Routing in Large-Scale Wireless Mesh Networks Using Temperature Fields

Rainer Baumann, Simon Heimlicher, and Bernhard Plattner, ETH Zurich

Abstract

Many wireless mesh networks are based on unicast routing protocols even though those protocols do not provide a particularly good fit for such scenarios. In this article, we report about an alternative routing paradigm, tailor-made for large multihop wireless mesh networks: field-based anycast routing. In particular, we present HEAT, a routing protocol based on this paradigm. In contrast to previous protocols, HEAT requires communication only between neighboring nodes. The underlying routing concept is a field similar to a temperature field in thermal physics. In extensive simulation experiments, we found that HEAT has excellent scalability properties due to a fully distributed implementation, and it provides much more robust routes than the unicast protocols, AODV and OLSR. As a consequence, in large-scale mobile scenarios, the packet delivery ratio with HEAT is more than two times higher, compared to AODV or OLSR. These promising results indicate that HEAT is suitable for large-scale wireless mesh networks that cover entire cities.

any IEEE 802.11 or WiFi access points are deployed on a daily basis, and most major cities already are covered by a dense mesh of such devices. Wireless access points provide a bridge between wireless and wired networks, typically connecting wireless equipment, such as laptops or PDAs to the world-wide Internet. Although some of these access points are installed as part of commercial wireless mesh networks, many access points are set up by private users and organizations striving for convenient Internet access. In the following, we refer to access points providing Internet access as gateways and to wireless devices demanding Internet access as mesh nodes or simply nodes. Because most privately-operated gateways are lightly loaded, their excess capacity could be leveraged to offer Internet access to other nodes that are in range at a negligible cost. Furthermore, the coverage of a gateway can be extended by having the nodes that are in range of the access point relay data on behalf of other wireless devices that are farther away. Such a scenario, where data is relayed among nodes to and from gateways, is called a multihop wireless mesh network. A multihop wireless mesh network can be an extremely cost-effective means to provide Internet access to wireless devices in cities. However, because the mesh nodes are commodity notebooks and hand-held or similar devices, carried and operated by humans, these nodes may move out of range or shut down at any time.

In general, mobile-to-mobile communication (i.e., communication among mobile devices) poses great challenges. Routing in wireless mesh networks is much easier if every mesh node is in range of at least one gateway and thus, only the last hop involves a human-operated device. In current wireless mesh networks, a dedicated wireless backbone network of a large number of stationary gateways provides this high level of coverage, albeit with a hefty price tag. We report about routing

protocols for such networks in the next section. Fortunately, novel routing paradigms, such as field-based anycast routing seem to make multihop wireless mesh networks feasible, and we present such a routing protocol. We report about our performance evaluation, and then we conclude the article.

Routing in Wireless Mesh Networks

Although routing in wireless networks has undergone extensive study, most wireless mesh networks are based on routing protocols that were originally designed for ad hoc networks, that is, small networks of mobile nodes that do not involve any infrastructure and where all nodes act both as routers and as end systems. Nordström et al. propose to use source routing based on the ad hoc routing protocol, dynamic source routing (DSR) [1]. Another popular ad hoc routing protocol, optimized link state routing (OLSR) [2] provides for interoperation with other networks by injecting external route information into the OLSR network. Because these protocols construct and maintain an individual unicast route from every mesh node to one of the gateways, the state information to be maintained increases with the number of nodes, as well as the number of gateways in the mesh network, and their scalability is limited.

The task group for mesh networking of the IEEE 802.11 working group also considers similar routing methods. In its first draft [3], it proposes to implement routing at the medium access control (MAC) layer. According to [4], the target size of an IEEE 802.11s wireless local area network (WLAN) mesh network is up to 32 static mesh gateways. In particular, the 802.11s task group specifies a default mandatory routing protocol called hybrid wireless mesh protocol (HWMP) that is inspired by a combination of the ad hoc routing protocol, ad hoc on-demand distance vector (AODV) [5] and tree-based

routing. In addition, the draft allows further standardized or vendor-specific path selection protocols. Up to now, the only alternative protocol described in the draft was the radio-aware optimized link state routing protocol (RA-OLSR).

Mosko et al. [6] propose to establish multiple non-disjoint paths. Although this may enhance the resilience against topology changes, this multipath unicast routing protocol is even less scalable than the single path unicast routing protocols discussed previously. In [7], scalability to the number of nodes is improved, based on geographical information; however, such information often is not available.

One challenging problem is that the scalability to the number of nodes of the described unicast routing protocols is limited. As we will see in the next section, the mesh network scenario lends itself well to anycast routing.

Anycast Routing

Anycast routing is aimed at networks where some client nodes require a route to any member from a certain group of service nodes. In the context of wireless mesh networks, the mesh nodes are the clients, and the gateways are the service nodes. Anycast routing was first proposed for IP networks; protocol implementations followed for wireless ad hoc networks [8]. However, these IP anycast routing protocols still are based on *unicast* routing techniques, such as link-state or distance vector routing, and as a consequence, they inherit the scalability problems of these protocols. IP anycast, in general, scales poorly to the number of groups because IP anycast addresses can not be aggregated into subnets.

However, in mesh networks, one anycast group representing the gateways to the Internet typically is sufficient, and scalability to the number of groups is not a concern.

To summarize, no established routing protocol tailor-made for large wireless mesh networks is readily available today. A routing protocol for wireless mesh networks should take advantage of the specific topology and traffic pattern of such networks. It must be scalable to the number of nodes, as well as to the number of gateways.

In the remainder of the article, we focus on multihop wireless mesh networks. Although the major portion of our elaboration applies to all mesh networks that involve mobile nodes, some of the characteristics are more accentuated in multihop mesh networks. In particular, multihop wireless mesh networks involve mobile-to-mobile communication, and it is imperative that the routing protocol be robust in the face of node mobility. Thus, an important feature of routing protocols for multihop wireless mesh networks is robustness against frequent changes in the topology of the network incurred by node mobility.

HEAT: Field-based Anycast Routing

The field-based or gradient-based routing paradigm, in general, has properties that are desirable in dynamic networks as it opens a greater design space than the traditional distance-vector or link-state routing paradigms. Lenders et al. proposed a model for anycast routing based on potential fields in [9] that uses flooding to establish the field. In this article, we have a closer look at *HEAT* [10], a protocol that aims to satisfy the requirements mentioned in the previous section. *HEAT* is a proactive field-based anycast routing protocol and shares with other field-based protocols the general forwarding principle. *HEAT* differs in how it establishes the field that defines the routes: its method to establish and maintain the field mimics heat dissipation in solids and is quite unique as it does not require flooding. Rather, nodes calculate their field value based solely on information from their immediate neighbors.

In multihop wireless mesh networks, multiple paths typically are available between a node and one or more gateways. It is the task of the anycast routing protocol to select a path according to a certain optimization goal. *HEAT* aims to select among the available paths the one that provides the maximal robustness against changes of the topology. Topology changes may be induced by node mobility or by temporary or permanent node failures. Furthermore, environmental influences also have a severe impact on the availability and lifetime of the wireless links between the nodes and the gateways.

HEAT is inspired by the physical laws that describe heat conduction. Similar to heat sources and surrounding particles, gateways and nodes define a field in the network. Gateways represent heat sources; nodes are assigned a temperature and conduct heat from the gateways to each other. The higher the temperature of a node, the closer it is to a gateway and the greater is the diversity of paths to this gateway. Based on this temperature field, a route is then defined as the path that follows the steepest gradient; in other words, packets always are forwarded to the neighboring node with the highest temperature and thus eventually reach a gateway. *HEAT* establishes the temperature field in the network based on purely local information, that is, every node calculates its own temperature solely by evaluating the temperature of its immediate neighbors. Because *HEAT* is an anycast protocol and only requires communication between immediate neighbors, it is well suited for large-scale applications, as it is scalable to the number of nodes and the number of gateways.

The Concept of HEAT

HEAT has two distinguishing features. First, in the routing decision, it considers both the length and the robustness of the available paths. Second, the field construction and maintenance mechanism of *HEAT* scales to the number of nodes and the number of gateways, as it only requires communication among *neighboring* nodes. These two features are tightly linked to the underlying routing concept that is inspired by temperature fields.

In brief, *HEAT* assigns a temperature value to every node in the mesh network. New nodes are assigned a value of zero; gateway nodes are assigned a well-defined maximum value. The temperature of nodes is determined, based on a simple yet effective that incorporates into the calculation the:

- Distances to the available gateways
- Robustness of the paths toward these gateways

That is, a path providing multiple alternative delivery opportunities along its way is preferred to a path over which packets cannot naturally be re-routed to an alternative path to one of the gateways. An example is depicted in Fig. 1. Note that there are many partly disjoint paths available leading to the gateway on the right-hand side, whereas only one path is available to the left-hand gateway. The temperature gradient as determined by *HEAT* is steeper toward the gateway on the right, and packets are routed in this direction, even if the network distance to the left gateway is shorter (measured in the number of hops). Only the packets from the two leftmost nodes are routed to the left-hand gateway.

Physical Analogy

A temperature field assigns a single scalar value to every particle in space. The temperature is higher in the vicinity of heat sources and then decreases with distance.

In a solid, heat is transferred by conduction. On a microscopic scale, conduction presents itself as hot, rapidly moving or vibrating atoms and molecules. By the interaction among neighboring atoms and molecules, heat is transferred. The

■ Figure 1. Example of a temperature field with a conductivity value of $\kappa = 1/4$ and areas of different link redundancies. The packets of the node with temperature value 0.051 are forwarded to the right, across an area of high redundancy instead of to the closest gateway at the left.

physical parameter *thermal conductivity*, κ , indicates its ability to conduct heat. The conduction of heat is governed by *Fourier's Law*. In essence, this law demands that the temperature of the field always decreases away from sources, resulting in a temperature gradient whose maxima are at the sources.

To map the properties of temperature fields to a given network topology, nodes in the mesh network are considered as particles, and gateways are considered as heat sources. In [9], Lenders et al. show that under the assumption that there are no local maxima in the field, following the path defined by the steepest gradient always leads to a gateway and that there are no loops in this path. However, not all policies for assigning scalars to nodes guarantee that there are no local maxima in the potential field. HEAT guarantees the absence of local maxima by adhering to the following policy: for every node, only neighbors with a higher temperature may contribute to the node's own temperature. This policy guarantees monotonicity of the field and thus ensures that there are no local maxima ([10]).

The HEAT Anycast Routing Protocol

According to the concept described previously, the gateways act as the heat sources of a temperature field, and the mesh nodes are assigned temperature values such that the optimal route toward any gateway is defined by the steepest gradient of the temperature field. To construct the temperature field starting from the initial temperature values of the gateways, the temperature values of neighbors are periodically exchanged between the gateways and neighboring mesh nodes through HEAT beacon messages. Based on these messages, every mesh node calculates its own temperature using the field calculation function.

After the field is constructed, the routing of packets from the mesh nodes to the gateways is straightforward and implemented on a hop-by-hop basis: A packet always is forwarded to the neighbor with the highest temperature, resulting in steepest-gradient routing. Routing back from the gateways to the mesh nodes is implemented as source routing from the gateways. All packets sent toward gateways record their route. Then, the source route for packets toward mesh nodes is constructed from the inverse of the path recorded by the last packet received from the destination node. A more thorough discussion of the backward path can be found in [11].

```

1:  $a = \text{sort}_{\text{ascending}}(\theta_0, \dots, \theta_n)$ 
2:  $j = 0$ 
3:  $t_j = 0$ 
4: while  $t_j < a[j]$  do
5:    $t_{j+1} = t_j + (a[j] - t_j) \cdot \kappa$ 
6:    $j = j + 1$ 
7: end while
8:  $t_{\text{final}} = t_j$ 

```

■ Algorithm 1. Temperature field calculation function.

Field Construction and Maintenance

As mentioned previously, the sources of the temperature field are the gateways. Therefore, each gateway initializes its temperature with a certain maximum value. For the heat propagation, every node (including the gateway nodes) periodically broadcasts its temperature value to its neighbors at a given HEAT beacon time interval. Based on these messages, all nodes build and maintain a data structure called a neighbor table that contains an entry for every known neighbor. Neighbor entries comprise the address, the last reported temperature, and a timestamp value of the corresponding node. Whenever an entry is added, removed, or changed, the temperature value is recomputed. In essence, we must differentiate among three cases:

- **New neighbor:** If a beacon from an unknown neighbor is received, a corresponding entry is added to the neighbor table.
- **Maintain neighbor:** If the reported temperature value of a known neighbor changes, the node recalculates its temperature value.
- **Missing neighbor:** If no beacon is received from a neighbor for a certain period, its entry is removed, and the temperature value is recomputed.

The detailed algorithm is described in Alg. 1. The algorithm calculates the temperature t_{final} of a node as follows: in a first step, the node sorts its neighbors, based on their temperatures θ_i , $i \in \{0, \dots, n\}$ in ascending order (line 1), into an array a . Then, it iterates over a accumulating the temperature of the next neighbor to the sum of the temperatures of the previous neighbors t_j until the temperature of the next neighbor is less than the accumulated temperature (line 4). In each step j , the value t_{j+1} is calculated as follows (line 5): the difference between the temperature of the currently considered neighbor, denoted by $a[j]$, and the temperature accumulated so far, t_j , is calculated. Then, this difference is multiplied by the conductivity parameter κ , and the result is added to the temperature accumulated so far, denoted by t_j .

As a result, nodes that have many neighbors that can reach one or more gateways obtain higher temperatures than nodes with only a small number of such neighbors. This effect is more pronounced the smaller the parameter κ that is chosen. Figure 1 illustrates an example temperature field with a rather small value of $\kappa = 1/4$. At this low κ , the greater link diversity in the right-hand side of the network has a considerable impact on the steepness of the temperature gradient. A step-by-step example of the field calculation function is given in Fig. 2 for $\kappa = 1/4$.

Expediting Convergence

A new node joining the network simply assigns itself a temperature of zero, broadcasts a HEAT beacon, and then waits for a beacon from one of its neighbors. The first arriving HEAT beacon provides the new node with a route to the Internet. As more beacons arrive, the node adjusts its temperature until the temperature converges to its final value.

Step 1		Step 2		Step 3		Step 4		Step 5	
NBR	FI	NBR	FI	NBR	FI	NBR	FI	NBR	FI
2	0.800	2	0.800	2	0.800	2	0.800	2	0.800
21	0.600	21	0.600	21	0.600	21	0.600	21	0.600
49	0.500	49	0.500	49	0.500	49	0.500	49	0.500
77	0.300	77	0.300	77	0.300	77	0.300	77	0.300
17	0.040	17	0.040	17	0.040	17	0.040	17	0.040
$t_0: 0.000$		$t_0: 0.000$		$t_1: 0.200$		$t_2: 0.300$		$t_3: 0.350$	
$t_1 = (0.800-0)/4+0$		$t_3 = (0.500-0.300)/4+0.300$		$t_2 = (0.600-0.200)/4+0.200$		$t_3 \geq 0.300$			

■ Figure 2. Example of the temperature field calculation with a conductivity value of $\kappa = 1/4$ for node 53: step 1, sort neighbors (nbr) by temperature value; steps 2–5 iterate down the table until the given temperature value of the node is higher or equal to the next neighbor; nodes 77 and 17 do not contribute to the temperature value of node 53: they will increase their values after the next HEAT beacon message of node 53 (node 77: 0.313 and node 17: 0.118).

In most cases, a node that disappears, for instance, by moving out of range, has only a local impact on the temperature field. As soon as a node detects that the neighbor with the highest temperature is no longer available, it selects the neighbor with the highest temperature among the remaining neighbors.

In rare cases, the disappearance of a node may cause network partitioning, and individual gateways may become unreachable. During the time it takes the temperature field to reconverge, some nodes may not be able to reach any gateways. To expedite convergence in such cases, HEAT uses so-called *early HEAT beacons*. When a node detects that a neighbor has disappeared and that this disappearance has a significant impact on its temperature, the node broadcasts an early HEAT beacon. To limit the overhead caused by early HEAT beacons, nodes that receive an early HEAT beacon wait for a short period (e.g., a few broadcast intervals) before forwarding it, allowing multiple messages triggered by the same event to be aggregated.

Load Balancing

Under the HEAT protocol, every gateway is assigned the maximal temperature at the beginning, but this value may be adjusted according to the load level of the gateway. This enables the gateway to avoid congestion among the mesh nodes in its vicinity and also to adjust the total traffic to the bandwidth of its Internet access link. Note that in the evaluation we present in the next section, we do not change the temperature values of the gateway because such dynamic adaptation may not be feasible in some scenarios.

Evaluation Methodology

To evaluate the performance of HEAT, we ran an extensive simulation study with Glomosim, a network simulator for

wireless networks. We compared the performance of HEAT with the popular AODV and OLSR ad hoc routing protocols in versions that were extended, such that they can be used in multihop mesh networks, as described next. Note that in the highly dynamic scenarios we consider, the link lifetimes are too short for advanced routing metrics such as ETX to be useful.

All simulations have a duration of at least 10,000 seconds; the reported values are averaged over at least 20 simulation runs with different random seeds.

Extended AODV

As a first reference for the performance of HEAT, we use an extended version of AODV. AODV is a reactive routing protocol and establishes a route to a destination only on demand. A node that requires a route broadcasts a route request to its neighbors, which forward this message and record the node from which they received it. This creates a number of temporary routes back to the requesting node. As soon as a node that already has a route to the destination node receives the route request, this node sends back a message through the temporary route to the node that requested the route, which then selects among the received replies the route with the least number of hops. To enable the use of AODV in the mesh scenario, we extended the standard implementation of AODV included in

Glomosim according to [12] to support gateway discovery in mesh networks. As proposed in the cited paper, all gateways are connected to a dedicated router that acts as a proxy to the Internet. This router has two tasks:

- On the forward path, it sends route replies on behalf of hosts in the Internet.
- On the backward path, it initiates route requests for nodes in the wireless mesh network.

Thus AODV does not have to distinguish between the different gateways, and only a common route to the Internet must be maintained, the route to the dedicated router.

OLSR

Second, we compared HEAT with the OLSR implementation from the University of Niigata. OLSR is a proactive link-state routing protocol, which means that it floods a complete topology description across the network, and every node computes the optimal forwarding paths locally. OLSR allows redistribution of routing information from so-called non OLSR interfaces as the gateway uplink interface to the Internet. Using simulation experiments, we found that the performance of OLSR drops quickly with increasing mobility. We assume that this is in part due to the long hello interval of two seconds. To achieve a fair comparison with HEAT that has a beacon interval of one second, we adjusted the hello interval of OLSR also to one second. With this adjustment, the performance of OLSR improves by roughly 10 percent, and we used this setting for all experiments presented in this article.

Simulation Settings

The simulation experiments are based on a WiFi network. All nodes were equipped with an IEEE 802.11b radio with a nominal bandwidth of 11 Mb/s and a maximum range of 250 m. As a MAC layer protocol, we used the 802.11 DCFMAC-DCF with the request-to-send/clear-to-send (RTS/CTS) handshake.

■ Figure 3. Impact of the network size: a) packet delivery ratio; b) routing overhead.

As the radio propagation model, we used the two-ray ground model.

Mobility Model

Because simple mobility models, such as the random waypoint or the random walk mobility models, often are reported to lack realism, we used our own, more realistic mobility model [10]. This model allows nodes to move only along roads defined by a road map of a real city. We extracted these road maps from the Swiss geographic information system, as this database includes vectored building and road maps that are accurate within less than one meter. Furthermore, this database also provides speed limit information for all roads.

The actual node movement is modeled according to the steady-state random trip mobility model [13] on the road maps. That is, a node chooses a random destination in the city and moves to this position at a constant speed along the fastest path. Note that we deliberately did not introduce any pausing of the nodes, and therefore, a node began to move to a new destination as soon as it arrived at the target position. The city mobility model was applied for pedestrians, as well as cars because the movements of both are constrained by the available roads. Cars are further restricted by the speed limits on all roads.

Traffic Pattern

Wireless mesh networks are used mostly for Internet applications, such as Web browsing, messaging, chatting, and so on. We used an Internet traffic model consisting of a mix of streaming and Web browsing traffic, as described in [10]. All traffic in the simulations was between nodes in the wireless mesh network and hosts in the Internet, and there was no communication among the mesh nodes.

Evaluation Results

We used the following metrics to compare the performance and scalability of HEAT with OLSR and AODV.

- The *packet delivery ratio* denotes the ratio between the number of packets that are successfully received and the total number of packets sent. This metric comprised the data packets sent from the mesh nodes to the gateways, as well as packets from the gateways back to the mesh nodes.
- The *routing overhead* refers to the average number of routing control messages sent per node and per second.

Scalability to the Network Size

In the first experiment, we evaluated how the performance is affected by increasing the network size at a constant average node degree. The node degree is kept constant by increasing the simulation area (the section of the maps we considered) in parallel with the number of nodes. The results for a *static* scenario are shown in Fig. 3. The nodes were placed randomly, 100 of the nodes were active, that is, they generated traffic, and five Internet gateways were available. The average node degree was approximately six. The upper part of this figure shows the delivery ratio. With HEAT, this ratio remained constant at almost 100 percent even as the network size increased to 2000 nodes. With OLSR, the delivery ratio decreased, but only marginally. With AODV, the delivery ratio dropped significantly at network sizes greater than 500 nodes. The routing overhead, as shown in the bottom plot of Fig. 3, indicates that the reason for the performance degradation of AODV was related to this metric. As the network size increased, the average distance between the data sources and the Internet gateways also became longer. This increase in the length of the shortest available path forced AODV to increase the scope of its route discovery procedure [5], and it ended up using the network capacity mostly for the flooding of control messages. The overhead under OLSR increased slightly because — being a link-state routing protocol — OLSR requires complete knowledge about the whole topology to calculate the shortest path. This result shows that the hierarchical flooding mechanism used by OLSR mitigated the scalability problem as compared to AODV. HEAT achieved the best result with a constant overhead per node, independent of the network size.

Effect of Node Mobility

In the second experiment, we investigated how node mobility affects the routing performance. We considered two scenarios:

- A scenario with mobile nodes moving at pedestrian speeds (i.e., node speeds that are uniformly distributed between 0.5 m/s and 3 m/s)
- A scenario including nodes moving at car speeds in a city (i.e., node speeds between 10 m/s and 20 m/s)

The results for the pedestrian scenario with a simulation area of 5 km by 5 km, five gateways placed at strategic positions, and 100 active nodes are given in Fig. 4. At this rather low node speed, the packet delivery ratio of HEAT was almost as high as in the static scenario (shown previously in Fig. 3), but the routing overhead was slightly higher. The increasing overhead originated from the early HEAT beacons.

■ Figure 4. Mobile scenario at pedestrian speeds: a) packet delivery ratio; b) routing overhead.

■ Figure 5. Mobile scenario at car speeds: a) Packet delivery ratio; b) routing overhead.

The results of OLSR revealed that its packet delivery ratio decreased slightly for nodes moving at pedestrian speed, particularly in larger networks with longer routes. AODV is most affected by the node mobility and its delivery ratio already dropped sharply at 1000 nodes.

Figure 5 shows the performance at car speeds using the same settings. At these node speeds, the performance of all three protocols was lower than at pedestrian speeds. The packet delivery ratio of HEAT remained above 70 percent for all network sizes we evaluated. OLSR, as well as AODV, suffered much more, and at a network size of 2000 nodes, neither of them delivered more than 50 percent of the packets. OLSR was affected heavily by mobility because it must propagate information about link state changes through the whole network. HEAT, in contrast, only requires local information exchange, and the early HEAT beacon mechanism accelerates the convergence of the temperature field. Again, AODV performed the most poorly of the three, presumably due to the high overhead of route discovery broadcast messages.

Effect of the Number of Gateways

To conclude our evaluation, we looked at the effect of the number of gateways in the mesh network. In Fig. 6, the packet delivery ratio and the routing overhead are plotted for the pedestrian scenario with 1000 nodes, 100 of which generated

traffic. A total of 1 to 30 gateways are placed randomly over the entire simulated area of 5 km by 5 km. The packet delivery ratio rose with the number of gateways. This is mainly because the average distance between mesh nodes and gateways became shorter when the number of gateways was increased. Therefore, the average path length is shorter, and the paths are less prone to link failures caused by mobility. Furthermore, when the number of gateways was too small (e.g., only one gateway), the capacity of the radio interface at the gateway(s) became a limiting factor. In other words, the available capacity of the gateway(s) was not sufficient to support all the traffic generated by the mesh nodes.

Considering the number of gateways required for an average packet delivery ratio of at least 99 percent, we found that with HEAT, five gateways are sufficient. OLSR achieved a packet delivery ratio of only 91 percent with this number of gateways; adding more gateways helped only slightly, because the limiting factor of OLSR in mobile scenarios is that routes fail frequently, and it does not discover and replace invalid routes quickly enough. With only five gateways, AODV achieved a delivery ratio of less than 50 percent; increasing the number to 30 gateways brought the delivery ratio to a still rather low 90 percent. We concluded from this experiment that in mobile scenarios, OLSR and AODV require many more gateways than HEAT to achieve a delivery ratio close to

■ Figure 6. The effect of the number of gateways (mobile scenario at pedestrian speeds): a) packet delivery ratio; b) routing overhead.

99 percent. Thus, HEAT appears particularly suitable for deployments where the number of gateways is a crucial figure.

Conclusion

In this article, we report about established and novel ways to route data in wireless mesh networks. We discuss the state of the art and then go on to report about a recently published routing protocol for multihop wireless mesh networks called HEAT. In contrast to established routing protocols, HEAT uses the field-based anycast routing paradigm. Using anycast makes HEAT particularly scalable and suitable for very large and dense mesh networks. The field-based routing algorithm of HEAT provides for robust routes even with mobile nodes moving at car speeds.

We compared the performance of HEAT with AODV and OLSR through extensive simulation experiments. In a large static mesh network scenario with 1000 nodes covering an area of 5 km by 5 km, we found that both HEAT and OLSR achieved packet delivery ratios above 95 percent, whereas AODV did not reach more than 30 percent. To evaluate the performance in mobile scenarios, we used a realistic mobility model based on road maps from real cities. With nodes moving at car speed, HEAT outperformed both AODV and OLSR in terms of packet delivery ratio by more than a factor of two. Because the number of gateways determines the cost of a mesh network to a large extent, we evaluated how many gateways are required to achieve a packet delivery ratio of at least 99 percent. Under HEAT, this delivery ratio is reached with five gateways. OLSR achieved a delivery ratio of only 91 percent with five gateways; with 30 gateways, it was still below 95 percent. AODV delivered only 42 percent of the packets with five gateways; 30 gateways raised the delivery ratio to 90 percent.

We conclude that novel routing paradigms, such as the field-based anycast routing concept employed by HEAT may contribute to more affordable wireless mesh networks in the near future. To what extent the results of our simulation experiments are applicable to real-world networks is difficult to determine because large-scale mobile testbeds are not available yet.

References

- [1] E. Nordström, P. Gunningberg, and C. Tschudin, "Gateway Forwarding Strategies for Ad Hoc Networks," *Scandinavian Wksp. Wireless Ad Hoc Networks*, May 2004.
- [2] T. Clausen and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)," RFC 3626 (experimental), Oct. 2003.
- [3] IEEE 802.11s TGs, "Draft Amendment to Standard IEEE 802.11: ESS Mesh Networking," Tech. rep. D0.01, 2006.

- [4] J. Hauser, D. Baker, and W. S. Conner, "Draft PAR for IEEE 802.11 ESS Mesh," IEEE P802.11 Wireless LANs, tech. rep. 11-04/0054r2, 2004.
- [5] C. E. Perkins, E. M. Belding-Royer, and S. R. Das, "Ad Hoc On-Demand Distance Vector (AODV) Routing," IETF Internet draft, Nov. 2002.
- [6] M. Mosko and J. Garcia-Luna-Aceves, "Multipath Routing in Wireless Mesh Networks," *Proc. IEEE Wksp. Wireless Mesh Networks*, Santa Clara, CA, Sept. 2005.
- [7] B.-N. Cheng, M. Yuksel, and S. Kalyanaraman, "Orthogonal Rendezvous Routing Protocol for Wireless Mesh Networks," *Proc. ICNP*, Santa Barbara, CA, Nov. 2006.
- [8] V. Park and J. Macker, "Anycast Routing for Mobile Services," *Conf. Info. Sci. and Sys.*, Baltimore, MD, Mar. 1999.
- [9] V. Lenders, M. May, and B. Plattner, "Density-based vs. Proximity-based Anycast Routing for Mobile Networks," *IEEE INFOCOM*, Barcelona, Spain, Apr. 2006.
- [10] R. Baumann *et al.*, "HEAT: Scalable Routing in Wireless Mesh Networks Using Temperature Fields," *IEEE WoWMoM*, Helsinki, Finland, June 2007.
- [11] R. Baumann *et al.*, "Routing Packets into Wireless Mesh Networks," *IEEE WiMob*, White Plains, NY, Oct. 2007.
- [12] M. Michalak and T. Braun, "Common Gateway Architecture for Mobile Ad-Hoc Networks," *Proc. 2nd Annual Conf. Wireless On-Demand Network Sys. and Svcs.*, Washington, DC, 2005, pp. 70–75.
- [13] J.-Y. L. Boudec and M. Vojnovic, "Perfect Simulation and Stationarity of a Class of Mobility Models," *IEEE INFOCOM*, 2005.

Biographies

RAINER BAUMANN (baumann@tik.ee.ethz.ch) obtained his M.Sc. degree in computer science with a major in software engineering and a minor in physics from Swiss Federal Institute of Technology (ETH), Zurich, in 2004. In parallel with his Master's degree, he finished his Master of Advanced Studies in Secondary and Higher Education. In autumn 2004 he joined the Computer Engineering and Networks Laboratory at ETH Zurich as a Ph.D. student. In autumn 2007 he received his D.Sc. degree from ETH Zurich. During his doctoral studies he focused his research on mobile, wireless, mesh, and ad hoc networking. Since autumn 2005 he has been a lecturer at the Applied University of Rapperswil, HSR, and recently was appointed as a lecturer at ETH Zurich.

SIMON HEIMLICH (heimlicher@tik.ee.ethz.ch) received a Master's degree in electrical engineering in 2005 from ETH Zurich. Currently, he is pursuing his Ph.D. degree at the same institution in the communication systems research group lead by Professor Dr. Plattner. His research interests are data transport in intermittently connected networks and delay- and disruption-tolerant networking. He is a member of the ACM.

BERNHARD PLATTNER (plattner@tik.ee.ethz.ch) is a professor of computer engineering at ETH Zurich, where he leads the Communication Systems Group. He has been the principal investigator (PI) or co-PI of numerous national and international projects in the area of computer networking. His current research interests are in self-organizing networks, mobile ad hoc networks, and practical aspects of information security. He has also directed research on active networks, starting as early as 1996, and multimedia applications for high-speed networks. From 1996 to 1998 he served as the head of faculty of electrical engineering at ETH Zurich. From 2005 to 2007 he was vice-rector of ETH for Bachelor/Master studies. He is a member of the ACM and the Internet Society. He has served as the program or general chair of various international conferences, such as ACM SIGCOMM '91, INET '94, and IWAN '02, and has served on the program committees of other major conferences, such as IEEE INFOCOM.