overhead (RO). Graphs denote the average of these metrics over all runs.

C. Experiments

Our initial set of tests involve subjecting KRP to a network with 30 active connections using different values for `BEACON_INTERVAL`, the interval between beacon packets. Figure 2 shows the results of these simulations (k =2). In general, the protocol performs well against the three established metrics. PDR stays close to 99% even in high mobility scenarios. When considering that our calculation of PDR labels packets that are *unroutable* (no path exists between source and destination) as *undeliverable* (the protocol could not deliver the packet successfully), the actual PDR is in fact much higher when subtracting these packets from the total (see [16]). Although not shown, the overhead of KRP at a beacon interval of 10 seconds and speed of 1 m/s is 17% of the data traffic load, rising to 30% at 20 m/s. At 1 m/s, the average end-to-end packet delay is 17 ms, and doubles by 20 m/s as indicated in the graph.

Of immediate notice in the figure is that the PDR *decreases* and the E^2PL *increases* as the beaconing interval decreases. This unanticipated behavior is a result of the beaconing process' bursty nature, as illustrated by Figure 3. This graph shows the time interval of 120 seconds to 126.5 seconds for a run with `BEACON_INTERVAL` set to 3. Around the time that beacons and replies are sent (at 120, 121.5, 123, 124.5 and 126 seconds) the end-to-end packet latency skyrockets to

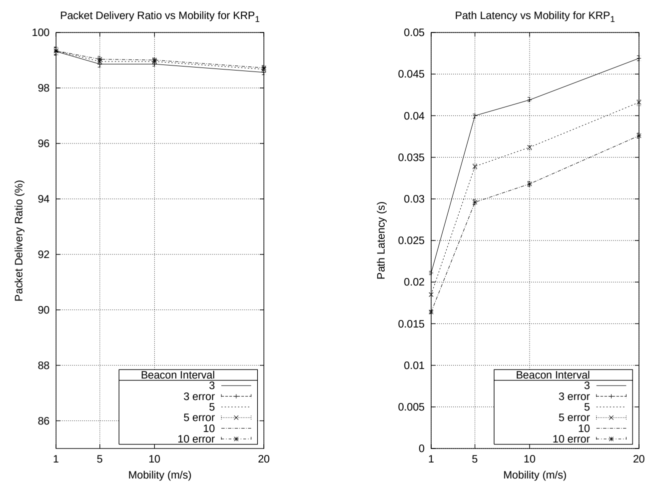

Fig. 2. PDR and E^2PL vs. Mobility for KRP (30 sources).

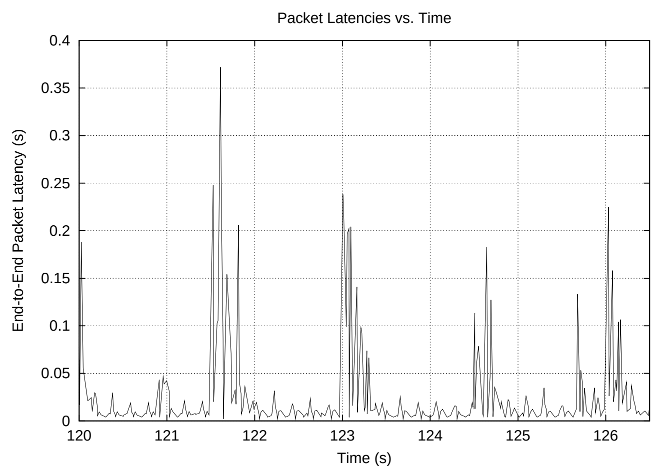

Fig. 3. E^2PL vs.time for KRP (`BEACON_INTERVAL = 3`, 30 sources).

nearly 10 times that of the average latency in-between the beacon periods. During the beacon and reply process, data packets collide with the broadcast BEACON packets, requiring retransmission. They also contend with the numerous replies that are sent. This short burst of periodic control overhead puts a strain on the network that is reflected in the data packet latencies.

To gain a better understanding of this phenomenon, the same test was performed using only 10 connections with `BEACON_INTERVAL` = 3. The results are plotted against the 30-connection scenario in Figure 4. With fewer connections, the performance of KRP shows significant improvement. The reduced contention is largely responsible for this change, and can be supported by looking at the graphs. The CO graph reveals that tripling the number of connections results in a more than 6-fold increase in overhead. The E^2PL graph confirms that the latency is considerably less with fewer traffic sources. The conclusion here is that *the beaconing mechanism*

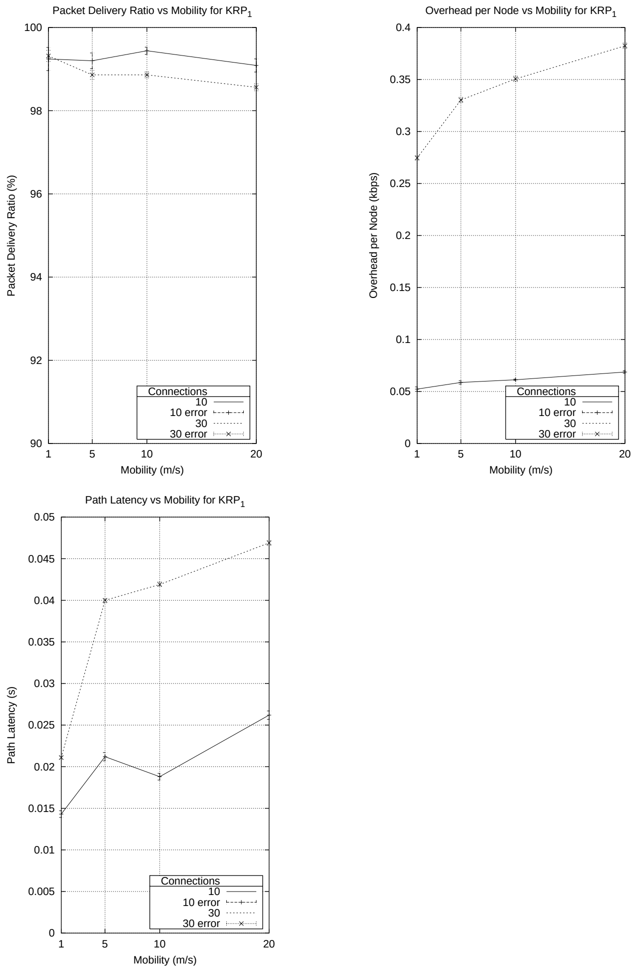

Fig. 4. PDR, CO, and E^{2}PL vs. Mobility for KRP (`BEACON_INTERVAL = 3`).

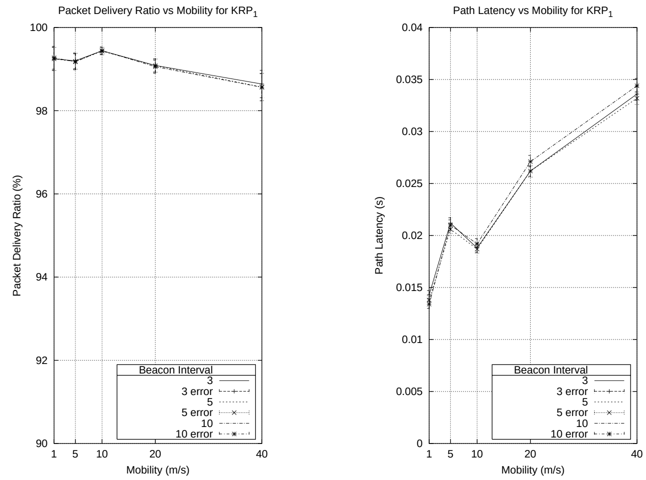

Fig. 5. PDR and E^{2}PL vs. Mobility for KRP (10 sources).

does not scale well as the number of connections increases.

Figure 5 shows the 10 connection scenario using different values for `BEACON_INTERVAL`. The PDR graph demonstrates that with fewer traffic sources the packet delivery ratio increases (rather than decreases as was the case for 30 connections) as the beacon frequency is increased. Likewise, the E^{2}PL for a beacon interval of 3 seconds is now less than that of the 5 and 10 second intervals at higher mobility speeds.

In addition to the tests described here, a study was conducted to measure the performance of KRP under different values of k. The results indicate that the overhead quickly out-paces the extra connectivity achieved as k increases. More details on this conclusion are presented in [16].

VI. Conclusions and Future Work

The main conclusion of the work discussed in this paper is that a periodic beaconing process can be useful for establishment and maintenance of paths, but must be carefully executed to avoid significant interference with the success and latency of ongoing data communications, as seen in Section V with KRP. For brevity, a more thorough discussion relating to the nature of beaconing is left in [16]. This paper also introduces our k-hop architecture, which leverages limited wireless connectivity in exchange for low overheads. A k-hop network has been implemented in our laboratory testbed (see [20] and [16] for details).

An area of future work involves the introduction of fixed *relays* that provide connectivity support for sparsely populated networks. Relays are positioned strategically throughout the network and serve as forwarding points for MHs that are more than one hop away from the GW. Relays take the burden of packet forwarding off of intermediate MHs, serving as an extension of the backbone infrastructure. Hence, they can employ more efficient routing techniques and provide a more accurate depiction of each MHs location in network.

References

[1] C. E. Perkins, Ed., *Ad Hoc Networking*. Addison Wesley, 2001, ch. 1, pp. 8–14.

[2] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388–404, March 2000.

[3] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Rome, Italy, July 2001.

[4] C. A. Santiváñez, B. McDonald, I. Stavrakakis, and R. Ramanathan, "On the scalability of ad hoc routing protocols," in *Proceedings of the IEEE Conference on Computer Communications (INFOCOM)*, New York, NY, June 2002.

[5] C. Tschudin and R. Gold, "LUNAR: Lightweight underlay network ad-hoc routing," University of Basel, Switzerland, Tech. Rep., January 2002.

[6] H.-Y. Hsieh and R. Sivakumar, "A hybrid network model for cellular wireless packet data networks," in *Proceedings of the IEEE Global Telecommunications Conference (Globecom)*, Taipei, Taiwan, Nov. 2002.

[7] H. Luo, R. Ramjee, P. Sinha, L. Li, and S. Lu, "UCAN: A unified cellular and ad-hoc network architecture," in *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, San Diego, CA, Sept. 2003.

[8] C. Qiao and H. Wu, "iCAR: An intelligent cellular and ad-hoc relay system," in *Proceedings of the IEEE International Conference on Computer Communications and Networks (IC3N)*, Las Vegas, NV, Oct. 2000.

[9] Y. Lin and Y. Hsu, “Multihop cellular: A new architecture for wireless communications,” in *Proceedings of the IEEE Conference on Computer Communications (INFOCOM)*, Tel Aviv, Israel, Mar. 2000.

[10] X. Wu, S. H. Chan, and B. Mukherjee, “MADF: A novel approach to add an ad-hoc overlay on a fixed cellular infrastructure,” in *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, Chicago, IL, Sept. 2000.

[11] J. Zhou and Y. R. Yang, “PARCelS: Pervasive ad-hoc relaying for cellular systems,” in *Med-Hoc-Net’02*, Sardegna, Italy, September 2002.

[12] G. Aggelou and R. Tafazolli, “On the relaying capacity of next-generation GSM cellular networks,” *IEEE Personal Communications*, vol. 8, no. 1, pp. 40–47, February 2001.

[13] Y.-B. Ko and N. H. Vaidya, “A routing protocol for physically hierarchical ad hoc networks,” Texas A&M University, Tech. Rep. TR97-010, Sept. 1997.

[14] K. Xu, X. Hong, and M. Gerla, “An ad hoc network with mobile backbones,” in *Proceedings of the IEEE International Conference on Communications (ICC)*, New York, NY, Apr. 2002.

[15] G. Pei, M. Gerla, and X. Hong, “LANMAR: Landmark routing for large scale wireless ad hoc networks with group mobility,” in *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Boston, MA, Aug. 2000.

[16] W. D. List, “Routing protocols for k-hop networks,” Master’s thesis, University of Illinois at Urbana-Champaign, Aug. 2003.

[17] K. Xu and M. Gerla, “A heterogeneous routing protocol based on a new stable clustering scheme,” in *Proceedings of the Military Communications Conference (MILCOM)*, Anaheim, CA, Oct. 2002.

[18] C. Perkins and E. Royer, “Ad-hoc on-demand distance vector routing,” in *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, New Orleans, LA, Feb. 1999.

[19] J. Broch, D. A. Maltz, D. B. Johnson, Y.-C. Hu, and J. Jetcheva, “A performance comparison of multi-hop wireless ad hoc network routing protocols,” in *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Dallas, TX, Oct. 1998.

[20] M. J. Miller, W. D. List, and N. H. Vaidya, “A hybrid network implementation to extend infrastructure reach,” `www.crhc.uiuc.edu/ nhv/publications.html`, University of Illinois, Champaign-Urbana, Tech. Rep., January 2003.