

Routing Packets into Wireless Mesh Networks

Rainer Baumann*, Simon Heimlicher*, Vincent Lenders†, Martin May*

*Computer Engineering and Networks Laboratory
ETH Zurich, Switzerland

†Department of Electrical Engineering
Princeton University, NJ, USA

Abstract—Wireless mesh networks are a promising way to provide Internet access to fixed and mobile wireless devices. In mesh networks, traffic between mesh nodes and the Internet is routed over mesh gateways. On the forward path, i.e., from mesh nodes to Internet nodes, for all mesh nodes only route information for one destination, the gateways, needs to be maintained. However, on the backward path from the Internet to mesh nodes, an individual route for every mesh node is required. In this paper we investigate protocols for backward path routing in wireless mesh networks. Using simulation experiments with realistic mobility patterns of pedestrians and cars in cities, we compare three protocols, each of which represents a routing protocol family: (i) AODV with an extension for mesh networks, a reactive routing protocol, (ii) FBR, a proactive routing protocol, and (iii) GSR, a source routing protocol. Our results indicate that FBR has the highest packet delivery ratio but is not scalable to the network size. The extended AODV seems to be neither scalable nor does it achieve a high packet delivery ratio. A good compromise is provided by GSR, which is the most scalable to the network size and still achieves a high packet delivery ratio.

I. INTRODUCTION

In recent years, the number of wireless-enabled devices such as laptops, PDAs, and mobile phones has been growing rapidly. At the same time, the Internet has turned into a critical networking infrastructure and continuous access is highly demanded from end users. Combined, these developments constitute a strong demand for Internet access from mobile devices anywhere, anytime. With current wireless technology, the gap between a set of wireless devices and the Internet is bridged by fixed wireless stations, so-called *access points*. Such access points establish a network among the wireless devices and bridge this network with the Internet. However, the communication range of current commodity wireless LAN technology is limited to a few hundred meters and complete coverage of a city would require a prohibitively high number of access points.

A solution to extend the coverage of the existing access points is to let wireless devices relay packets towards access points on behalf of their neighbors. Such multi-hop wireless access networks are called *wireless mesh networks (WMN)* in the remainder of this paper. An example of a wireless mesh network is depicted in Fig. 1, where intermediate fixed and mobile wireless devices (*mesh nodes*) relay data between Internet access gateways (*mesh gateways*) and more distant devices.

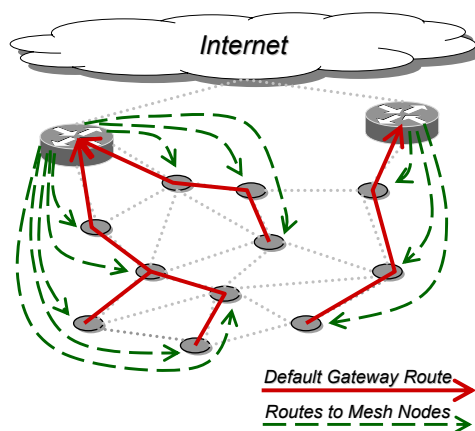


Fig. 1. Wireless mesh network with one default route out of the mesh and a single route to every mesh node.

The majority of Internet services, such as web surfing, video and radio streaming, are hosted on servers located somewhere in the Internet; hence the majority of the traffic in a WMN flows between the Internet and the mesh nodes and we assume that the traffic between mesh nodes is negligible.

In related studies, routing protocols were proposed and evaluated for wireless mesh networks. Most of these protocols are designed to route in the direction from the mesh nodes to the Internet (e.g. AODV [1] or OLSR [2]). However, most services generate asymmetric traffic and the amount of downstream traffic from servers in the Internet to the mesh nodes far exceeds upstream traffic. Consequently, in this paper we study, how to route data from the Internet to the mesh nodes. We will refer to this direction as the *backward path*.

Backward path routing is more involved than routing on the forward path, due to multiple reasons. While data addressed at any Internet host only needs to be forwarded to *any* gateway, the backward path routing protocol needs to be able to address every mesh node separately (see Fig. 1). Additionally, mesh nodes may be mobile, leading to frequent changes in the topology of the wireless mesh network. Consequently, the backward path routing information needs to be updated frequently.

As with forward path routing, optimizing the performance of the backward path is a trade-off between accuracy of the routing state information and control traffic overhead.

We compare three backward path routing protocols, each of which represents one of the three most widely used families of routing protocols: reactive hop-by-hop routing (AODV-CGA), proactive hop-by-hop routing (FBR), and proactive source routing (GSR). One major difference between GSR and the other two protocols is that with GSR, all routing state information is kept on the mesh gateways, while with the other two protocols; the state information is distributed on the mesh nodes. Furthermore, GSR re-uses forward paths collected by data packets while the other protocols depend on dedicated routing control packets to determine routes.

AODV-CGA is a version of AODV extended for mesh networks as proposed by Braun et al. [3]; the other two protocols, FBR and GSR, we have developed ourselves.

- 1) *Reactive hop-by-hop routing (AODV-CGA)*: We use an extended version of AODV specialized for wireless mesh networks. This extension extends the AODV routing domain to include a network border router that is connected to all gateways. Upon receiving a packet from the Internet addressed at a mesh node, the border router floods a route request to all gateways from where a route is then determined in the mesh network.
- 2) *Proactive field-based routing (FBR)*: We propose a proactive field-based routing protocol similar to HEAT [4]. With this protocol, every mesh node maintains a scalar field that is propagated by beacon messages through the mesh network. Routing towards a specific mesh node is achieved by forwarding along the steepest gradient of the field of the destination node.
- 3) *Gateway source routing (GSR)*: We present a source routing protocol that re-uses the forward paths that are recorded by data packets and stored on the gateways. These paths are then used for source routing on the backward path.

Based on the experiences we made during the evaluation, we introduce modifications to the FBR and the GSR protocols that mitigate some of their limitations and improve their performance.

With the Glomosim [5] network simulator, we compare the performance as well as the communication overhead of these three routing protocols. Our mobility patterns are based on the vectorized street map of the city of Zurich, Switzerland. In order to account for worst-case situations, we assume that all nodes except the gateways are moving constantly. We simulate an application similar to radio streaming, where traffic is sent by a few Internet servers to mesh nodes. In order to account for the frequently required “keep-alive” feedback packets, the receiving mesh nodes periodically send a data packet to the Internet server. Note, that there is no traffic on the forward path except these feedback packets to ensure that our evaluation only considers the performance of the backward path protocols.

In our simulation experiments, FBR outperforms the other

protocols with respect to the packet delivery ratio. However, FBR does not scale well with the network size, especially if routes are updated frequently. GSR scales much better to the network size. For instance, with an update rate of one update per second, the packet delivery ratio of GSR is almost as good as with FBR and the communication overhead is much lower. We show that, besides offering very good performance, GSR offers the best scalability properties with respect to the network size. For instance with a route update frequency of five seconds, GSR has a somewhat lower delivery ratio than FBR but still outperforms AODV-CGA in both delivery ratio and communication overhead.

The rest of this paper is organized as follows. In the next section, we highlight related work. Then, we present the three routing protocols we compare in this paper. In Section IV, we specify our simulation environment while in Section V, we present our evaluation results. Finally, we discuss the results and conclude the paper in Section VII.

II. RELATED WORK

Routing in wireless networks has undergone extensive study. A fundamental property of the mesh networks for Internet access that we consider in this paper is that all traffic flows between the mesh nodes and a small number of gateways.

A. MANET and mesh network routing protocols

A multitude of unicast routing protocols for MANETs have been proposed (e.g. [1], [2], [6]–[8] to name just a few), and these protocols have been extended with gateway discovery functionality [3], [9]–[16] for wireless mesh networks. Most of these protocols use hop-by-hop routing. For example Nordström et al. proposed in [15] to use source routing based on DSR to route packets between wireless mesh nodes and multiple gateways. There, the wireless mesh gateways specify the routes using previously gained knowledge on the network. In OLSR [2], it is mentioned that interoperation with other networks should be possible by injecting external route information into the OLSR network. In [3], Braun et al. propose to extend AODV for use in mesh networks in a setup called “Common Gateway Architecture (CGA)”. We compare our two protocols with an implementation of this architecture and refer to it as AODV-CGA.

B. Field-based routing

Field-based or gradient-based routing has been proposed in the past for various types of applications, including routing in MANETs [4], [17], load balancing in the Internet [18], data collection in sensor networks [19], [20], sensor node placement [21], guided navigation [22], or service discovery in MANETs [23]. These routing schemes all share the same design idea: routing is based on a scalar field that assigns a scalar value to every node in the network. The destinations are represented as maximum values and packets are always forwarded along the steepest gradient towards the destination.

The concept of field-based routing is fairly simple yet it provides a very versatile bias for routing decisions. For instance, modeling shortest-path routing with a field-based scheme is straightforward and has been demonstrated in [24]. Owing to the fundamental properties of fields, loop freedom of routes is ensured, and it is guaranteed that packets are routed towards the destination. By design, a field may comprise multiple routes, thus if the link to the next hop breaks, a successor can easily be determined. As soon as the routing protocol has re-established the field, an alternative next hop is determined. The proactive maintenance of the field state requires the exchange of periodic beacon messages among neighbors. Hence, in contrast to AODV-CGA and GSR, the FBR protocol incurs a small communication overhead even if there is no data traffic in the mesh network.

III. ROUTING PACKETS FROM THE INTERNET TO WIRELESS MESH NODES

In this section, we briefly introduce the three backward path routing protocols that we compare in this paper: AODV-CGA, FBR, and GSR. Subsequently, we introduce FBR-GW, an enhancement to FBR that improves its scalability to the network size. Then, we present GSR-PN, an enhanced GSR protocol with higher packet delivery ratio.

A. The Protocols under Study

1) *AODV-CGA (Reactive Hop-by-hop Routing)*: As the representative of the class of reactive routing protocols, we use an extended version of AODV that we call AODV-CGA. This extended AODV protocol allows the use of multiple gateways to the Internet and was proposed by Braun et al. [3]. AODV-CGA shares most mechanisms with the well-known AODV protocol and thus may be seen as a benchmark in our comparison. The basic principle of operation of AODV is as follows. AODV constructs routes *on demand* by flooding route requests using an expanding ring search mechanism. If a request reaches a destination, the destination node answers with a route reply message that is forwarded using temporarily stored information from the request. The reply concurrently establishes the route by registering it at every intermediate node. In AODV, route error messages and sequence numbers are used to deal with broken routes.

As proposed by Braun, all gateways are connected to a dedicated router that acts as a proxy to the Internet. This router has two tasks: (i) on the forward path, it sends route replies on behalf of hosts in the Internet; (ii) on the backward path, it initiates route requests for nodes in the wireless mesh network. To improve scalability with respect to the network size, expanding ring search [25] is used for the route requests, which slightly increases the route set-up time.

2) *FBR (Proactive Field-based Routing)*: As the representative of the proactive routing protocol family, we propose a field-based routing protocol similar to HEAT [4]. In HEAT,

wireless mesh nodes periodically exchange beacons. These beacons contain a list of all known destinations with their respective field value. When a new destination appears, it announces its presence with beacons to its neighbors in order to establish a field. With this mechanism, a field on the network is constructed for every destination. This field assigns a value to every node in the network; the destination bears the maximum value. Packets are then routed along the steepest gradient towards the destination.

Owing to the fundamental properties of fields, loop freedom of routes is ensured, and it is guaranteed that packets are routed towards the destination [26]. Field-based routing enables nodes to consider multiple routes to the destination, thus if the neighbor with the highest field value disappears, an alternative route can be determined easily if one is available.

Note, that FBR—being a proactive routing protocol—incurs a small communication overhead since it proactively maintains all routes, regardless of whether there is data traffic or not.

3) *GSR (Gateway Source Routing)*: With gateway source routing (GSR) [4], we propose to reuse the forward path information from the packets that arrive at the gateways. In the routing header of every packet, the intermediate hops from the mesh node to the gateway are recorded. These paths are then stored in the gateways. To route packets to a mesh node, the mesh gateway inverts the recorded forward path and copies it to the packet header. The gateway then sends the packet to the first node of the backward path. Each node updates the path in the header by removing its entry and forwards the packet to the given next hop until the packet reaches the destination.

By design, this approach is scalable to the number of mesh nodes as it imposes no overhead that depends on this number. Only the gateways have to maintain up-to-date routes to individual mesh nodes. Also, this approach does not increase the number of control packets exchanged between the mesh nodes, and thus reduces the chance of collisions.

Obviously, GSR requires that a packet towards a host in the Internet is first sent by a mesh node in order to establish the backward path. However, since the majority of communication is initiated by mesh nodes, we consider this to be only a small limitation.

Should a mesh node act as a server, a dedicated addressing mechanism (e.g., HIP [27], [28]) would probably be used. HIP and most other addressing mechanisms require periodic registration messages from the mesh node towards a gateway. Those periodic registration messages serve also to initiate and maintain the path at the gateway. Should traffic be unidirectional from an Internet node towards a mesh node, this poses a problem, as the backward path cannot be maintained. This problem can only be solved by requiring the mesh node to periodically send some sort of ping packets to a gateway. Note, however, that most applications have a feedback channel (i.e., TCP acknowledgments or RTCP messages for streaming applications) and hence generate bidirectional traffic. In our

evaluation study in Section V, we evaluate a broad range of feedback intervals to shed some light on the trade-off between communication overhead and the quality of the backward path.

B. Enhancements

Based on simulation experiments, we propose enhancements to our protocols. With FBR-GW we aim to improve the scalability to the network size of FBR and with GSR-PN we strive to improve the packet delivery ratio of GSR.

1) *FBR-GW (Enhancement for Scalability)*: The FBR protocol is designed for maximum packet delivery ratio, but it scales poorly with the network size. Routing from the mesh nodes to the Internet gateways requires only a single field. In contrast, routing from the gateways back to the mesh nodes requires a separate field for every mesh node. Thus, a field of every node is propagated through the entire network.

Assuming that there are no connections among the mesh nodes, the scalability to the network size can be improved as follows: We propose to let the field information of mesh nodes only be propagated towards the gateways. Implementing this enhancement is straightforward. Instead of establishing the “per-node” field with all neighbors, only the neighbors with an increasing “gateway field” value are used to establish a mesh node field. As a result, state information about individual mesh nodes is only established in nodes that might be used for packet forwarding.

2) *GSR-PN (Enhancement for Performance)*: The GSR protocol is designed for scalability to the network size. However, its packet delivery ratio drops rapidly if the routes are not frequently updated by feedback packets. Most packets are lost due to paths that contain links that are broken because the nodes moved away from each other. Such link breaks happen mostly between nodes that are almost at the maximal transmission range to each other.

In order to reduce the probability of mobility-induced link breaks, we propose to add a preferred neighbor (PN) mechanism to GSR that is similar to the mechanism we presented in [29]. With this mechanism, links between nodes at a preferable distance are used whenever possible. To this end, nodes are classified into three groups, based on the received signal strength value (RSSI), see Fig. 2:

- *Preferred Neighbor (PN)* group: nodes with a signal level in the preferred range;
- *In* group: nodes with a signal stronger than the preferred level;
- *Out* group: nodes with a signal weaker than the preferred level.

To classify a node, the power of a received signal is compared to two values: Inner Threshold (IT) and Outer Threshold (OT). More details can be found in [29].

IV. SIMULATION SETUP

We perform our simulations with Glomosim [5], a network simulator for wireless networks. With Glomosim, we evaluate

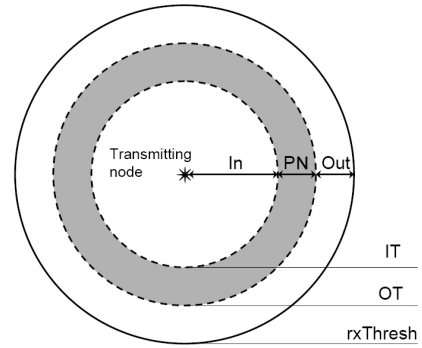


Fig. 2. Preferred Neighbor Selection. Note: For simplicity, we depict the communication range of a node and the signal strength levels for the *in* and *out* groups as circles, but our algorithm by no means assumes that the range is indeed circular.

the performance, overhead, and set-up time of the three presented protocols. All simulations run for a duration of at least 10000 seconds. Note also that the results presented in the next section are always an average over at least 20 runs with different random seeds.

A. City mobility model

We use a mobility pattern that models the mobility of pedestrians and cars in several Swiss cities. The road maps of these cities are extracted from the Swiss geographic information system (GIS) [30] and they include vectorized building and street maps along with speed limit data. As an example, the vectorized map of the city center of Zurich is shown in Fig. 3.

The actual movements of the nodes are generated according to the steady-state random trip mobility model [31], as follows: A node chooses a random destination in the city and moves to this position at a constant speed along the shortest path, always following the proper road. For instance, pedestrians do not use highways. In order to model worst-case mobility, nodes never pause. Hence, a node moves to a new destination as soon as it reaches the previous destination.

We run simulations with three scenarios that differ in the node speed: (i) a scenario with nodes moving at pedestrian speeds (uniformly distributed between 1 m/s and 4 m/s), (ii) a scenario with nodes moving at car speeds (i.e., 10 – 20 m/s), and (iii) a static scenario for model verification and benchmarking. In each scenario, there are 1000 nodes and 5 gateways. The simulation area is 5 km by 7 km .

B. Radio settings

The radio propagation is modeled by the two-ray ground propagation model. All mesh nodes and gateways are equipped with a 802.11b radio with a bandwidth of 11 $Mbps$ and a nominal range of 250 $meters$. As MAC layer protocol, we use the 802.11 DFWMAC-DCF w/RTS/CTS.

C. Traffic pattern

We are interested in the suitability of the routing protocols for Internet traffic, which is typically a mix of streaming and

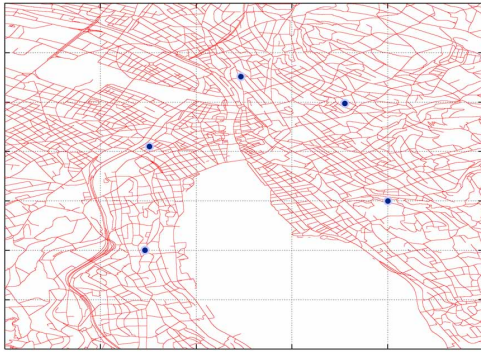


Fig. 3. Vectorized street map of the city of Zurich, Switzerland (5km by 7km). The dots indicate the positions of the mesh gateways.

request-response traffic [32], [33]. In this paper, we focus on the backward path and we use highly asymmetric streaming traffic that corresponds to radio streaming. Symmetric traffic as produced by VoIP or peer-to-peer streaming applications might lead to different results. In this paper, however, the focus is on the backward path performance and therefore, we only consider backward path traffic.

We assume high quality audio with a constant bit rate of 96 *kBit/s* [34], [35]. From the 1000 nodes present in the simulations, 200 nodes act as streaming clients. The durations of streams are exponentially distributed with an average of 480 seconds and an exponentially distributed pause time of 120 seconds.

The streaming traffic originates at Internet servers and is forwarded to mesh nodes. The mesh nodes regularly send feedback packets of 100 Bytes to the Internet server they receive data from. In order to examine the stability of the routes, we vary the interval of feedback packets from 1 to 40 seconds.

V. SIMULATION RESULTS

We use the following metrics to compare the performance and overhead of the three routing protocols:

- *Packet delivery ratio*: The ratio between the number of packets that are received and the number of packets sent. This metric only considers backward path traffic, i.e., the data packets from the gateways to the mesh nodes.
- *Routing overhead*: We use the following metrics to capture several aspects of routing overhead:
 - The average number and size of routing control messages sent per second by a mesh node.
 - The average additional header space required for routing information in a data packet.
- *Route set-up time*: The time elapsed until a demanded route is available. Unsuccessful route establishments are ignored.

A. Packet delivery ratio

When we compare the packet delivery ratio in the static scenario, all routing protocols deliver over 99.6% of the data packets. However, already if nodes move at pedestrian speed, the results differ. In Fig. 4, we plot the delivery ratio of the different routing protocols vs. the feedback packet interval. Increasing the feedback interval decreases the freshness of the routing information. Due to its proactive route maintenance mechanism, the FBR protocol is almost independent of the feedback interval. The enhanced FBR-GW protocol shows similar behavior at a lower ratio of roughly 90%. In the simulation log file, we found that the packet losses occur when a node that is directly connected to a gateway moves away. During the period where the node re-establishes its field, packets are lost since there is no alternative path to such a node.

The AODV-CGA protocol performs very well if the feedback interval is long and achieves a packet delivery ratio of up to 95%. However, with shorter feedback intervals, the routing packet broadcasts interfere with the data traffic and the packet delivery ratio decreases to 84% at a feedback interval of 1 second. Note, that this feedback interval corresponds to less than 1 *kbit/s*. If the traffic were symmetric, the delivery ratio of AODV-CGA would presumably drop further. Apparently, AODV-CGA is very sensitive to the network load. A detailed study of the routing overhead is presented in the following subsection.

The GSR protocol performs almost as good as the FBR protocol when the feedback interval is short. But when the paths stored at the gateways are updated less frequently, the packet delivery ratio drops quickly. The enhanced protocol, GSR-PN, reduces this problem. With a feedback interval of 5 or even 10 seconds, GSR-PN still achieves a higher packet delivery ratio than all other protocols of our study except FBR.

In Fig. 5, we plot the packet delivery ratio for nodes moving at car speed. The results are similar to those of the pedestrian model. The FBR protocol still outperforms all other protocols, but, due to the high mobility, the delivery ratio is only around 90%. Note that the GSR-PN protocol performs almost as good as the AODV-CGA protocol at a feedback interval of 5 seconds.

A profile over all scenarios with a feedback interval of 5 seconds is provided in Fig. 6. Obviously, all routing protocols are affected by the higher node speed, but the FBR-GW protocol is particularly susceptible to this parameter.

B. Routing overhead

Besides packet delivery, routing overhead is the most important property that we evaluate. In our simulation setting, we compare the following routing overhead metrics:

- 1) Number of routing packets sent per second by a mesh node
- 2) Average size of routing packets

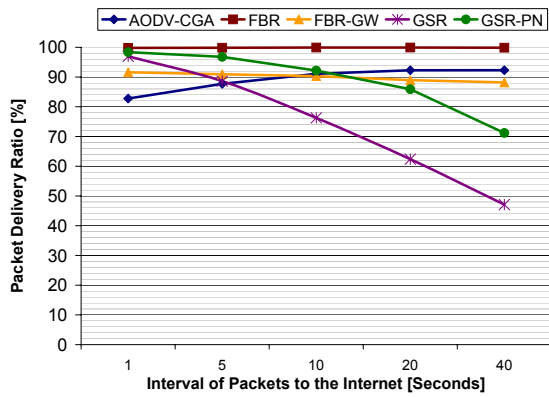


Fig. 4. Packet delivery ratio in the pedestrian scenario.

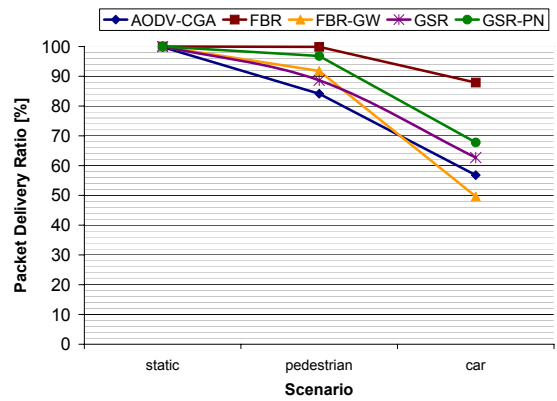


Fig. 6. Packet delivery ratio with a feedback interval of 5 seconds.




Fig. 5. Packet delivery ratio in the car scenario.

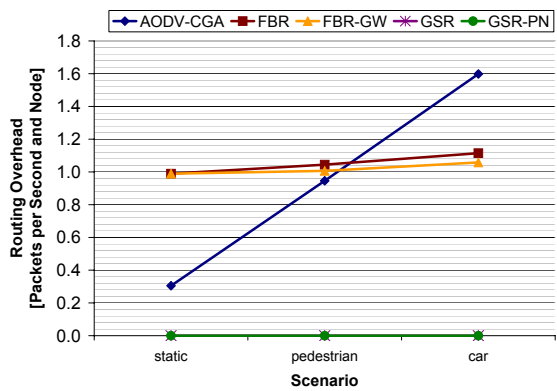


Fig. 7. Average number of routing control packets sent per second and per node.

3) Average additional header space required per data packet
We evaluate these metrics with a feedback packet interval of 10 seconds.

(i) Figure 7 shows that for AODV-CGA, the number of routing packets per data packet raises steeply with increasing mobility while it remains almost constant for the proactive protocols. This is not surprising, since AODV-CGA is a reactive routing protocol. Its advantage lies in the low communication overhead if there is little or no traffic.

(ii) Figure 8 shows the size of routing packets in a logarithmic scale. While both GSR and GSR-PN do not require additional routing packets, the other protocols are based on dedicated routing control packets. The routing packets of AODV-CGA are small and have a constant size. The FBR protocol produces larger routing packets. As the plot shows, the limitation of field information propagation used in FBR-GW helps to reduce the packet size and thus increases the scalability to the network size. However, with increasing speed, the packet size increases considerably.

(iii) GSR and GSR-PN require additional space for routing information in the header of each data packet. This overhead slightly increases for scenarios with higher mobility because the average routes tend to become longer (see Fig. 9). Of note,

with GSR-PN, the space required is higher than with GSR because the Preferred Neighbor enhancement leads to paths with higher hop count that provide shorter and more reliable links.

C. Route set-up time

We compare the route set-up time for the investigated routing protocols. By design, the proactive protocols, i.e., FBR, FBR-GW, GSR, and GSR-PN do not require any route set-up time. With the AODV-CGA protocol, when a route is required that has not recently been used, the protocol initiates the route discovery process. Table I shows the route set-up time for different mobility schemes. At a first glance, the set-up time seems to decrease with increasing node speed. However, by analyzing the simulation logs in more detail, we found that this is due to the fact that the discovery of longer routers fails more frequently. In order to eliminate this effect, the average set-up time only comprises route discovery processes that succeed.

VI. DISCUSSION

In Tab. II, we summarize the strengths and weaknesses of the evaluated routing protocols. Reactive routing protocols are

	AODV-CGA	FBR	FBR-GW	GSR	GSR-PN
Packet delivery ratio	-	++	-	0	+
Routing packets per node and second	-	0	0	++	++
Average size of routing packets	0	--	-	++	++
Average additional space for routing in each data packet	++	++	++	-	-
Routing set-up time	--	++	++	++	++
Scale with the network size	0	--	0	++	++

TABLE II
QUALITATIVE COMPARISON OF EVALUATED ROUTING SCHEMES.




Fig. 8. Average size of routing packets.




Fig. 9. Average routing overhead encapsulated.

scenario	static	pedestrian	car
routing set-up time	0.50s	0.34s	0.28s

TABLE I
ROUTE SET-UP TIME OF THE AODV-CGA PROTOCOL.

widely used in the ad hoc network area. However, in our mesh network scenarios, the AODV-CGA protocol exhibits poor performance.

If the goal is to maximize the packet delivery ratio, the

proactive field-based routing protocol, FBR, should be considered. FBR inherently does not scale to the network size, but for smaller networks it seems to be a good fit.

Considering scalability with respect to the network size, the gateway source routing protocol, GSR, outperforms the other protocols. GSR is scalable to the network size and still achieves a very high packet delivery ratio given frequent feedback packets from the mesh nodes. With the proposed enhancement that prefers neighbors at a certain distance, GSR-PN, the protocol performs better even if feedback packets are sent less frequently.

VII. CONCLUSION

Wireless mesh networks are a large-scale solution to provide Internet access to mobile users. In such a network, mesh gateways provide Internet access to nodes in their vicinity; data from and to mesh nodes further away is relayed through the mesh network. Routing data to the gateways is fundamentally different from routing towards specific mesh nodes.

In this paper, we compare three protocols for routing from the Internet to mesh nodes in static and mobile scenarios. These protocols all represent different classes of routing strategies. Based on our findings, we then propose and evaluate enhancements of two protocols.

In the simulation experiments, AODV-CGA, a reactive hop-by-hop routing protocol based on AODV, exhibits limited scalability to the network size. Furthermore, due to the high route setup time of AODV-CGA, the packet delivery ratio is rather low.

FBR, our proactive field-based routing protocol, outperforms all others with respect to the packet delivery ratio. However, FBR is not scalable to the network size. Furthermore, compared to the other protocols, FBR incurs the highest communication overhead. With FBR-GW, an enhanced version of FBR, the communication overhead decreases slightly.

GSR, our gateway source routing protocol, delivers promising results. Its source routes gained from recorded routes of packets destined to the Internet prove to be quite reliable. GSR is scalable to the network size and has no route set-up

delay (assuming that the receiver has sent at least one packet). We find that if the receiver sends a feedback packet towards the Internet host every five seconds, the packet delivery ratio remains high, even in scenarios where nodes move at car speed. With an enhanced version called GSR-PN, we achieve a higher packet delivery ratio in highly mobile scenarios.

We conclude that gateway source routing is a promising routing approach, since in our study, it delivers the best trade-off between packet delivery ratio, routing overhead, and scalability to the network size.

REFERENCES

- [1] C. Perkins, "Ad-hoc on-demand distance vector routing," in *MILCOM '97 panel on Ad Hoc Networks*, 1997.
- [2] T. Clausen and P. Jacquet, "Optimized Link State Routing Protocol," IETF Internet Draft, draft-ietf-manet-olsr-11.txt, July 2003.
- [3] M. Michalak and T. Braun, "Common gateway architecture for mobile ad-hoc networks," in *WONS '05: Proceedings of the Second Annual Conference on Wireless On-demand Network Systems and Services (WONS'05)*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 70–75.
- [4] Rainer Baumann and Vincent Lenders and Simon Heimlicher and Martin May, "HEAT: Scalable Routing in Wireless Mesh Networks using Temperature Fields," in *Proceedings of the IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2007*.
- [5] X. Zeng, R. Bagrodia, and M. Gerla, "Glomosim: A library for parallel simulation of large-scale wireless networks," in *Workshop on Parallel and Distributed Simulation*, 1998, pp. 154–161.
- [6] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing*, Imielinski and Korth, Eds. Kluwer Academic Publishers, 1996, vol. 353.
- [7] C. Perkins and P. Bhagwat, "Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers," in *ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, 1994, pp. 234–244.
- [8] Z. J. Haas and M. R. Pearlman, "The performance of query control schemes for the zone routing protocol," in *SIGCOMM*, 1998, pp. 167–177.
- [9] J.-C. Chen and S. Li and S.-H. Chan and J.-Y. He, "WIANI: Wireless Infrastructure and Ad-Hoc Network Integration," in *Proceedings of IEEE International Conference on Communications (ICC)*, 2005.
- [10] Ren-Hung Hwang and Chiung-Ying Wang and Cheng-Ying Li and Yuh-Shyan Chen, "Mobile IPv6-based Ad Hoc Networks: Its Development and Application," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 11, pp. 2161–2171, Nov 2005.
- [11] C. Ahlund and A. Zaslavsky, "Extending Global IP Connectivity for Ad Hoc Networks," *Kluwer. Telecommunication Systems, Modeling, Analysis, Design and Management*, vol. 24, no. 2-4, 2003.
- [12] Josh Broch and David A. Maltz and David B. Johnson, "Heterogeneous Interfaces in Multi-Hop Wireless Ad Hoc Networks," in *In Proceedings of The International Symposium on Parallel Architectures, Algorithms and Networks (ISPAN99), Workshop on Mobile Computing*.
- [13] P. Ratanchandani and R. Kravets, "A hybrid approach to internet connectivity for mobile ad hoc networks," in *Proceedings of IEEE WCNC*, 2003.
- [14] Y. Sun, E. Belding-Royer, and C. Perkins, "Internet connectivity for ad hoc mobile networks," in *International Journal of Wireless Information Networks, special issue on Mobile Ad hoc Networks*, 2002.
- [15] E. Nordström and P. Gunningberg and C. Tschudin, "Gateway Forwarding Strategies for Ad hoc Networks," in *Scandinavian Workshop on Wireless Ad hoc Networks*, May 2004.
- [16] R. Baumann, "VANET: Vehicular Ad Hoc Networks," Master's thesis, ETH Zurich, 2004.
- [17] V. Park and S. Corson, *Temporally-Ordered Routing Algorithm (TORA)*, IETF Internet Draft, July 2001.
- [18] A. Basu, A. Lin, and S. Ramanathan, "Routing Using Potentials: A Dynamic Traffic-Aware Routing Algorithm," in *Proceedings of the ACM annual conference of the Special Interest Group on Data Communication (SIGCOMM'03)*, Karlsruhe, Germany, August 2003.
- [19] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," in *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (Mobicom '00)*, Boston, USA, 2000.
- [20] J. Faruque and A. Helmy, "RUGGED: RoUting on finGerprint Gradients in sEnsor Networks," in *Proceedings of the IEEE International Conference on Pervasive Services (ICPS)*, Novi Sad, Serbia and Montenegro, July 2004.
- [21] S. Toumpis and L. Tassiulas, "Packetostatics: Deployment of Massively Dense Sensor Networks as an Electrostatic Problem," in *IEEE INFOCOM*, Miami, USA, March 2005.
- [22] Q. Li, M. D. Rosa, and D. Rus, "Distributed Algorithms for Guiding Navigation across a Sensor Network," in *Proceedings of ACM MobiCom*, San Diego, California, September 2003.
- [23] Vincent Lenders and Martin May and Bernhard Plattner, "Service Discovery in Mobile Ad Hoc Networks: A Field Theoretic Approach," in *Proceedings of the IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Taormina, Italy, June 2005.
- [24] V. Lenders, M. May, and B. Plattner, "Density-based vs. Proximity-based Anycast Routing for Mobile Networks," in *IEEE INFOCOM*, Barcelona, Spain, April 2006.
- [25] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing," RFC 3561 (Experimental), Jul. 2003. [Online]. Available: <http://www.ietf.org/rfc/rfc3561.txt>
- [26] Vincent Lenders and Martin May and Bernhard Plattner, "Density-based vs. Proximity-based Anycast Routing," in *Proceedings of the IEEE Infocom 2006*, Barcelona, Spain, April 2006.
- [27] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson, "Host Identity Protocol (HIP)," IETF draft-ietf-hip-base-00, October 2003.
- [28] C. Perkins, "IP Mobility Support," RFC 2002 (Proposed Standard), Oct. 1996.
- [29] V. Naumov, R. Baumann, and T. Gross, "An evaluation of intervehicle ad hoc networks based on realistic vehicular traces," in *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Florence, Italy, May 2006.
- [30] Federal Office of Topography, "Swiss Geographic Information System." [Online]. Available: www.swisstopo.ch
- [31] J.-Y. Le Boudec and M. Vojnovic, "Perfect Simulation and Stationarity of a Class of Mobility Models," in *IEEE Infocom*, 2005.
- [32] A. Feldmann, A. C. Gilbert, P. Huang, and W. Willinger, "Dynamics of IP traffic: A study of the role of variability and the impact of control," in *SIGCOMM*, 1999, pp. 301–313.
- [33] U. Fiedler, "Evaluating performance in systems with heavy-tailed input - a quantile-based approach," Ph.D. dissertation, ETH Zurich, Aug. 2003.
- [34] S. A. M. Corporation, "About Windows Media Audio Codec," October 2000. [Online]. Available: <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnwm/html/msaudio.asp>
- [35] P. Stokas, "OGG vs. MP3 vs. WMA vs. RA," 2002. [Online]. Available: <http://ekei.com/audio/>