Available online at www.sciencedirect.com

SCIENCE DIRECT®

Ad Hoc Networks 3 (2005) 221–233

Ad Hoc Networks

www.elsevier.com/locate/adhoc

Lessons from experimental MANET research

Christian Tschudin [a,*], Per Gunningberg [b], Henrik Lundgren [b], Erik Nordström [b]

[a] Computer Science Department, University of Basel, Bernoullistrasse 16, CH-4056 Basel, Switzerland
[b] IT Department, Uppsala University, Box 337, SE-751 05 Uppsala, Sweden

Abstract

After almost a decade of research into ad hoc networking, MANET technology has not yet affected our way of using wireless networks. In this paper we discuss lessons to draw and back them with experiences from our experimental work. We find that simulations have to be complemented to a much higher degree by real-world experiments, that there is a lack of mature implementations and integration and that efforts should be focused on more realistic settings inside the "ad hoc horizon" where decent network services still can be provided.

Keywords: Wireless ad hoc networks; MANET; Simulation; Real world experiment; Internet integration

1. Introduction

Since the initial work on an IETF charter for "Mobile Ad hoc NETworks" started around 1995/1996, the IETF MANET working group [14] has been working towards a proposed standard for IP level routing in wireless ad hoc networks. The IETF MANET working group charter as of 2004 has proposed four protocols for experimental RFC status, namely AODV [17], DSR [8], OLSR [5] and TBRPF [15]. This would indicate a new level of maturity of multihop ad hoc networking.

However, although the IETF MANET research effort has achieved respectable progress, we believe half of the job still remains. There are no clear results that show how well MANETs will work in reality. Simulations have not been conclusive on clearing the field of IETF MANET protocol candidates. We point out the fact that MANET solutions, if they exist, are not available for use in our daily network experience. Furthermore, reports from the users are missing because there is virtually no deployment of multihop ad hoc networks.

In this paper we argue that more realistic ad hoc networking research efforts should be conducted that focus on small scale and complete solutions. We organize our arguments along the work carried out at Uppsala University. Many other research

---

[*] Corresponding author.
E-mail address: christian.tschudin@unibas.ch (C. Tschudin).

1570-8705/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.adhoc.2004.07.007

groups have also done extensive systems research on wireless ad hoc networking. What is still lacking though, and what we address in this paper, is a more comprehensive review of the original MANET goals and the results obtained so far. Often, research results are focused on small aspects of the big MANET picture. These isolated advances have to be contrasted with the realistic assessment that MANET results have had little impact on the networking mainstream.

For example, we observe that IEEE 802.11 has become ubiquitous, but that users seldomly operate this hardware in ad hoc mode. The lack of robust ad hoc routing software and drivers in mainstream operating systems, along with the lack of reliable ad hoc solutions add to the acceptance problem of this technology. There seems also to be a mismatch between what end users might find useful and what research problems are currently being addressed. As a consequence, the cost of using ad hoc networking is bigger than the potential benefit.

A simple use case is web browsing and email download in multihop access networks based on IEEE 802.11. However, this is not the prominent model that is valued in the competition among routing protocols. Instead, proposals how to change the media access control scheme (MAC) (and thus to invalidate the deployed hardware) and 1000 node scenarios with constant bit rate traffic (CBR) are prevalent. Moreover, as we point out in this paper, current technology forms an ad hoc horizon at two to three hops and 10–20 nodes where the benefit from wireless multihop ad hoc networking virtually vanishes. The focus should therefore be to get the most out of ad hoc networking *within* these limits. A scenario consisting of a few people wanting to form an ad hoc network and sharing access to the Internet is simple, but much more probable and feasible. Managed mode 802.11 already anticipates this use and provides the possibility to extend the reach of access points by using a wireless repeater. MANET style multihop forwarding would permit people to extend the range of hardware that they already own, without depending on a network operator to deploy such relays.

Working towards lowering the barriers for ad hoc networking has to become a prime objective of MANET research. By hammering out ‘‘good enough’’ solutions it should be possible to get a first foot into the end user devices. Later on, more sophisticated solutions can be envisaged that also include progress as it becomes available in the area of security, QoS and power consumption. However, we argue that even for a first ‘‘good enough’’ solution, the hurdles are still high.

The rest of this paper is organized as follows. In Section 2 we start by outlining the key areas of focus that we believe are important for advancing MANET research. In Section 3 we present lessons learned from our experimental research. We revisit some of the underlying ideas and premises behind MANET thinking in Section 4, where we present work that we have done outside the scope of the IETF MANET working group. This work emphasizes simplicity, easily deployable software and a focus on small ad hoc networks. Section 5 continues with discussing challenges in MANET research before we conclude this paper in Section 6.

2. MANET research approach

While several thousand research papers on wireless ad hoc networking and MANET have been published, there has been virtually no transfer from research into products at a large scale. Only very few prototype implementations have been made available. The tool set of exploration consists mostly of simulations and little experimentation. The integration into the networking environment of end users has not been adequately addressed till today. These are the three main deficits that we identify and discuss in this section.

2.1. Implementation

The adoption of wireless ad hoc solutions has been slow. A major reason is the lack of quality real world implementations. When we started our Ad hoc Protocol Evaluation testbed activity in the year 2000 [9,1], we had a hard time finding any implementations that worked properly. Out of three openly available protocol implementations (OLSR, AODV, TORA), only one was able to route packets over multiple hops. One implementation even had problems in a single hop scenario.

Today, AODV and OLSR are the most mature protocols whilst other MANET protocols either lack updated implementations or have intellectual property issues and suffer from the reluctance of their developers to share work and promote open development. In the case of TBRPF there are no implementations freely available. [1]

2.2. Experimentation

Experimentation is needed to move technology from paper specification to functional systems. It allows drafts to be closely examined and increases the likelihood that invalid assumptions are discovered. Although MANET research has been ongoing for some time, there is little valuable "hands-on" experience. Instead, a large portion of protocol development is done in unrealistic simulation settings. As an example, routing protocols in simulation typically implement layer interactions (e.g., feeding link layer failures up the networking stack) that are not readily available in existing real world implementations.

Another problem is that simulations are almost never validated. It is tempting to do all research through simulations because they save time compared to real experiments and are easily reproducible. But the protocol implementations that come with the simulation packages are seldom assessed and verified independently. Furthermore, simulation results are rarely compared or calibrated with findings from actual measurements.

2.3. Integration

A consequence of the currently narrow implementation base is that the integration of protocol logic into actual operating systems is similarly scarce. Moreover, deploying ad hoc routing protocols is not just about implementing a new algorithm: there are also operational matters to consider and configuration issues. Unfortunately, questions such as address assignment or Internet connectivity and hand-over between gateways—although important for the end user's service satisfaction—are beyond routing protocol specifications. Self-configuration solutions for MANETs still need to be identified and to be proven in the field.

These three innocuous looking problem areas—implementation, experimentation and integration—are preventing wireless ad hoc networking to reach the maturity needed for a broader acceptance. We believe that putting more emphasis in these three areas will help to produce more focused and applicable research results, as we argue in the following two sections.

3. Lessons from experimental research within the IETF MANET area

In this section we present our lessons learned from conducting real world experimental research within the IETF MANET area. First, we discuss issues with designing testbeds and conducting large scale tests of MANET routing protocols. Second, we address the importance of implementing protocols in real systems. Third, we describe why we believe that real-world experiments and simulations must complement each other. Finally, we point to usability limits in form of an "ad hoc horizon" which highlights the importance of focusing on achievable goals in MANET research.

3.1. APE testbed—repeatable real world testing

Most wireless testbeds [13,12,20,18] target a single protocol, and do not represent a software product which can be used by others for reproducing experimental results. By developing the Ad hoc Protocol Evaluation testbed (APE) we aimed at creating a tool that ultimately would be used to compare all MANET implementations on a large scale.

To this end we adopted a modular architecture that makes it possible to plug in *different* protocols and traffic generators. Because of logistic difficulties when setting up experiments with 10 or more human participants, APE is designed for easy installation, deployment and handling. Installation is as simple as downloading and extracting a single

---

[1] A TBRPF implementation was previously available from SRI, but has since then been retracted.

file. This package contains a self-extracting stripped-down Linux system that can be installed and booted on machines running Windows or Linux. It is also possible to run APE from a bootable CD. This has shown to be a successful approach, as it has allowed us to scale experiments to 37 individual participants with hardware and people being the limiting factors.

The intrinsic difficulty with a real testbed compared to simulation is the repeatability of the experiments. The nodes need to move in the same pattern, during the same time span, while running the same application from one experiment to another. To address repeatability we use strict choreography instructions to the participants who carry the laptops. These instructions are given in real time on the screen and constitute the movements. The participants only need to follow the instructions.

During the experiments all measurement data are collected at all nodes, including the signal quality of received packets. This data is uploaded to a central server after test-run and allows us to build a time dependent connectivity map of nodes. By comparing signal quality data from one experiment to another we can assure ourselves that the movements are sufficiently similar for comparison purposes. In our experiments, we have successfully used signal quality based "fingerprints" to assess the repeatability.

One such fingerprint is "virtual mobility". The idea behind this metric is that the signal quality as measured at reception time can be related to the (virtual) distance to the sender. To this end we use a simple path loss model for the far field which we calibrated using actual measurement data from Orinoco cards. The distance D in the range of 0.5–65 m is defined as

$$D_j(node_i) = 4 * 10^{(40-0.9*Q_j(node_i))/33}, \tag{1}$$

where Q is the signal quality $(0\ldots75)$ for a packet received from node j at node i. The derivatives of the virtual distances between all node pairs are then aggregated and form a network wide "virtual mobility": It represents how an average node moves during a test. The upper and lower quartiles of the mean value reflect the movement heterogeneity and can reveal different movement patterns within the network. Although the computed "distances" do not reflect true geographic distance during an experiment, as the signal quality depends on many other factors like walls or persons moving by, it captures the world as the routing protocol perceives it and is therefore a good measure for assessing repeatability.

Figs. 1 and 2 visualize fingerprints from two runs which are compared to ensure the fidelity of movements. "Virtual mobility" (Fig. 1 showing average, lower and upper quantile) captures the continuous changes in signal quality between nodes, while the "link change" metric (Fig. 2) relates to the routing level view of connectivity. The movement scenario "Double Lost and Found" consists of three groups of nodes initially positioned at the same location. At a given time

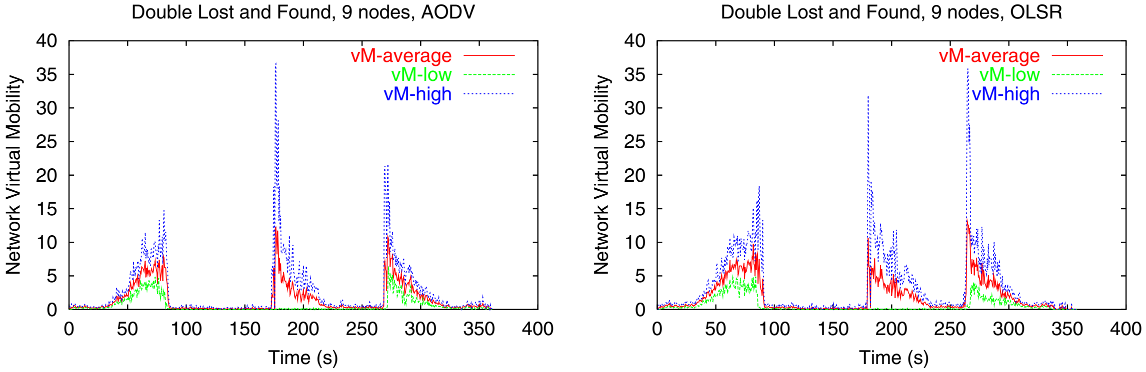

Fig. 1. Virtual mobility fingerprints from two test-runs (AODV and OLSR) of the "Double Lost and Found" scenario, where three groups of nodes first split up and later on re-join at different times.

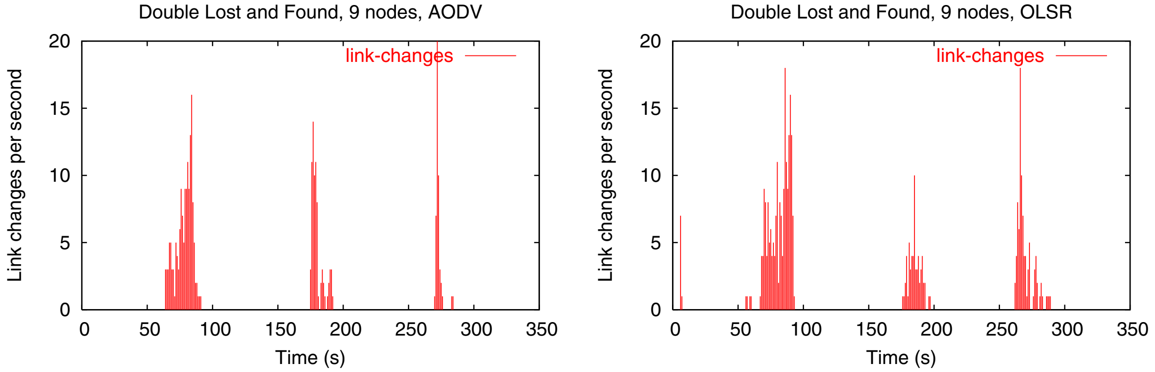

Fig. 2. Link change fingerprints for the "Double Lost and Found" scenario.

two groups move away to a location outside the other groups' radio range. Later on, the two groups re-join to their initial position at *different* points in time. The figures consistently show three peaks that relate to these movements of nodes. As timing and extent of the peaks correspond well in all four graphs we considered these two test-runs as a valid base for comparing the two routing protocols.

This test-run was used to compare the Ping success ratios of AODV and OLSR. The Ping traffic in this "Double Lost and Found" scenario is as follows: One third of the Ping sessions are between nodes in the two mobile groups, while the other Ping sessions are between nodes in a mobile group and nodes in the stationary group. During separation of the two mobile groups there is a period with two-hop connectivity via the stationary group before the mobile groups become isolated. At first sight, the graphs in Fig. 3 seem to indicate similar behavior for both protocols. However, closer analysis of the path lengths showed that the AODV implementation, Madhoc-AODV, had serious problems with multihop paths, wherefore conclusions about the relative performance of the two protocols could not be drawn from this experiment. Ultimately, this triggered our own implementation work on AODV that we describe in the following section. The APE testbed itself remained an important evaluation tool and, for example, enabled us to isolate the Gray Zone problem described in Section 3.3.

3.2. *AODV-UU—protocol implementation*

Most protocols are first implemented in simulation environments and therefore run the risk of being tailor made to the particular environment

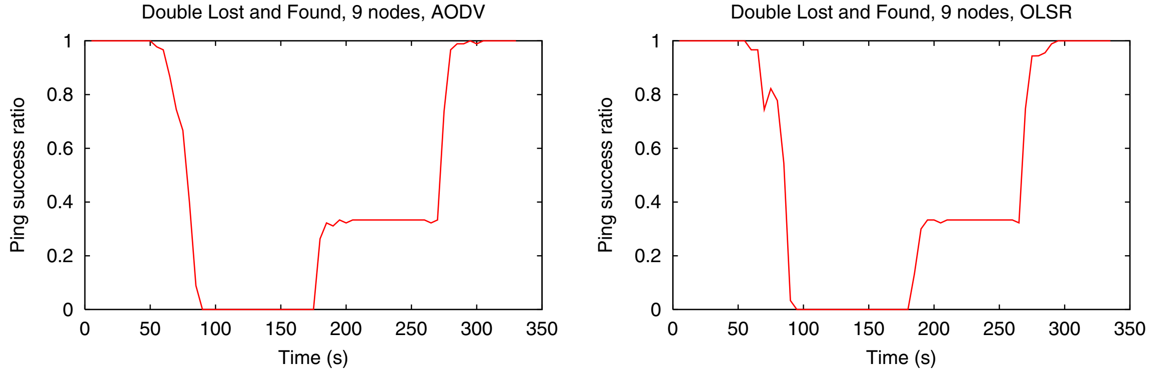

Fig. 3. Ping success ratios in the "Double Lost and Found" scenario for AODV and OLSR.

at hand. This means that the impact of the implementation environment is quite often set aside in a rush to reach the standards body. All protocols need to interact with the existing networking stack, the operating system and the actual physical media. All these components are simplified in a simulation environment.

Our goal with AODV-UU [4] was to create an up-to-date and stable implementation of the AODV routing protocol that we could use in the APE testbed. Later we wanted to bridge the gap between real world tests and simulations, and decided to port AODV-UU to the ns-2 simulator. Having the same protocol logic (i.e., the actual code) running in both real world and in simulation has proven very successful. Bugs and problems with larger scale networks are easily handled in a fully controlled simulation environment. At the same time, we can validate that the protocol also functions well in a real environment.

The lesson from the work with AODV-UU is that lack of implementation experience and unrealistic assumptions may lead to protocol designs that are not feasible to implement in a real system. We have encountered several such designs. For example, the AODV specification suggests to use link-layer feedback for monitoring the link status of active routes. If that is not possible, it should rely on periodic HELLO beacons. In simulation, link-layer feedback is easy and therefore commonly used. But in real implementations of IEEE 802.11 it is rather difficult for upper layers to acquire such information reliably. This observation impacts the conclusions from protocol simulation studies that use link-layer feedback. To more accurately model reality, protocols should be re-evaluated and possibly re-simulated with only HELLO messages.

3.3. Gray zones—reality complements simulation

Exposing protocol designs to the real world will not only highlight practical limitations, but will also reveal protocol behavior that might not show up in simulations due to inaccurate models. During experimentation, we observed spurious performance degradations with the AODV-UU implementation. In a simple scenario with three static and one mobile node (see Fig. 4), we experienced "communication gray zones" [11,10] which were several meters wide. In these zones, the routing protocol would report a valid route to the destination node, but virtually no data packets would pass through. This effect did not show up when we recreated the setting inside the ns-2 simulator, using the same AODV-UU code as in the real world test.

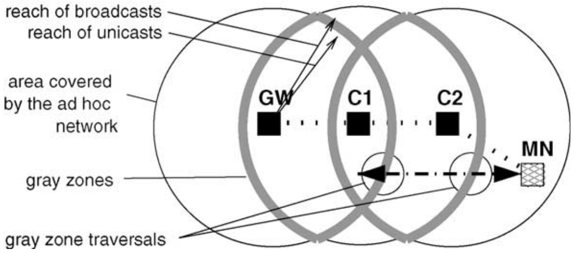

Fig. 4. Communication gray zones in a scenario with three stationary nodes and one mobile node which roams the the network back and forth. The mobile node (MN) communicates with the gateway node (GW).

Further experiments revealed that these gray zones are caused by the difference in transmission range for different packet types. In IEEE 802.11b, broadcast packets are sent at 2 Mbit/s while unicast packets can be sent at speeds up to 11 Mbit/s. The transmission output power is constant and therefore lower transmission speed means more energy per bit, which translates to longer transmission range, as was recently confirmed in [2]. As a consequence, broadcasted HELLO packets (used for neighbor sensing) will reach further than data packets. The routing daemon believes that the other node is reachable, thus it will *not* try to discover a new route although regular data packets might not reach the other node. This gray zone effect did not show up in the simulations simply because ns-2 uses a flat 2 Mbit/s model which gives data packets the same range as HELLO packets.

As a result of our findings, we started to think about work-arounds. In [10] we showed that artificially limiting the range of AODV HELLO packets by filtering on the SNR (signal to noise ratio) value eliminates the gray zone effect. In fact, setting the SNR acceptance level such that HELLO packets have slightly shorter transmission range than the data packets will force AODV to pick a more robust link earlier and increase the overall performance. Table 1 summarizes our results from experiments where the mobile node has two-way (Ping) and one-way communication (MP3 stream) with the gateway node (GW). The routing protocols tested were OLSR [5], LUNAR [22] and the two variants of AODV. Adding the SNR filter to AODV increased the Ping success ratio from 91.9% to 99.1%.

The accuracy of the simulation model for the physical layer is equally important. Such a model may include SNR calculation, signal reception, fading and path loss. The study in [19] reveals how important the accuracy is. Its authors examined the interaction between two routing protocols and the physical layer models used in ns-2 and GloMoSim. They observed that the physical model had great influence on the result; a plausible change in the simulation environment could even change the relative ranking of the compared protocols.

In conclusion, it is important to realize that simulation models are just models. Although simulation studies are useful in the protocol development process they do not replace real-world experimentation. In fact, we need more real-world experience earlier in the process, both to understand the interaction with a real environment as well as for validation of simulation models.

3.4. The ad hoc horizon—focus on achievable goals

We believe that there exists an "ad hoc horizon", i.e., a range that limits the number of hops and the number of nodes for functional ad hoc networking. Beyond this point, network services fail and the wireless network is not usable anymore. In real applications, it may therefore not be useful to provide MANET services when this range is exceeded. Our conservative estimate is that this horizon, using IEEE 802.11 technology and standard Internet applications like web browsing, is as narrow as 2–3 hops and a dozen of nodes. The horizon hypothesis stems from experimental work: The challenge was to demonstrate and quantify this effect using the field's standard investigation tool—simulations.

That TCP traffic suffers in a wireless multihop environment has been documented several times [6,7]. Typically, these studies do not show any abrupt service degradation. For example, it is reported that TCP throughput is inversely proportional to the hop distance for a CSMA/CA style MAC layer. However, most studies work with simple network topologies like a linear sequence of nodes. In order to cover a broader set of topologies, we created a family of static "beam star" scenarios [23] (see Fig. 5). The beam star topologies let us study the effect of increasing path lengths and growing number of nodes in a controlled manner. We expect both to have a negative impact on TCP performance. Thus, we can assess the combined limit for effective data delivery in multihop ad hoc networks.

Common to all beam star scenarios is that each beam is hosting one TCP session where the beam's end node communicates with the center node. Typically, the center node would be a gateway to the fixed Internet. By varying the number of beams and hops we can plot a service quality surface as in Fig. 6. In this diagram, the x-axis (one beam) represents all "string topologies" from 1 to 10 hops, while the y-axis (one hop) represents topologies where all nodes are within a one-hop distance. Inside the quadrant, we have mixed topologies with a cluster around the center node from which beams emanate.

The service quality criteria used in our simulations is the worst accumulated no-progress time among all TCP sessions of the beam star scenario in relation to the simulation's duration. Assuming some user tolerance to intermittent delays in data delivery, we only considered intervals with stalling TCP sequence numbers that are longer than 3 s.

Table 1
Real world measurements for the gray zone scenario (Fig. 4) of successful Ping and MP3 packet delivery ratios

Success ratio	OLSR	LUNAR	AODV-UU	AODV-UU + SNR
Ping	89.0%	96.5%	91.9%	99.1%
MP3	91.9%	96.8%	97.9%	99.7%

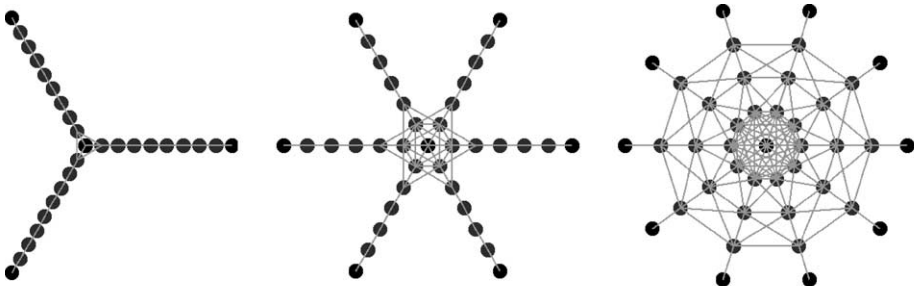

Fig. 5. "Beam star" network topologies for various numbers of beams and beam lengths (3 × 9, 6 × 6, 10 × 4) shown with potential connectivity. (The figures have different scales: the inter-node distance along a beam is identical for all topologies.)

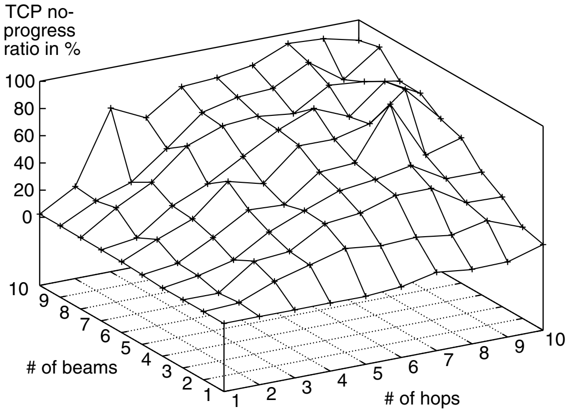

Fig. 6. "Worst accumulated TCP no-progress time" as percentage of total run time for various number of beams and hops (beam star scenarios, one FTP session per beam, no mobility, AODV-UU with link-layer feedback, local repair and gratuitous RREQ).

These no-progress gaps are added up and the worst number from all beams is retained i.e., the no-progress ratio is defined as

$$r_{\text{noprogress}} = \frac{\max_i(\sum\{t_{\text{interseggap}}(beam_i) | t_{\text{interseggap}} > 3\,\text{s}\})}{t_{\text{testrunlength}}}.$$

From the data that underlies Fig. 6, we have observed that already with two beams and three hops (totaling seven nodes) there is up to 30% accumulated no-progress time, which includes a 17 s stall (not visible from the graph). That is, at least one user in the system fetching WEB pages or downloading emails had no useful TCP service for approximately one third of the time! Extended no-progress times of this magnitude are not tolerable, even for web browsing. In this case, users would probably not adopt ad hoc technology. The reason why such bad performance figures do not show up in simulation studies is mainly due to the reporting of *average* throughput numbers where unfairness due to TCP capture is not detectable. The other reason of course is that almost all MANET performance comparisons are carried out with CBR traffic rather than TCP.

Referring to Fig. 6, which was obtained by using the most optimal settings for AODV (no node mobility, availability of link-layer feedback), we suggest that two hops and five beams (11 nodes) define a safe area for the combination of plain TCP, AODV and 802.11: acceptable network performance beyond this range can only be achieved if the horizon can be extended. The beam star topologies form a synthetic benchmark for advances at the (combined) level of TCP, ad hoc routing and media access control. At the same time, the ad hoc horizon—as defined over these beam star scenarios—defines a target area on which routing protocols should focus on.

4. Lessons from research outside the IETF MANET area

While we see the APE testbed and the AODV-UU implementation as our contributions to the IETF MANET working group effort, we have also investigated research questions outside the

MANET borders. One such activity relates to simplifying the routing logic and sidestepping IP interfacing problems. In this section we give a brief introduction to the LUNAR routing protocol and highlight some of our results.

4.1. LUNAR—targeting simplicity

LUNAR, the "Lightweight Underlay Network Ad hoc Routing" protocol [22,21], is the result of an implementation experiment where the ad hoc routing logic was reduced to a minimum. Instead of having routing protocol logic that actively maintains and repairs existing routing paths, LUNAR always establishes paths from scratch. That is, as long as there are packets to be sent for a source–destination pair, LUNAR floods the network every 3 s for discovering a route. During the 3 s interval, all data packets are shipped over the same path and no attempts are made to discover lost packets or to monitor link changes. Another departure from MANET protocols, briefly described below, is the underlaying of IP such that multihop paths are masked away. LUNAR presents to IP a virtual subnet which leads to a clean separation of IP level routing and the wireless multihop forwarding.

To our surprise, this simple approach performed better than OLSR and the original AODV-UU implementation in real-world experiments [22] (see also Table 1). By keeping the rediscovery interval small enough, we could match the reactiveness of HELLO-based protocols which—in absence of link-layer feedback—typically need two HELLO rounds (of 1 s each) to decide that a neighbor is not reachable anymore.

Because of the repeated full discovery of each routing path, it was decided to limit LUNAR's range to 3 hops, i.e., to target LUNAR specifically to small networks of approximately a dozen nodes. When LUNAR was ported to the ns-2 simulator [16], we found that LUNAR works well beyond this self-imposed three hop limit. In a constant density scenario, where the simulation area grows with the number of nodes, and a random waypoint mobility model, LUNAR(maxhop = 15) works comparably to MANET protocols in settings with 100 nodes and paths up to 15 hops. Under CBR traffic LUNAR outperforms OLSR, TORA and DSR over most network sizes although it does not match AODV's performance (see left graph in Fig. 7).

Regarding the routing protocol overhead (and the doubt, that the forced rediscovery of routing paths does not scale well), we found that LUNAR(15) generates routing traffic comparable to DSR (right graph of Fig. 7: The normalized routing load is the amount of routing bytes per each delivered data byte). At 40 nodes, LUNAR(15) still has the same routing overhead as AODV with

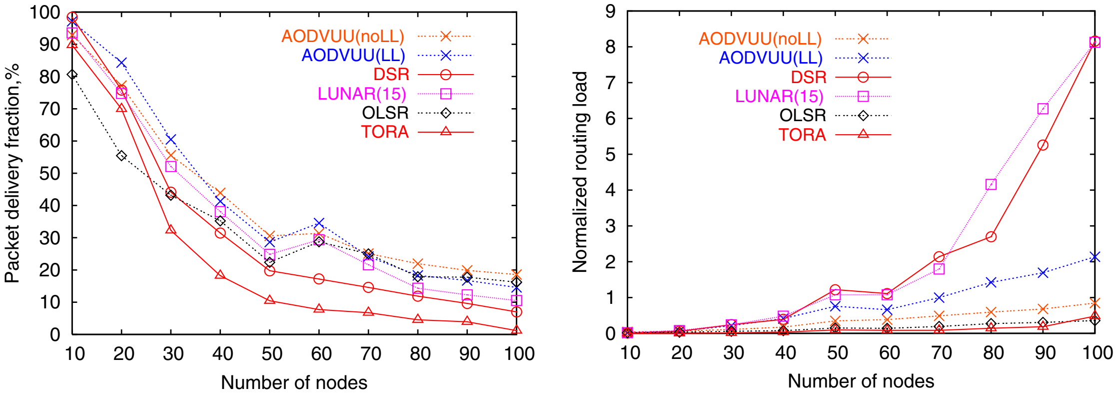

Fig. 7. Comparison of CBR delivery fraction (left figure) and routing overhead (right figure) in a constant density/random waypoint scenario for LUNAR and major MANET protocols.

link-layer feedback enabled (AODV(LL)), but quickly becomes rather inefficient at higher node numbers. Still, this is beyond the network size for which LUNAR was designed for. As a side remark we note that so far we could not fully explain the surprisingly high amount of overhead for DSR. A partial explanation relates to the length of paths (and thus increased chance of link breaks), where our trace data shows that the longest routes established by DSR are twice as long as for AODV and LUNAR. Another potentially related element could be DSR's small amount of jittering for forwarding route requests, which aggravates congestions inside the network. This performance aspect of DSR needs further investigations, including the examination of its ns-2 implementation.

A useful architectural element of LUNAR is the underlaying of IP (see Fig. 8). LUNAR creates a virtual (Ethernet) subnet where all hosts attached to this subnet seem to be reachable by one hop. All hosts are addressed through classic Ethernet numbers wherefore LUNAR can be implemented as a device driver. LUNAR "intercepts" address resolution (ARP) requests that the IP stack emits for discovering a destination's Ethernet number. The ARP request then triggers a discovery wave. After discovery, the established forwarding path is associated with a virtual Ethernet number which is returned to the IP stack as an ARP reply. This creates a subnet illusion to IP (Fig. 8) which makes integration with the IP stack much simpler: only a single static subnet entry has to be added to the IP routing table.

Inside the underlay we keep independent forwarding paths for each active node pair. This has the architectural advantage that there can be no false sharing of routing state. For example, in an ad hoc network with several gateways to the fixed Internet, each LUNAR node will be able to pick and stick to its preferred default gateway while in other routing protocols like AODV it can happen that downstream nodes "repair" a route and switch between gateways, which creates problems with NAT and MobileIP state residing in the gateway nodes.

The simplicity of LUNAR also helped us to use it in environments as diverse as Infrared-based embedded systems, Bluetooth-enabled devices and both the Linux and Windows XP/2000 Operating Systems for IEEE 802.11. LUNAR was ported to the LEGO mindstorm micro-controllers with 32K of RAM. We managed to fit the LegOS, a TCP/IP stack, a tiny web server as well as the LUNAR protocol into this limited environment.

4.2. Cross-layer optimization between link and transport layer

The significance of LUNAR is that it permits more experimentation to be done outside the classic proactive and reactive protocols of MANET, and inside protocol stacks.

One observation on LUNAR is that it blurs the reactive/proactive categorization. LUNAR is reactive as it does all routing on demand, but has at the same time a proactive attitude as it re-samples the full network topology at fixed intervals.

Another aspect of LUNAR is that it is well positioned for extending the ad hoc horizon through cross-layer optimization. More specifically,

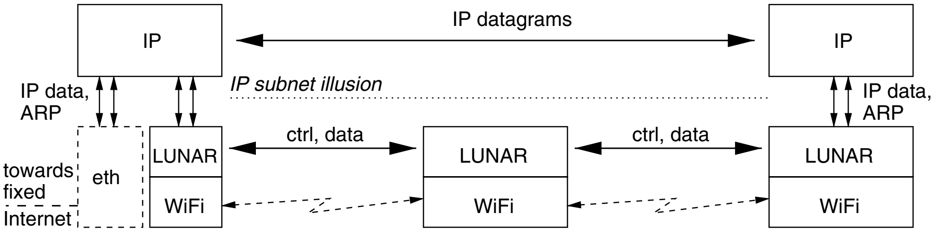

Fig. 8. The LUNAR ad hoc routing underlay for IP.

we exploit the fact that LUNAR can mediate between link layer technologies and the transport layer. Because LUNAR is an underlay network (instead of a IP level add-on routing protocol), it sees all TCP traffic and can react in smart ways based on an internal model of TCP's behavior by scheduling TCP messages. At the same time, LUNAR knows the behavior of the link layer; It can sense the network condition and has detailed knowledge about routes, like their length or congestion indications that are collected along the LUNAR forwarding path. First results with an enhanced version of LUNAR indicate that we will be able to double the ad hoc horizon for IEEE 802.11 based IP networks without changing neither the widely deployed MAC nor the ubiquitous TCP layer.

5. Challenges to ad hoc networking

During the course of our experimental MANET research, we have identified two key challenges to the future success of ad hoc networking. Namely, how to make MANET routing keep up with developments in link layer technology and how to target MANET research on realistically achievable goals.

5.1. Dependence upon link layer technology

The clear separation between link, network and transport layer is difficult to maintain in a wireless environment. The IETF MANET working group itself points out that routing in mobile ad hoc networks is different from traditional fixed infrastructure Internet routing. Handling cross-layer interactions, as is frequently proposed these days, is hard and there is a danger to build solutions for specific combinations, e.g., "802.11 + ad hoc routing". New combinations like "Bluetooth + ad hoc routing" would then need an independent reevaluation, considerably slowing down the progress of the field.

There is a risk that lower layer progress will outpace the slow progress of MANET routing. Lower layer progress is driven by technologies for infrastructure settings. This may create more link layer variants for MANETs to cope with which—in the worst case—work poorly in ad hoc mode. We believe that we should provide even limited, but robust MANET solutions in order to benefit the users and to have an impact on the lower layer technology development.

5.2. Limiting the problem space

One unrealistic goal of the MANET work is to aim at networks of hundreds of mobile nodes. Many simulations use such scenarios and see a high node number as a goal of its own. Recently a special IRTF interest group (ANS [3]) was formed to look into scaling issues of ad hoc networks. We have doubts that results from this effort will be helpful for the standard user in his daily wireless Internet usage. Still, we acknowledge that large scale ad hoc networking makes sense for sensor networks with special purpose transport protocols and reduced mobility, etc.

More artificial barriers are raised that hinder the development of pragmatic solutions. Security, QoS and power awareness are the well known arguments in this context. The reaction to the growing complexity should be to concentrate on workable solutions and convincing cases. It is better to have a working solution for a majority of wireless interfaces that might be limited in scaling and that does not address all possible concerns, but which gives an added value right now. As Metcalfe's law tells us it is a matter of increasing the number of users in order to stimulate the application scenarios and the attractiveness of the technology.

The case today is that everybody uses IEEE 802.11 hardware, but virtually nobody runs it in multihop ad hoc mode. Kernels are not shipped with ad hoc routing logic and software configuration is complex and arcane. One crucial element in raising the penetration rate of ad hoc functionality is that configuration management must become really simple. The real value of ad hoc networking will only be evident when every device is ad hoc enabled, which is why resolving penetration hurdles should be the prime target for the ad hoc community.

6. Outlook

In this paper we discussed lessons learned from our experimental research in ad hoc networking. We make the point that some targets and directions of MANET research should be debunked and revisited.

One concern is that research is expanding into so many new areas and technologies that a simple solution for the common case is effectively ignored or will come too late. Another—more methodological—observation is that ad hoc network research is too narrowly conducted through the means of simulation. We have found that many assumptions in simulation are not valid in the real world. A more focused approach guided by "real" MANET problems and available solutions should be adopted.

An additional guideline should be to study feasible problems, i.e., to limit the problem space and to optimize for constraint but useful settings. The ad hoc horizon of three hops and 10–20 nodes shows a suitable first target area. Improving performance inside this range and expanding its limits would be the next logical step before pursuing artificial and extreme problems.

Finally, we point out that solving routing in MANETs is only a first step towards deployable ad hoc networks. Integrated self-configuration solutions that cover everything from address assignment to automated gatewaying with the fixed Internet are needed to push ad hoc functionality into main stream operating systems and into our daily network experience.

Acknowledgements

We wish to thank Richard Gold for his careful editing help and Evgeny Osipov for the port of LUNAR to the ns-2 simulator.

References

[1] The APE testbed. http://apetestbed.sourceforge.net, 2001.
[2] G. Anastasi, E. Borgia, M. Conti, E. Gregori, Wi-Fi in ad hoc mode: a measurement study, in: IEEE International Conference on Pervasive Computing and Communications, PerCom'04, March 2004.
[3] Ad hoc Network Systems Research (IRTF group). https://www1.ietf.org/mailman/listinfo/ans-research.
[4] AODV-UU implementation. http://www.docs.uu.se/scanet/aodv, 2004.
[5] T. Clausen, P. Jacquet, Optimized Link State Routing Protocol (OLSR), IETF RFC 3626, Octotober 2003.
[6] M. Gerla, K. Tang, R. Bagrodia, TCP performance in wireless multi-hop networks, in: Proceedings of IEEE WMCSA'99, New Orleans, LA, February 1999.
[7] G. Holland, N.H. Vaidya, Analysis of TCP performance over mobile ad hoc networks, in: Proceedings of IEEE/ACM MOBICOM '99, August 1999, pp. 219–230.
[8] D. Johnson, D. Maltz, Y.-C. Hu, The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR), IETF Internet draft, draft-ietf-manet-dsr-09.txt, April 2003.
[9] H. Lundgren, D. Lundberg, J. Nielsen, E. Nordström, C. Tschudin, A large-scale testbed for reproducible ad hoc protocol evaluations, in: 3rd Annual IEEE Wireless Communications and Networking Conference, WCNC 2002, March 2002, pp. 412–418.
[10] H. Lundgren, E. Nordström, C. Tschudin, Coping with communication gray zones in IEEE 802.11b based ad hoc networks, in: 5th ACM International Workshop on Wireless Mobile Multimedia, WoWMoM 2002, September 2002, pp. 49–55.
[11] H. Lundgren, E. Nordström, C. Tschudin, The gray zone problem in IEEE 802.11b based ad hoc networks, in: ACM SIGMOBILE Mobile Computing and Communications Review, vol. 6, July 2002, pp. 104–105.
[12] D. Maltz, J. Broch, D. Johnson, Experiences designing and building a multi-hop wireless ad hoc network testbed, Carnegie Mellon University Technical Report CMU-CS-99-116, March 1999.
[13] D. Maltz, J. Broch, D. Johnson, Quantitative lessons from a full-scale multi-hop wireless ad hoc network testbed, in: 2nd Annual IEEE Wireless Communications and Networking Conference, WCNC 2000, September 2000.
[14] The Official IETF working group MANET webpage. http://www.ietf.org/html.charters/manet-charter.html.
[15] R. Ogier, F. Templin, M. Lewis, Topology Dissemination Based on Reverse-Path Forwarding (TBRPF), IETF RFC 3684, February 2004.
[16] E. Osipov, C. Tschudin, LUNAR vs. MANET—performance comparison of routing protocols for wireless ad-hoc networks, in: Proceedings of the 4th Scandinavian Workshop on Wireless Ad-hoc Networks ADHOC'04, Johannesberg, Sweden, May 2004.
[17] C. Perkins, E. Belding-Royer, S. Das, Ad hoc On-demand Distance Vector (AODV) Routing, IETF RFC 3561, July 2003.
[18] R. Ramanathan, R. Hain, An ad hoc wireless testbed for scalable, adaptive QoS support, in: 2nd Annual IEEE

Wireless Communications and Networking Conference, WCNC 2000, September 2000.

[19] M. Takai, J. Martin, R. Bagrodia, Effects of wireless physical layer modeling in mobile ad hoc networks, in: ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2001, October 2001.

[20] C.-K. Toh, M. Delwar, D. Allen, Evaluating the communication performance of an ad hoc wireless network, IEEE Transactions on Wireless Communications 1 (3) (2002) 402–414.

[21] C. Tschudin, Lightweight Underlay Network Ad Hoc Routing (LUNAR) Protocol, IETF Internet draft, draft-tschudin-manet-lunar-00.txt, March 2004.

[22] C. Tschudin, R. Gold, O. Rensfelt, O. Wibling, LUNAR—a lightweight underlay network ad-hoc routing protocol and implementation, in: Next Generation Teletraffic and Wired/Wireless Advanced Networking, NEW2AN'04, St. Petersburg, Russia, February 2004, pp. 38–43.

[23] C. Tschudin, E. Osipov, Estimating the ad hoc horizon for TCP over IEEE 802.11 networks, in: Proceedings of the 3rd Annual Mediterranean Ad Hoc Networking Workshop, Med-Hoc-Net'04, Bodrum, Turkey, June 2004, pp. 255–262.

Christian Tschudin is a Professor in Computer Science at the University of Basel, Switzerland. He holds a Diploma degree in Mathematics (1986, Basel) and a Doctor of Science in Informatics (1993, Geneva). Before joining the University of Basel, he has been doing research in computer networks at the University of Geneva, the University of Zurich and at the International Computer Science Institute (ICSI) in Berkeley. From 1998 to 2002 he was an associate professor at the Uppsala University, Sweden. His research interests are in mobile code and active networks, security, routing protocols and performance evaluation for wireless networks, network architectures and bio-inspired networks.

Per Gunningberg is Professor in Computer Communication at Uppsala University, Sweden. He holds Doctor of Technology from Uppsala University from 1983, and Master of Science in Computer Science, from University of California Los Angeles, from 1981. He is a member of IEEE and ACM. Since 1985 Gunningberg's research activities have been within the Computer Communication field, with emphasis on performance and higher layer protocol issues. Current activities include ad hoc networking, network support for overlays, wireless performance measurements and dynamic proxies for supporting wireless thin clients.

Henrik Lundgren received his BSc degree in Computer Technology with specialization in telecommunication from Kalmar University, Sweden in 1997. Between 1996 and 1998 he was working as an R&D engineer at Ellemtel Utvecklings AB, a joint Ericsson/Telia research company in Stockholm, Sweden. Since 1999 he is a PhD student in computer communication at Uppsala University, Sweden. His research includes mobile and wireless communication systems. He currently focuses on experimental evaluation of mobile ad hoc networks.

Erik Nordström obtained his MS degree in 2002 from Uppsala University, Sweden, where currently he is a PhD student at the Department of Information Technology. His research in ad hoc networking focuses on self-configuration, Internet integration and real-world testbeds. He is the main author and maintainer of the Uppsala implementation of the Ad hoc On-demand Distance Vector routing protocol, AODV-UU, and has contributed major parts of the APE testbed. Erik has also been active in the IETF MANET working group.