Embedding MANETs in the Real World

Christian Tschudin[1], Henrik Lundgren[2], and Erik Nordström[2]

[1] Computer Science Department, University of Basel
Bernoullistrasse 16, CH – 4056 Basel, Switzerland
[2] IT Department, Uppsala University
Box 337, SE – 751 05 Uppsala, Sweden

Abstract. We believe that research on ad hoc networking so far has focused on too broad a range of scenarios and optimization issues and we argue that limiting the settings to bare bones services and small networks will allow for quick MANET deployment. Especially if these solutions interoperate easily with the Internet and offer the possibility to grow the MANET networks from there. Starting at this end will also help us to identify and possibly enlarge the "ad hoc horizon", beyond which ad hoc networking is no longer useful. In this paper we report on our experiences from five years of real world ad hoc network experimentation. Our approach is pragmatic and includes several implementations and a testbed software package. We highlight issues that we have encountered and identify crucial problems and potential dangers that deserves special attention in order to make ad hoc networking usable.

1 Introduction

Since the initial work on an IETF charter for mobile ad hoc networks started around 1995/1996, the MANET working group [13] has been working towards a proposed standard for IP level routing in ad hoc networks. Recently, the charter was narrowed as a result of the increasingly broad scope touching on problems like scalability, power awareness and security. Most of these research issues will now be addressed in the context of the Internet *Research* Task Force [8] . The MANET charter as of 2003 will propose four protocols for experimental RFC status, namely AODV, DSR, OLSR and TBRPF. This would indicate a new level of maturity of multi-hop ad hoc networking. However, we believe that the "race for multi-hop" is far from over and point to the fact that MANET solutions, if they exist, are not embedded in our daily network experience.

In this paper we argue for going simple and real. We believe that solving simple scenarios, i.e. small 10-20 node networks with one to three hops together with Internet connectivity would spur innovation and make ad hoc networking go from being a research-only phenomena to being utilized in real world settings. Our suggestion is to focus on implementation, experimentation and integration.

1.1 Implementation

Unfortunately, the deployment of MANET functionality has been slow. The first and main reason is the lack of quality real world implementations. When we started building the Ad hoc Protocol Evaluation Testbed in the year 2000 [9], we had a hard time finding (a) any implementations, and (b) implementations that worked properly. Out of three 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. This contrasts sharply with the number of related Internet Drafts and revisions, which seems to indicate that the field is making good progress.

AODV has seen most maturity in this respect whilst other major MANET protocols either lack updated implementations, have intellectual property issues or suffer from the reluctance among researchers to share work and promote open development. In some cases there are no implementations available at all, as is the case with TBRPF[3]. But even AODV is not a product yet: Self-configuration and seamless interfacing with the fixed Internet is still unsolved.

1.2 Experimentation

Experimentation is the bridge between paper specification and real world deployment. This is where the quirks that the draft authors missed and the invalid assumptions they made are discovered. However, although MANET research has been ongoing for some time, little valuable "hands-on" experience is available. It is alarming that protocols are mainly developed and updated in response to findings from simulations which are run using optimized settings that are not available in real life. For example, in simulations, routing protocols typically implement layer interactions (e.g. feeding link layer failures up the stack) that are not available in real world settings.

Another problem is that simulations are almost never validated. It is very tempting to do research through simulations because they are cheap when compared to setting up real experiments and are nicely repeatable. A first aspect of this is the trust which is put into the protocol logic that comes with the simulation packages: The protocol implementations are seldomly assessed in an independent way. The second aspect is that there is rarely an effort to calibrate or directly compare the simulation with findings from actual measurements.

1.3 Integration

Integration is the next step towards deployment. How do we make sure that the functionality we need in operating systems are exposed through standard APIs? How do we make sure that new functionality is integrated?

A consequence of the narrow implementation basis is that the integration of the protocol logic into actual operating systems is similarly scanty. The manipulation of IP routing tables is a good place to point to: "Default IP routes" make little sense in a MANET and the notion of a subnet to which an interface is attached becomes similarly void. This requires "creative solutions" and kernel

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

workarounds for mapping transient routing paths to IP routing table entries. (At this point one can literally feel the conceptual stress between MANET routing and classical IP routing concepts.)

The other point that integration demonstrates is that routing is more than just an algorithm. Routing is inherently linked to operational matters (e.g., address assignment) and therefore configuration issues. Unfortunately, such questions, although crucial for the end user's comfort, are beyond routing protocol specifications. It is fair to say that we do not have good answers ready for extending Internet connectivity into the MANET realm.

1.4 Half Time for the Multihop Game

Although respectable progress has been achieved we think that research in MANET multihop routing has only done half of the job. Clearly there are no final results. Simulations have not been conclusive on clearing the field of protocol contenders, and hard reports from the user front are missing because there is virtually no deployment of wireless multihop functionality.

So the second part of the MANET job will be decisive for whether wireless multihop routing manages to find a user basis. It could well be that infrastructure solutions (e.g., managed mode IEEE 802.11), which are actually responsible for the deployment of the radio hardware, will simply occupy the market space and that real MANET scenarios are too insignificant in order to create a shift.

In this paper we discuss our pragmatic approach to MANET research and how we believe that it can help turning the MANET concepts into a deployable technology. In section 2 we start with presenting the experimental research we have done in our group during the past years. In section 3, we revisit some of the underlaying ideas and premises behind MANET thinking and present work that we have done outside MANET borders. Section 4 continues with prospects for the future before we conclude this paper in section 5.

2 Reality Checks

Most wireless testbeds [12, 11] [16] [14] could be called "measurement setups"; They targeted a single protocol and do not represent a software product which can be used by others for repeating experiments (see table 1 below). By developing the Ad hoc Protocol Evaluation testbed (APE) we aimed at creating an investigation tool that ultimately would be used to compare all MANET implementations at a 1-to-1 scale.

To this end we adopted a modular architecture allowing an experimentalist to plug in *different* protocols and traffic generators. Currently, the APE testbed is prepared for installation of most available protocol implementations running under Linux. Because of "scaling concerns" when setting up experiments with ten or more human participants, APE is designed for easy installation, deployment and handling. Installation is as simple as "downloading and extracting" a single file. It contains a self-extracting stripped-down Linux system that can

	CMU	*GA tech*	*BBN*	*APE*
Real-world	yes	yes	yes	yes
W-less emul.	yes	yes	yes	yes
Hardware	laptops + GPS	laptops	Nokia Rooftop	laptops
OS	FreeBSD	Linux	DAWN/FreeBSD	Linux
Routing Prot.	DSR	ABR	NLS	**Independent**
Visualization	yes	no	no	yes
Analysis tools	yes	yes	yes	yes
Software Distribution	no	no	no	**Yes**
Reproduc.	no	no	no	**Yes**

Table 1. Comparison of various ad hoc network testbeds.

be easily installed and booted on machines running Windows or Linux, or run using a "LiveCD" on a PC compatible machine. During the testruns the participants only need to follow on-screen movement instructions and afterwards the collected measurement data is uploaded to a collect-node, all with minimal user intervention.

To address repeatability we use strict choreography instructions to the participants in conjunction with a new mobility metric, derived from measured signal quality. This gives us a way to compare the "mobility fingerprint" of each testrun and to assess whether testruns were performed in a sufficiently similar manner, independent of the routing protocol used. We have run experiments where we scaled up to a 37 node setting and successfully used the new mobility metric to assess the reproducibility.

2.1 Implementation Issues

As the first thing a routing protocol is usually exposed to is a simulation environment, the implementations tend to be more or less tailored to the particular simulation environment at hand and neglect the interactions with the operating system and its protocol stack. However, sooner or later the protocol will have to be implemented and integrated within an OS.

Although OS kernel interactions are quite complex and make the implementation time-consuming, the most striking experience during our AODV-UU [2] development was that some protocol design assumptions might not be feasible to implement in a real system. For example, the AODV specification suggests to primarily use link-layer feedback for monitoring the link status of active routes and only if that is not possible to rely on periodic HELLO beacons. In simulation, link-layer feedback is easy and therefore commonly used. In reality, there is to this date still no way for upper layers to acquire such information from IEEE 802.11 interfaces. This observation impacts those protocol simulations studies which used link-layer feedback. Repeating them using the HELLO beacons would give a better view of how the protocol would behave and perform in reality.

2.2 Simulation Model Accuracy

Exposing implementations to the real world will not only highlight practical limitations, but will also help reveal protocol behavior that might not show up in simulations due to inaccurate simulation models. We became subject to this statement when we observed spurious performance degradations during experiments with our AODV-UU implementation. Already in a basic scenario with one mobile node we experienced "communication gray zones" [10]. In these gray zones of several meters the routing protocol reports a valid route to the destination node but virtually no data packets would pass through. This effect does not show up when replicating the setting inside the ns-2 simulator (using the same AODV-UU code as in the OS setting) and thus it was clear that this is a discrepancy between the simulation environment and the real-world because data transmission works well in all other areas.

Further real-world experiments revealed that these gray zones form because of the difference in transmission range for specific packet types. In IEEE 802.11b broadcast packets are sent at 2Mbit/s and unicast can be sent at speeds up to 11Mbit/s. The transmission output power is independent of transmission speed and thus lower transmission speed means more energy per bit, which translates to longer transmission range.

The consequence is that at the transmission range border between two nodes the data packets will not reach the destination while the broadcasted HELLO beacons (used for neighbor sensing) will be delivered and make the routing daemon believe that the other node is reachable, thus it will *not* try to discover a new route. This gray zone effect did not show up in the simulations simply because ns-2 uses a flat 2Mbit/s model.

Based on our implementation we then started to test work-arounds. In [10] we showed that artificially limiting the range of the HELLO packets by filtering on the SNR (Signal to Noise Ratio) value eliminates the gray zone effect. In fact, setting the acceptance SNR level such that HELLO packets have slightly *shorter* transmission range than the data packets will force AODV to pick a more robust link earlier, which decreases the inherent packet loss during route switching.

Accuracy of the physical layer model is equally important. This model includes e.g., SNR calculation, signal reception, fading and path loss. However, the study in [15] is revealing: 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 such that the corresponding choice could even change the relative ranking of the compared protocols.

In conclusion, it is important to realize that the simulation models are just models. Although simulation studies are very useful in the protocol development process they can never totally replace real-world testing. In fact, we need more real-world experience earlier in the process.

2.3 The Ad Hoc Horizon

There are many questions still open when it comes to real world experiments. For example, little has been reported in the research literature about problems of TCP over wireless multi-hop *implementations*, although there are some discomforting results from simulations [4, 6]: TCP throughput is inversely proportional to the hop distance for a RTS/CTS style MAC layer (i.e., 25% after 4 hops) and even worse if the TCP window is not artificially limited. Another way of grasping this is to say that there exists an "ad hoc horizon" beyond which it is not worth trying to offer MANET services. Our conservative estimate is that using IEEE 802.11 technology this horizon is as narrow as 3 hops.

How can such a horizon form? Reachability information, on which routing protocols base their decisions, loses its value the further it travels in a mobile environment. Also, a loss rate factor applies for each hop that a data packet has to traverse: Crossing 4 hops with an average delivery ratio of 95% brings us down to almost 80% (0.95^4) end-to-end, probably making TCP based applications unusable. Finally, the longer the path is the more places there are where TCP data packets and backward traveling ACK packets can collide. The overall effect, even if the average drop rate is low, is that TCP becomes an unusable transport protocol because it is tricked into seeing a permanently congested network.

In [5] a measurement study is presented for determining the "scaling law" of (static) wireless networks. It demonstrates a dramatic decrease in (UDP-) throughput as the number of nodes goes from 2 to 10 and compares this to theoretical figures. It seems that the IEEE 802.11 standard is far from using the available capacity in an efficient way. Here, we present a simulation study that shows similarly severe throughput degradation when going from 10 to 100 nodes and when adding mobility and ad hoc routing.

Using ns-2 (v2.26) we compared TORA, DSR and two AODV implementations whereof one is bundled with ns-2 (AODV) and the other is our own AODV-UU implementation. The goal of the simulation is to compare network scalability with respect to the number of nodes ranging from 10 nodes to 100 nodes while keeping the density constant. In our simulation this is done by increasing the network area size accordingly, using a rectangular area with a x:y ratio of 1:4, to promote increasing path lengths as the network grows[4].

[4] Node placement and movement is random, using the random waypoint mobility model with a max speed of 20 m/s and pause time of 0.2 s. Traffic is generated using the CMU traffic generation tools. We use CBR traffic with a packet size of 512B, 0.5 connections per node (which translates to $0.5 \times \#nodes$ connections max), with a rate of 10 pkt/s. We experimentally found a node density of $5e^{-5}$, by aiming for a delivery ratio slightly under 100% for the 10 node setting. This translates to an initial area size of $223.6 \times 894.4m^2$. For the 100 nodes setting, the network size has increased to $707.1 \times 2828.4m^2$. The protocols are run with default settings for TORA and DSR that corresponds to ns-2 version 2.26. The ns-2 AODV implementation also had default settings with the exception that HELLO messages were used instead of link layer feedback. AODV-UU was run using default values for version 0.7.1, but also with HELLO messages.

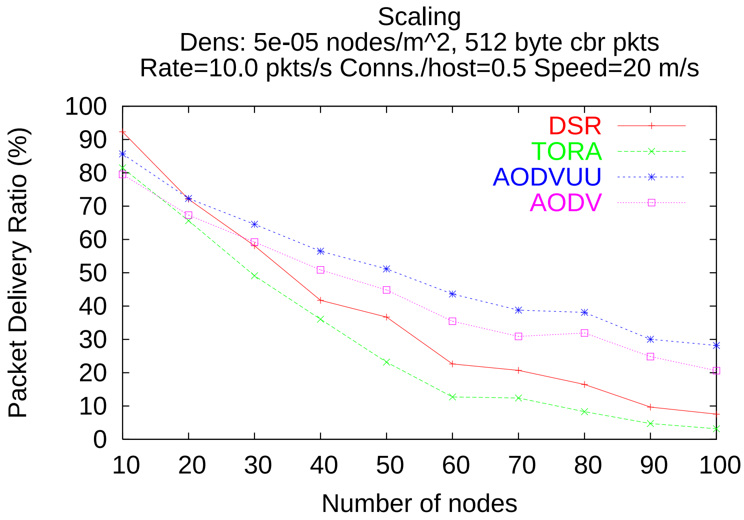

Fig. 1. The effect of scaling where node density is kept fixed.

Figure 1 shows the performance results averaged over 10 simulation runs. Although working with optimistic assumptions (simple UDP traffic, uniformly spread sources and sinks) the delivery ratio is in the best case inversely proportional to the maximal path length: Remember that since density is fixed, increasing the number of nodes means increasing the area size and producing longer routes. An examination of the trace files indicate that the maximum path length increases from 5 to 15 hops when growing the network from 10 to 100 nodes. It is obvious that the performance of *all* protocols decrease rather ungracefully, although at different rates. Our conclusion is that independently of which of the three routing algorithms we chose we will hit the ad hoc horizon well below 100 nodes, that is 15 hops. Adding TCP to the picture we deem it reasonable to suspect the horizon to be 10 nodes and 3 hops – a hypothesis to be further studied with experiments and simulations.

3 Out of the MANET Box Thinking

While we see the APE testbed and the AODV-UU implementations as our contributions to the MANET effort, we also wanted to investigate opportunities outside these borders. One direction uses an active networking approach where protocol logic resides in control packets instead of the nodes (hence one obtains a multi-protocol environment). The other direction is on sidestepping some of the IP interfacing problems and on simplifications of the routing logic. In this section we give a gentle introduction to these approaches and highlight some of our results.

3.1 Active Ad Hoc Routing

The difficulty to agree on "the" set of scenarios and border conditions for a MANET protocol relates to the assumption that a single ad hoc routing protocol could satisfy all (or most) requirements. Our approach relies on mobile code for ad hoc routing implementations. Such an *active networking* approach permits to deploy the routing logic at run time either in the scope of a network or within a single session. This extreme customization also comes in handy for the debugging and testing of protocol variants, as there is no need to visit nodes for installing a new software version.

We gained several insights from this "implementation experiment", which we have documented in [18]. First, routing logic does not need to be executed very fast. The transmission time in the range of milliseconds and the low accuracy of HELLO beacons (multiple seconds) leaves enough room to start up complete command shells for processing routing logic. We used simple Bash scripts which could forward themselves through the network for implementing the routing protocol. Of course, handling the data packets in the same manner would not work well. To this end we created a separate forwarding layer below IP where data packets are forwarded in the kernel without IP processing. This separation of pure data forwarding (without any intervening routing or packet salvation logic) and active routing control packets worked well enough to demonstrate the streaming of MP3 audio packets in a wireless ad hoc network!

Second, we used this setup to experiment with proactive as well as reactive routing schemes i.e., both routing styles can be expressed equally well with active packets. We "scripted" a proactive protocol loosely inspired by Cellular IP [3]: Three different populations of active packets were responsible for a) periodically gathering neighborhood information, b) for building a common delivery tree and c) for establishing delivery paths along this tree. Later on, we also scripted a reactive protocol called AORTA with the classical flood and reverse path pattern. This was a mere matter of using another script, under the full control of the initiating node. Each routing path then operates independently and could use its own variant of a reactive protocol, if desirable.

Third, when looking back we think that current ad hoc routing logic is not complex in its core: Most difficulties stem from optimization that one applies on top of the basic functionality. One example is the sharing of forwarding state where one has to decide which control packet is allowed to modify state created by other (end) nodes' packets. Another example is the salvation of data packets which considerably complicates the simple forwarding mechanism by adding timeouts and new dynamic elements.

3.2 LUNAR – Simple Routing Below IP

Even plain ad hoc routing can be simplified, as we tried out with LUNAR, the "Lightweight Underlay Network Ad hoc Routing protocol" [17]. The main question we stumbled over was: Once we have set up a routing path, why should we bother to maintain it? Protocols using HELLO messages will typically need two HELLO rounds to decide that a link is broken before the protocol starts reacting. So we decided to do a plain discovery every 3 seconds and to drop HELLO messages altogether.

To our surprise this simple approach performed better than OLSR and the original AODV implementation (before applying the "grey zone patch" described above). Admittedly, LUNAR's "periodic-full-rediscovery" will not scale well. However, the crucial question is whether more complex approaches will ever be operated in large networks, as inherent scaling problems as reported in figure 1 degrade their usefulness.

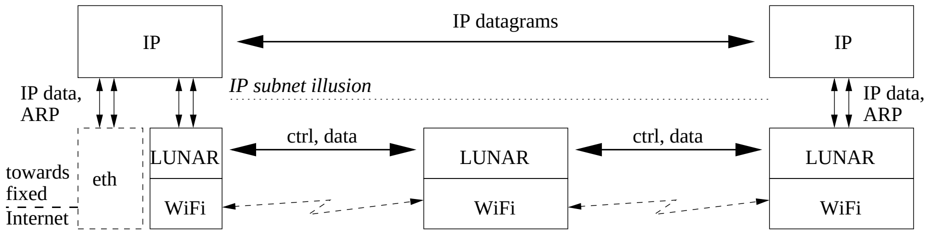

Fig. 2. The LUNAR ad hoc routing underlay for IP.

Another useful simplification of LUNAR is the underlaying of IP (see figure 2). LUNAR works by trapping ARP requests which are then forwarded through the wireless network. Once the destination is found a delivery path is set up end-to-end and mapped to a virtual ethernet number. This creates a subnet illusion to IP which makes integration with the IP stack much simpler, as only a single entry has to be added to the IP routing table.

A fun project turned out to be the demonstration of leveraging LUNAR's simplicity. We rewrote LUNAR for embedded devices and ported it to the LEGO mindstorm microcontrollers with 32K of RAM, wherein we managed to fit the LegOS, a TCP/IP stack, a tiny web server as well as the LUNAR protocol and demonstrated it at Mobicom'2002. All code is GPL and available at http://www.docs.uu.se/selnet/lunar/.

4 Ad Hoc Prospects

4.1 Danger of Point Solutions and too Slow Progress

The basic assumption of research in ad hoc routing is that conventional (IP) routing algorithms behave poorly in wireless environments. This means that the layering approach is flawed because one cannot hide the link layer characteristics from the networking layer: the network layer has to choose different routing strategies depending on the fixed or wireless link layers used. Ironically, this observation will fall back on the ad hoc research community as IEEE 802.11a, b, g and Bluetooth have already sufficiently many differences that conclusions drawn for the "802.11+ad hoc" combination will be invalid for the "Bluetooth+ad hoc" pair.

One possible (worst case) outcome is that new developments and progress at the link and physical layer will outpace the slow progress at the routing level, first by creating more variants which the routing layer has to cope with, second by producing technology which performs well in an infrastructure setting but badly in ad hoc mode, and finally by raising the level of expectations so that the "imperfect" ad hoc solutions never will have no chance to catch up. So users should have a chance to experience the advantages of ad hoc functionality early on: The quicker even limited (but robust) MANET solutions become available, the better.

4.2 Shooting for the Moon instead of Mars

One impressive goal of MANET work is to aim at networks of 100s and more nodes. Many simulations work with such scenarios and see a high node number as a goal of its own. Recently a special interest group (ANS [1]) was formed to look into scaling issues of ad hoc networks. The problem with such large numbers is that they represent a fake goal. First, such a point of operation is rather arbitrary on the spectrum from a single-site spontaneous network with a handful of mobiles to the Terminodes vision [7] that aims at operating at the global level. Second, such a number is again dependent on link layer technology as well as the scenarios. The HELLO beacons of 100 mobiles in the cocktail party room of a networking conference would kill AODV. Are the protocols at fault? The problem is rather with the scenario as well as the "one-protocol-should-fit-all" assumption. Thus it is wise to aim at a more modest goal.

4.3 Constantly Raising the Barrier

Progress in the lower layers is one front where MANET has to fight, progress at the problem front is the other. In the past couple of years new problems have become fashionable while working on core ad hoc problems has lost attractivity. This pattern certainly applies to the topic of security – the prime knock out candidate – and also extends to QoS and power awareness. The net effect is that the innovation hurdle has been raised dramatically and that the solution to problem ratio is getting smaller and smaller. Maybe we should learn from the sensor net community and declare some network attributes as being currently "out of reach" because of self-chosen constraints. Probably more balanced would be to accept these new research topics, of course, and to make sure that we do not loose focus on hammering out pragmatic solutions.

4.4 What is the Value of Ad Hoc, Really?

The reaction to the growing "grand complication" should be to concentrate on simple but convincing cases. Better 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 getting full coverage which gives value to the network users (and much less the perfection by which this happens). The case today is that although everybody has and uses IEEE 802.11 hardware, virtually nobody is able to run it in multihop ad hoc mode, as kernels are not shipped with ad hoc routing logic and configuration is a nightmare. The real value of ad hoc networking would be that every device is ad hoc enabled, wherefore resolving penetration hurdles should be the prime target for the ad hoc community.

4.5 Ad hoc for the Lazy

One crucial element in raising the penetration rate of ad hoc functionality is that management must become really simple. We all are lazy and even nerds will become bored after having configured their WiFi cards for the 42nd time. Clearly, there is a management issue here that has to be solved. One could establish a list of the other ten related problems to be looked into too... Again: Cutting a thin slice through the problem space and going for it till the end seems to us the more promising way of getting us lazy people to use multihop ad hoc.

5 Outlook

This papers tries to debunk some of the glory targets of wireless ad hoc networking research. Our main concern is that research on wireless multihop ad hoc networking could look into so many new and important problem areas and technologies that when "acceptable" solutions will eventually emerge nobody is interested anymore. Also, we are concerned about the lack of confrontation of simulation studies and conceptual contributions with actual implementations. A first point to work on is in our opinion the selection of a modest but realistic target e.g., a dozen nodes operating inside the ad horizon. Second, the integration of ad hoc routing with the Internet has to become a no-brainer. And only then, as a third step, we have a chance to push ad hoc functionality into the operating system main stream and the real world.

Acknowledgment

We wish to thank Richard Gold for his careful reading and editing help.

References

1. Ad hoc Network Scaling Research (IRTF group). `https://www1.ietf.org/mailman/listinfo/ans-research`.
2. AODV-UU implementation, 2003. `http://www.docs.uu.se/scanet/aodv`.
3. Andrew T. Campbell, Javier Gomez, and András G. Valkó. An overview of cellular IP. *IEEE Wireless Communications and Networks Conferance 1999 (WCNC'99)*, 2:606–611, 1999.
4. M. Gerla, K. Tang, and R. Bagrodia. TCP performance in wireless multi-hop networks. In *Proceedings of IEEE WMCSA'99, (New Orleans, LA).*, Feb 1999.
5. Piyush Gupta, Robert Gray, and P. R. Kumar. An experimental scaling law for ad hoc networks. Unpublished submission, May 2001, http://citeseer.nj.nec.com/gupta01experimental.html.
6. Gavin Holland and Nitin H. Vaidya. Analysis of TCP performance over mobile ad hoc networks. In *Proceedings of IEEE/ACM MOBICOM '99*, pages 219–230, Aug 1999.
7. J. Hubaux, J. Le Boudec, M. Vojnović, S. Giordano, M. Hamdi, L. Blazević, and L. Buttyán. Toward mobile ad-hoc WANs: Terminodes. Technical Report DSC/2000/006, Swiss Federal Institute of Technology, Lausanne, Jan 2000.
8. Internet research task force. `http://www.irtf.org`.
9. H. Lundgren, D. Lundberg, J. Nielsen, E. Nordström, and C. Tschudin. A large-scale testbed for reproducible ad hoc protocol evaluations. In *3rd annual IEEE Wireless Communications and Networking Conference (WCNC 2002)*, pages 412–418, Mar 2002.
10. H. Lundgren, E. Nordström, and 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)*, pages 49–55, 2002.
11. D. Maltz, J. Broch, and D. Johnson. Experiences designing and building a multi-hop wireless ad hoc network testbed, 1999.
12. D. Maltz, J. Broch, and 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)*, Sep 2000.
13. The Official IETF working group MANET webpage. `http://www.ietf.org/html.charters/manet-charter.html`.
14. R. Ramanathan and R. Hain. An ad hoc wireless testbed for scalable, adaptive qos support. In *2nd annual IEEE Wireless Communications and Networking Conference (WCNC 2000)*, Sep 2000.
15. M. Takai, J. Martin, and 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)*, Oct 2001.
16. C-K. Toh, M. Delwar, and D. Allen. Evaluating the communication performance of an ad hoc wireless network. In *Transaction on Wireless Communications*, Jul 2002.
17. C. Tschudin and R. Gold. LUNAR: Lightweight Underlay Network Ad-hoc Routing. *Uppsala University Technical Report 2003-021*, Jan 2002.
18. C. Tschudin, H. Gulbrandsen, and H. Lundgren. Active routing for ad-hoc networks. *IEEE Communications Magazine, Special issue on Active and Programmable Networks*, Apr 2000.