



Problems When Realizing Ad Hoc Networks: How a Hierarchical Architecture Can Help

Stefan Bouckaert, Dries Naudts, Ingrid Moerman, and Piet Demeester

Ghent University – IBBT vzw – IMEC vzw
Department of Information Technology (INTEC) – IBCN
Gaston Crommenlaan 8/201 – B-9050 Gent—Belgium.
sbouckae@intec.ugent.be – WWW: <http://www.ibcn.intec.ugent.be>

Abstract. Despite the fact that many electronic devices are equipped with wireless interfaces and a lot of publications on wireless ad hoc and mesh networking exist, these networks are seldom used in our everyday life. A possible explanation is the fact that only few of the numerous theoretically promising proposals lead to practical solutions on real systems. Currently, wireless network design is mostly approached from a purely theoretical angle. In this paper, common theoretical assumptions are challenged and disproved, and key problems that are faced when putting theory to practice are determined by experiment. We show how these problems can be mitigated, and motivate why a heterogeneous hierarchical wireless mesh architecture can help in making wireless ad hoc networking a reality.

1 Introduction

Wireless ad hoc networks based on the IEEE802.11 standard have been a popular research topic for many years. Despite the fact that a rising number of electronic devices are being equipped with wireless interfaces as technology gets cheaper, ad hoc networks are rarely used in our everyday life.

How can this be explained, even though it is easy to think of scenarios that could benefit from ad hoc technology?

In [1], the authors answer this question by suggesting that most ad hoc network design focuses on military or specialized civilian applications, making the solutions impractical for everyday life. This is an important observation, however, this paper addresses another, perhaps more fundamental problem. We feel that a lot of research gets stuck in a crucial phase of development: while there are a massive number of initiatives to design wireless ad hoc solutions, few ideas are implemented on actual systems. Unfortunately, as promising as some ideas may be, they do not always lead to good or practical solutions. We have found that, using a purely theoretical top-down approach where implementation is the last step in designing wireless network protocols, architectural decisions are often made which, after months of research, turn out to be impossible to realize because of unforeseen implementation problems.

Wireless research focused on single-interface homogeneous ad-hoc networks for a long time. Recently, an evolution in wireless networking research is observed: more researchers start focusing on multiple interface nodes in mesh topologies. Additionally, several aspects of cross-layer research, such as power control, are gaining popularity. While there is a growing awareness within the research community that simulation of protocols might not be sufficient in order to declare wireless networking protocols working and stable [2, 3], a lot of assumptions are still made while studying old and new topics in wireless research. In this paper, we will verify the validity of common assumptions by experiments, and formulate lessons learned while evaluating several experimental set-ups and real life implementations over the past months.

2 Observations and Experiments

When collecting data from test set-ups using any wireless driver, one should always keep in mind that the driver could be causing errors at a node. We have tried, to the extent possible, to exclude driver-caused errors from the observations below. The observations follow from several real-life experiments and testbeds, using a broad range of hardware. A first testbed consists of 18 Linux nodes, equipped with D-Link AG520 wireless a/b/g PCI cards and external antennas. In addition, a lot of research on mesh networks at our labs is done using several modified 4G Meshcubes with up to 4 wireless mini-PCI a/b/g cards per node. Other tests are performed using Linksys WRT54GL wireless routers with modified firmware. All devices can be powered using batteries, allowing to test real mobility.

2.1 Single Frequency, Single Hop Networking

Assumption: *In an isolated situation, an 802.11 wireless link between two nodes will be of better quality if transmission power is increased.*

Observation:

Consider a test set-up where three identical network nodes, each equipped with a single wireless network interface are stacked on top of each other in a rack. The external antennas are positioned in a triangle, the antennas about 1.5 meters apart. Although at some times data can be sent at the theoretical speed limit, results are very unpredictable. Even if a link seems to be stable for a certain period of time, data rates can drop below one third of the stable rate, seemingly without a reason. In addition, links are not always symmetrical: changing the direction of the traffic flow can result in degraded throughput.

In this specific setup, one solution seems to solve most of these problems. *Reducing* the transmit power of the wireless cards results in highly increased stability. In this case, a transmit power of $10mW$ gives the best results.

Experiment:

When putting single interface Meshcubes to test in an RF-shielded box [4], the same observation is made in a controlled environment: reducing power when sender and transmitter are at close distances, increases networking quality.

Placing the antennas relatively close to each other, as in this case, might seem artificial—there are however a lot of situations imaginable where two single interface nodes are placed at comparable distances: e.g. a user can carry a wireless-enabled PDA and a laptop, or, during meetings, there is a high laptop density.

2.2 Multiple Frequencies, Multihop Networking

Assumption: *when configuring the different wireless interfaces of a multi-interface integrated node to theoretically non-overlapping channels, these different links will not interfere. Capacity of a wireless network can be increased dramatically by adding multiple interfaces.*

Observation:

Reading through literature, lots of innovative solutions involving wireless networking can be found, and numerous protocols are designed to support systems with multiple network interfaces [5, 6]. Many of these solutions are based on the assumption that multiple theoretically non-interfering channels can be operated simultaneously at a node's different interfaces. This seems obvious when considering theory and simulations. Unfortunately, when these solutions are deployed in a testbed, it turns out that this assumption is not necessarily true.

When a two-hop path is created using a traditional single-interface approach, the wireless medium has to be shared between two links. The maximum reachable throughput of the total path thus roughly halves because of a single additional hop. Adding a second wireless interface to the middle node and choosing orthogonal frequencies for the first and second link solves this problem, in theory, at the cost of adding a single extra interface. As simple as this may be, different results can be observed: even when two interfaces are available at the middle node and a two-hop path is constructed using two theoretically non-overlapping frequencies, the throughput does not rise considerably.

Still, in theory, devices supporting the 802.11b and 802.11g (802.11a) standard should be able to use, depending on the region, three (twelve) fully non-overlapping frequencies. Recently, other researchers described the same effects and concluded that wireless nodes with multiple interfaces can suffer severely from self-generated interference between the different interfaces [7].

Experiment: Measurements done at our lab (cf. Fig. 1) show that these effects of self-generated interference can be severely reduced by limiting the transmit power used at the different interfaces and physically separating the antennas of a node. On the other hand, these results also reveal the sad truth that—at least when using off the shelf hardware—a single wireless interface using high transmit power can severely degrade the performance of the other interfaces and surrounding nodes, even when they are set to operate on non-overlapping channels.

When using integrated 802.11a devices such as the Meshcubes in a two-hop test, interference problems are still observed, even when lowering transmit power. The first and second link can be set up separately on orthogonal channels with




Fig. 1. Throughput measurement. Using two distant groups of two, 1.5 meter separated, ethernet-linked WRT54GL routers-representing two nodes with two interfaces each-, two connections are set up in parallel. The Y axis shows the throughput, relative to the maximum throughput of a single, non-interfered stream.

a goodput of over 30Mbps . However, if both links are used at the same time in a multihop configuration, the goodput drops to about 16Mbps even though the processor of the middle node can handle the packet forwarding. We have found that in this case, the problem essentially follows from the fact that the Meshcubes are highly integrated: it's not only the fact that two mini-PCI cards are located right on top of each other that causes problems, but especially the fact that the distance between the antennas is too small. This is not surprising: the antennas in the integrated devices are located very close to each other. At a frequency of 5GHz , the wavelength used for transmission is about 6cm , thus in our integrated devices, up to 4 antennas are separated by less than a single wavelength, resulting in a very unpredictable system. Increasing the distance between two antennas to about 15cm alleviates this intra-node interference and goodput rises considerably. In [7], the authors conclude that these interference effects can be solved by providing an antenna separation of 35dB . However, this separation is hard or even impossible to obtain in mobile devices with a small form factor.

Figure 2 shows additional measurements using WRT54GL routers at a transmission power of 1mW , quantifying the effect of antenna distance on interference between two non-overlapping 802.11g channels: channel 1 and 11. In scenario A, two separate UDP streams (starting at 30Mbps , decreasing until packet loss is minimal) with a UDP packet size of 1470 bytes are set up, sequential at first, then simultaneous. The graphs with circles and squares show the according sum of the average throughput observed at the receiving interfaces. The individual throughput is not shown on the graph, as bandwidth is equally divided between the two flows. The figure shows that the sum of throughputs is reduced by 16.8% when the flows are set up simultaneously with a distance d of 1 meter between the antennas. At a distance of 5cm , throughput is reduced by over 45% compared to non simultaneous transmission. The sum of throughputs reduces both in the sequential and parallel tests when the antenna distance is decreased.

In scenario B of Fig. 2, a single UDP stream with same characteristics is set up. The test packets are now transferred over wire between interface B and D , recreating a typical multihop situation where every receiving interface




Fig. 2. Two scenarios quantifying the effect of the antenna distance on throughput. Link A-B operates on channel 1, link C-D on channel 11. Measurements at $d = 1cm, 5cm, 10cm + k \cdot 15cm, k = 0..6$. Txpower = 1 mW.

has a sending interface nearby. Looking at the graph with the triangles, a first observation is that the end-to-end throughput is higher than half the aggregate throughput from scenario A, for antenna distances d larger than $55cm$: the single flow is transferred more efficiently than two traffic flows originating separately. However, performance drops faster with the decrease of antenna distance. It is also observed that unlike in the previous situation, throughput is very unstable when the antennas are moved more closely to each other: moving the antennas slightly can result in a serious performance drop or increase. For this reason, the graph with the triangles shows the maximum values that were obtained from a large amount of tests. Minimum throughput results are shown with error bars. When the distance between antennas drops below about $40cm$, communication at reasonable throughputs is no longer possible. It is clear that the consequences of interference between neighboring interfaces with nearby antennas are worse in scenario B than in scenario A. The receiving interfaces are disturbed by a signal of the nearby transmitter, even though it is sending on a non-overlapping channel, rendering successful reception of the UDP test packets impossible. Note that in both situations, 802.11ACK messages are sent in the opposite direction of the UDP flows. However, these small packets suffer less from the interference of neighboring interfaces.

This observation raises questions about using multiple interfaces in integrated devices: even if perfect algorithms for using multiple interfaces on devices can be thought of and simulated, it is very likely that the result will not be as expected when deploying them in real integrated systems. We also believe that a testbed with 'full size' computers and wireless PCI cards with external antennas, can not fully represent an integrated end-user device with multiple interfaces.

2.3 Power Adaptation

From previous paragraphs, we have learned, that changing power levels can lead to a more reliable link, but increasing transmit power does not necessarily increase communication quality. Furthermore, it was shown that theoretically non-interfering channels will interfere when using off-the-shelf hardware at mid

to high transmit power, and that integrated systems with multiple interfaces will suffer from self-generated interference if there is no adequate antenna separation. Consequently, a single device transmitting at a relatively high output power may render all surrounding communication virtually impossible.

Lowering transmission power is often considered as a measure of freedom, in order to decrease interference and thus increase the number of possible simultaneous transmissions in a certain area, or to increase the lifetime of battery powered devices [8]. Our experience shows that when using today's 802.11 hardware as a base for multi-interface nodes, power adaptation is not a measure of freedom but rather a necessity in order to guarantee network operation. End-user hardware can and probably will improve in quality over time, however, we predict that it is very unlikely that in the near future palm-size systems will be able to fully take advantage of using multiple interfaces at relatively high output powers, if the interfaces are tuned to neighboring channels. As frequency regulations only allow a small part of the spectrum to be used for unlicensed civilian WLAN communication, it is not easy to provide the required channel separation. When using only two interfaces, putting one interface in the $2.4GHz$ range and the other one in the $5GHz$ might turn out effective with some types of hardware, but it is a mere palliative as this solution can currently not be scaled.

2.4 Hardware Issues

A recurring problem faced during tests with hardware from different vendors, is that changing an interface to some specific channels can result in a wireless link of bad quality. As communicating on those channels is possible using hardware from a different vendor, and bad channels stay bad when replacing hardware with an identical spare, the problem is most likely hardware related.

In general, there is a big difference in stability and performance between hardware from different vendors. In Fig. 3, the spectrum of two mini-pci cards from different vendors, operated at the same power level and frequency are shown. The figure clearly shows that the spectrum of the first card suffers a lot more from frequency leakage than the second card. Consequently, when building a testbed with hardware of the first type, test results will be much more pessimistic than test results with hardware from the second vendor, especially when operating at multiple theoretically non-interfering frequencies.

Although these problems can be solved by replacing faulty hardware with hardware from a different vendor, ad hoc networking protocols will only become successful if a large group of end users can use them instantly without problems, regardless of their choice of vendor. Unfortunately, in a traditional ad hoc network where nodes can join freely, it is wrong to assume that all nodes will react identically to a specific algorithm's action. If the quality of the hardware cannot be guaranteed, control loops should be provided within algorithms to verify whether a certain action has the desired effect. These effects are very hard to model in a simulator and can only be discovered by putting algorithms to test on real-life testbeds.




Fig. 3. Spectrum measurements of two different 802.11a mini PCI-cards operating at same Txpower (15dBm), both using channel 40.

2.5 A Lab Environment is not a Real Environment

Assumption: (i) *If it works in simulation, it will work on a testbed.* (ii) *If it works on a testbed, it will work in real life.*

From previous paragraphs, it is clear that designing algorithms and protocols for wireless systems should preferably not be done solely by considering theory or simulations. Creating a wireless test environment in a laboratory is not an easy task. Not only does it require a lot of -sometimes costly- hardware and space, it is also time consuming. On the other hand, creating these test environments and implementing developed algorithms on actual hardware forces the researcher to develop a system close to reality.

However, there is always a risk that the testing environment itself will lose its value as 'real test case', as over time wireless systems could be tuned—unintentionally—to work great in the testing environment only. This way, a solution that evaluates positive in a testing environment can at the same time be useless when deployed in a random situation. This is a frustrating experience that was witnessed before at our lab: a demonstrator which transmits video over a self-forming and self-recovering multihop mesh/relay network was developed. After the demonstrator had proved to be working perfectly in the lab environment—even when moving the battery fed relay stations through the building—it was taken to a large hangar. Surprisingly, even with relatively short distances between the relaying hardware and line of sight communication, link breaks occurred frequently and maximum throughput was low. In this case, the set-up probably suffered from the absence of the waveguide effect described in [9]: there are circumstances where a wireless signal does not degrade as fast when using devices indoor, compared to using them in an open space.

This example, amongst others, shows that a system should not be declared stable based on a single test environment, and certainly not based on simulations. We believe that wireless ad hoc networking protocols and systems will only be used in our everyday life if their use is not limited to a specific scenario or environment—however, today, a lot of algorithms are evaluated only in simulators using a very specific test scenario.

3 A Heterogeneous Hierarchical Architecture

As shown in previous paragraphs, a lot of problems were observed while bringing wireless ad hoc and mesh networking algorithms to real systems. Some of these problems are hardware related, and one might argue that it is not the task of networking algorithms to correct them. However, wireless systems will always be more unreliable than their wired counterparts, and therefore, algorithms must be able to detect anomalies and react appropriately.

First, transmission power should be chosen wisely: neither too low nor too high. While in a static set-up, transmission power can be set manually by trial and error, there is need for automatic tuning in dynamic environments.

Second, because of interference and hardware related issues, the channel choice has a great impact on the wireless link quality. Problems occur at various layers of the protocol stack when wireless links break due to changing channel conditions or failing hardware. Cross-layer interactions might be required to cope with these complexities.

Third, as small devices with multiple interfaces suffer from self-generated interference, we should focus our research on an architecture which takes this fact into consideration. An algorithm which presupposes perfect separation between multiple interfaces at end-user nodes will most likely never be able to achieve its claimed results when used in real systems.

In a heterogeneous architecture, devices have distinct capabilities and technologies. In a hierarchical architecture, different nodes can belong to different logical groups, for example, backbone nodes and clients. Heterogeneous hierarchical architectures (Fig. 4) have been described in the past, however, we believe that their true potential has not been discovered yet. In [10], the authors describe a (hybrid) wireless mesh architecture. In a wireless mesh network (WMN), two types of nodes are distinguished: mesh routers and mesh clients. Mesh routers hold superior properties concerning processing power, interfaces, available power and memory, enabling them to perform more complex functions. In addition, they have limited mobility compared to the clients, resulting in a wireless mesh backbone. Mesh routers can be added or removed at any time and act as a gateway to other networks such as the internet. In a *hybrid* WMN, mesh clients can connect to the backbone network either directly, or by using a multi hop path through other clients.




Fig. 4. Heterogeneous hierarchical mesh network, i.e. a Hybrid Wireless Mesh Ad Hoc network. Clients connect either directly or through another client to a mesh backbone.

Some benefits of heterogeneous hierarchical networks have been described in the past, such as an increase in coverage, or the (theoretical) ease of set-up.

However, we believe that there are more reasons why hierarchical heterogeneous architectures can help to realize robust wireless networks, and that a conscious choice of networking architecture can help to make certain assumptions that are invalid for homogeneous wireless networks become valid.

Most small and mobile end user devices such as PDAs or smartphones will probably only have a single (high speed) wireless network interface, using the unlicensed bands, enabled at a time, as adding interfaces is suboptimal due to the described interference problems and other limitations such as power and cost. On the other hand, the mesh routers in the backbone can and should have multiple interfaces: they can be bigger in size and antennas can adequately be separated, thereby reducing the interference problems.

Faulty hardware will always be used within a cooperative wireless network, resulting in decreased performance and satisfaction for the end-users. In a traditional ad hoc network, even if one user invests in high-quality hardware, he can still experience bad performance if the person he is connecting through uses faulty hardware. In a heterogeneous architecture, end users can e.g. connect to a mesh backbone which is constructed with hardware of better quality. The nodes which are higher in the hierarchy can be more expensive, as less nodes of higher hierarchy are needed.

In a hierarchical network architecture, operators can invest in high quality wireless backbone nodes. In addition to increasing the number of interfaces, more expensive network nodes could also use alternative technologies such as WiMAX. By using more interfaces or better technologies at higher hierarchical layers, wireless networks become more scalable.

Changing protocols or interface configuration on all end user devices, constructed by different vendors, is harder to achieve than making changes to a smaller group of devices at a higher hierarchical layer, all the more since it is more likely that wireless devices at a higher hierarchical layer are controlled by a single administrator. For example, multi-interface wireless 802.11a backbone routers with (proprietary) cross-layer optimizations can easily provide wireless coverage within a building, or at a fair or festival. By adding an extra wireless interface to every backbone router, configured as an 802.11g access point, end users are able to connect to this backbone with hardware from any vendor, without compromising the quality of the backbone network. If a user has the right hardware and chooses to function as a relaying node and extend the network, he can do this by voluntarily installing the required protocols.

4 Conclusion

In this paper, assumptions that are commonly made when researching wireless ad hoc networking protocols were challenged. It was shown that, whether installing a single interface or multiple wireless interfaces at a node, real-life performance is always worse than can be expected from theoretical models or simulations. We raised questions about the usefulness of embedding multiple interfaces of the same type in a palm-size device, and argued that, in contrast to what is

believed in many research papers, adjusting transmission power is not a measure of freedom but a necessity. We described how a choice of hardware affects the efficiency of algorithms, and how this influences the stability of wireless networks.

In order to test the robustness of algorithms, testing on one or multiple testbeds is a necessity. However, one must keep in mind that positive testbed results do not always imply a stable system under all circumstances. Finally, we argued that a heterogeneous hierarchical wireless mesh network architecture can help solving the observed problems, by reducing the need for miniaturization and providing incentives for network operators and businesses. We believe that, if protocols are developed closer to reality, and more realistic architectural choices are made at the start of a design process, the usability of ad hoc and mesh networks can drastically be improved.

Acknowledgment

S. Bouckaert thanks the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen). This research is also partly funded through the IBBT-Cisco Eye-Sense project.

References

1. Bruno R., Conti M., Gregori E.: *Mesh networks: Commodity multihop ad hoc networks*, IEEE Communications Magazine (2005) 123–131.
2. Tschudin, C. F. Gunningberg P., Lundgren H., Nordström E.: *Lessons from experimental manet research*, Ad Hoc Networks **3** (2005) 221–233.
3. De P., Raniwala A., Sharma S., Chiueh T.: *MiNT: a miniaturized network testbed for mobile wireless research*, INFOCOM 2005. 24th Annual Joint Conference of the IEEE ComSoc. Proceedings IEEE **4** (2005) 2731–2742.
4. Qosmotec: Air Interface Simulator. <http://www.qosmotec.com/>
5. Tang J., Xue G., Zhang W.: *Interference-aware topology control and qos routing in multi-channel wireless mesh networks*, In: Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing, Urbana-Champaign, IL, USA (2005) 68–77.
6. Kyasanur P., Vaidya N. H.: *Routing and interface assignment in multi-channel multi-interface wireless networks*, Wireless Communications and Networking Conference, 2005 IEEE **4** (2005) 2051–2056 Vol. 4.
7. Robinson J., Papagiannaki K., Diot C., Guo X., Krishnamurthy L.: *Experimenting with a multi-radio mesh networking testbed*, In: WiNMee—1st workshop on Wireless Network Measurements, Italy (2005).
8. Jung E. S., Vaidya N. H.: *Power aware routing using power control in ad hoc networks*, Technical report, CSL, University of Illinois, Urbana (2005).
9. Lu D., Rutledge D.: *Investigation of indoor radio channels from 2.4 ghz to 24 ghz*, In: IEEE Antennas and Propagation Society International Symposium. (2003) 134–137.
10. Akyildiz I. F., Wang X., Wang W.: *Wireless mesh networks: a survey*, Computer Networks **47** (2005) 445–487.