Technische Universität Berlin
Seminar "Internet Measurement"
Betreuer: Bernhard Ager and Harald Schioeberg

Architecture and Evaluation of an Unplanned 802.11b Mesh Network

Summary

This text tries to clarify what a mesh network is and which differences and advantages it introduces compared to other kinds of wireless community networks. It is a summary of the publication "Architecture and evaluation of an unplanned 802.11b mesh network", written by John Bicket, Daniel Aguayo, Sanjit Biswas and Robert Morris, from Computer Science and Artificial Intelligence Laboratory of M.I.T.. The text is one of the activities in the seminar "Internet Measurement" of the department Intelligent Networks and Management of Distributed Systems, Research Group Prof. Anja Feldmann in "Technische Universität Berlin".

SS 2007/2008 version of 4 July

Roberto Martín Martín
(rmartinmar@gmail.com)

1. **Introduction**

A mesh network is a kind of network where the nodes work as routers and as "hot-spots" to provide connection to the finals users. A "hot-spot" can be seen as a point of a wireless network where the users can connect. In a mesh network, one node can connect with another node through several hops, not only one hop as it is in a point-to-multipoint network.

The principal aim of this kind of network is to make the deployment and management easy for the users (and, in my opinion, for the internet companies too), without losses of throughput, connectivity or other network properties, compared with other kinds of wire/wireless community networks.

The development of this technology helps to spread quickly new computer networks at low cost, providing a good quality. That is why a lot of mesh network associations are starting to appear, taking advantage of all of these properties [4].

It can be used too in those highly populated areas where a high speed wire network is not yet implanted, like in some cities of the developing countries. It can be useful there to give new kinds of users access to the new technologies.

In chapter two we see how a mesh network works and which characteristics the mesh network measured by M.I.T.[1] group has. In chapter three the procedures and results of the measurements are shown, and in chapter four we can read its most important conclusions.

[1] Massachusetts institute of technology (U.S.A.)

2. Characteristics of the studied mesh network

Traditionally we have two approaches to the community wireless network architecture:

- Multihop network, a well planned network that works with nodes in chosen locations and unidirectional antennas. To this purpose an expertise group is needed, who studies where and how the links between nodes must be placed. This kind of network gives a very good quality in terms of connectivity, throughput and robustness, but has the disadvantage that all the architecture has to be planned and kept by technical experts, and a normal user cannot easily connect without their help.

- “Hot-Spot” network, where all the nodes operate independently. They are not or only loosely connected with each other. The users connect directly to the hot-spots access points and it makes the connection easier. This kind of network gives a lower coverage per wired connection than multi-hop networks.

The mesh networks try to combine the best characteristics of these two architectures: unplanned architecture, easy to deploy and to manage, but with an acceptable performance and coverage.

The studied and measured network has to fulfill a series of design decisions, in order to make it similar to the normal mesh networks:

- Random node placements, which are only determined by where the users want to be connected. The planning of nodes placement is not necessary.

- The antennas used in the node-links are omnidirectional, in other words, they are not focused on other defined points like in planned networks. A normal user must be able to find a connection without the knowledge of where the other nodes are and how to focus an antenna. The antennas have to reach other nodes in their covered area.

- Multihop routing, rather than single-hop routing, to provide better coverage and performance. That means that the packets of information can “jump” between nodes several times till they reach the node with internet access.

- The routing algorithm has to reach the optimal route to provide the highest throughput, but must be ready to find another route if the old connection fails.

2.1. Physical part

We will call our mesh network “Roofnet” due to the place where all the nodes are placed. This Roofnet consists of all the nodes which act as routers and the nodes which act as gateways to the Internet. A gateway is a special node that takes part of the community network and provides Internet connection to it.

The Roofnet is located in a portion of Cambridge, Massachusets, over an area of about four square kilometers. This area is a typical urban area, densely populated and most of the buildings are two- or three-story apartments, but there are taller buildings, too, where eight of the nodes are placed. In this scenario the lines-of-sight of the antennas are often blocked, that is, the antennas of one node cannot directly see the other one [5].

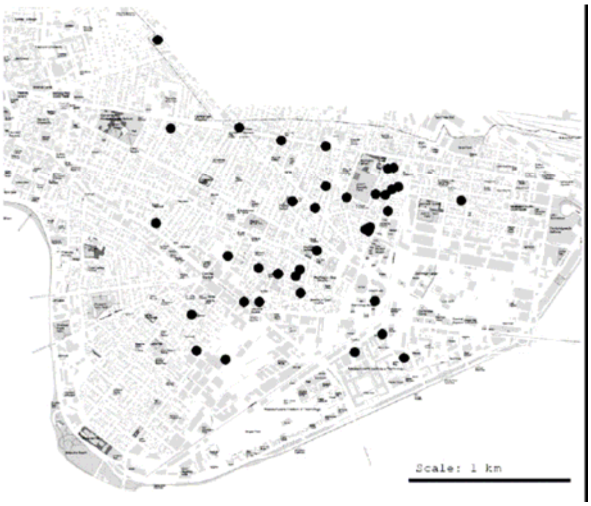

Figure 1: Node’s placement in Cambridge

2.1.1. Hardware

The Roofnet nodes are PCs with an 802.11b card and a roof-mounted omnidirectional antenna. The PC has a CD reader for updating in case of malfunction of the online-update and a hard drive for collecting traces. The user connects directly to the Ethernet port of this PC (wire connection).

The omnidirectional antenna has a –3 dB vertical beam width of 20 degrees, and they have 8 dBi of gain. In the figure two the beam of the antennas is shown.

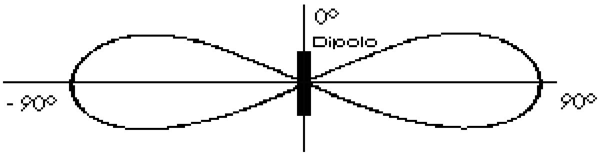

Figure 2: Beam width

This characteristic allows that the antenna has not to be perfectly vertical to obtain a connection, in exchange of smaller gain.

Three nodes in taller buildings have a Yagi directional antenna with 12 dBi of gain and 45 degrees horizontal and vertical beam widths. I think that these nodes are deployed by M.I.T., because one of the requirements was that the users do not have to know how to focus an antenna. We will see in the measurements that they don't modify the architecture so much, because they can be eliminated and the mesh network keeps its properties.

All the radios of the 802.11b wireless card transmit 200 milliwatts of power in the same frequency channel. They operate with RTS/CTS (Request To Send/Clear To Send) disabled.

2.1.2. Software

The nodes run identical software on Linux: a routing software implemented in Click[2], a DHCP server and a web server that allows users to monitor the network status. The DHCP (Dynamic Host Configuration Protocol) permits the node to obtain all the configuration parameters automatically and to give it to the users to configurate the connection parameters (IP address, subnet mask, ...).

The node acts as a DSL modem to the user, and he only has to connect his PC or laptop through the Ethernet interface of the node to obtain Internet connection (transparent service). The user could connect his own wireless access point to the node, in order to make his own private wireless network with Internet access through the mesh network.

The upgrades of the software are made via the mesh network and sometimes via mail, with an upgrade-cd.

The node's software has to resolve a series of requirements:

-Managing addresses to route the information.
-Finding/Reaching a node-gateway which provides Internet connection.
-Choosing the best route to this gateway.

[2] Click is a software architecture for building exible and configurable routers. A Click router is assembled from packet processing modules called elements, which implement simple router functions like packet classification, queueing, scheduling, and interfacing with network devices.[2]

2.2. *Routing part*

In order to send the packets of info to the correct receiver address by the best way, we need a “routing system”. The packets will travel through different networks (Ethernet, Roofnet and the Internet) and we have to define the addresses in each one and how they interact. We need to find the best route for these packets, so we need to say which link is the best (metric) and how we can find it (routing protocol). In every link we will have a optimal throughput (bit-rate selection).

2.2.1 Addressing

Each node needs its own Roofnet address and an IP address in order to run IP applications between Roofnet nodes (for example, IPphone). An IP address is a 32 bit number, which defines uniquely a user in a net (public address in Internet or private address in a private network). The address of each node is assigned by itself, using the 24 lowest bits of its Ethernet address (or MAC address, every Ethernet card has its own unique 48 bit address) as the 24 lowest bits of the IP address. The 8 highest bits are an unused class A IP address block. In the Roofnet, the IP address and the Roofnet address of a node are the same number.

The nodes act as DHCP to the users connected on its Ethernet interface. It means, that it has to assign automatically a IP address to each user when they connect, and the assigned addresses are in the form of 192.168.1.x.

The nodes use NAT (Network Address Translation) between its Ethernet (where the users are connected) and the Roofnet. NAT is a mechanism that allows to connect several users of a private network to another network through only one address. If a user of a node’s Ethernet sends a packet to the Roofnet (or to Internet via Roofnet), the node overwrites the sender IP address and the sender port with its own address and port and it saves a register of the original ones. The answers will be sent to this node’s IP address and to this node’s port. The node only has to see to which port the answer has arrived to “translate” the receiver IP address into the address of the correct user and send it to him. This way the packets in the Roofnet seem to be originated by the node or routed to it, instead of the user.

In the Roofnet we need another addressing system in order to route the packets in every hop. The Roofnet layer encapsulates the IP packets into Roofnet packets in the nodes, that is to say, it writes the Roofnet address of sender and receiver node in every IP packet. Then, the user sends the IP packets to IP addresses, but the nodes send these packets as Roofnet packets to Roofnet addresses (other Roofnet nodes), hop by hop, until the packets reach a gateway.

A diagram of the different networks involved in our scenario is shown in the figure three.

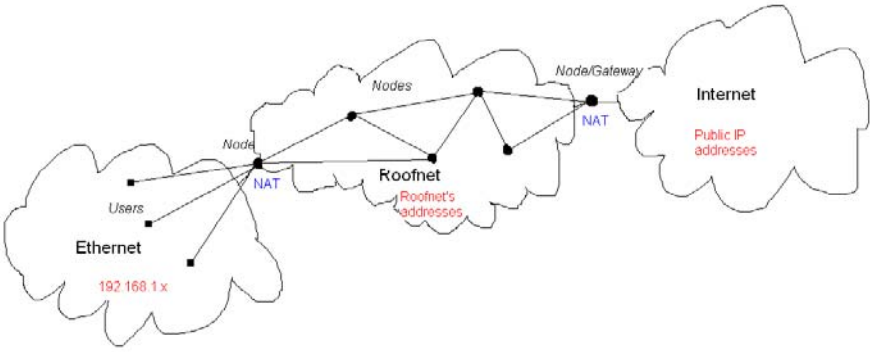

Figure 3: Architecture of the networks

2.2.2 Gateways and Internet Access

Some of the nodes of the Roofnet are gateways too. It means that they have a direct connection to the Internet and a public unique IP address. These gateways share their Internet access with the rest of the nodes of the Roofnet, and with their users.

First of all the nodes have to verify if they are gateways too, so they check if they have Internet access through their Ethernet interface. They ask for an IP address for the Internet via DHCP and try to connect with a well-known Internet server. If that succeeds, the node reports to the Roofnet, otherwise the node acts as a DHCP server and router for the users connected on its Ethernet port and tries to find a gateway-node in the Roofnet.

The gateways use NAT too, for connections from Roofnet to Internet. That works the same way like NAT between Ethernet and Roofnet: the gateway stores the address and port of the sender node and replaces them with its public IP address and port. The answers that the gateway receives from the Internet to this port will be resent to the original node. There will be two different “network address translations” in the packets between a user and the Internet: one between Ethernet and Roofnet and another one between Roofnet and the Internet.

When a user sends a packet, its node selects the best route to reach a gateway and stores which gateway is used by each TCP[3] connection. The route for each TCP connection is decided only one time, when the connection starts. If this route breaks down, the TCP connection fails.

To find the best route to reach the Internet, the nodes use the routing protocol that estimates which route is the best: the one with the optimal metric.

[3] Transmission Control Protocol. It provides a secure connection free of errors through retransmissions and messages of acknowledgment (ACK). TCP is used by the applications and it uses the IP protocol.[6]

2.2.3. Routing Metric

A routing metric is a way to say “how long” is the path between two nodes, and it can be defined by different ways: time to transmit a packet, time until an answer to a sent packet is received, etc.

The routing protocol of Roofnet works with ETT, Estimated Transmission Time. ETT estimates the time spent to send a data packet through a link, taking into account the highest throughput on this link and the delivery probability at this bit-rate. The link with lowest ETT will be chosen by the routing protocol, because it is the fastest link. In a mesh network a single-hop link with high bit-rate and losses is more useful than a slower single-hop link without losses.

To determinate the routing metric each node sends periodic 1500-byte broadcasts to all the reachable neighbors at each available 802.11b bit-rate (1, 2 and 5.5 megabits/second).It also sends periodic minimum size broadcasts at one megabit/second. The receivers keep a trace of the fraction of these information that they receive and send the statistics to the emitter. With this info the emitter node will decide which link is the best.

To determinate the delivery probability at each bit-rate, the protocol multiplies the fraction of 1500 byte broadcast packets delivered and the fraction of 60-byte one megabit/second broadcast packets delivered in the reverse direction, to account for lost 802.11b ACKs. ACK is a small packet to confirm the reception of an info packet.

Finally, the ETT metric for a link is the expected time to successfully send a 1500-byte packet at that link´s highest throughput bit-rate, including the time for the number of retransmissions predicted by the measured delivery probabilities. The ETT metric for a route is the sum of the ETTs for each of the route´s links.

2.2.4. Routing Protocol

Srcr is the Roofnet’s routing protocol. It tries to find the best route between any pair of Roofnet nodes, the route with highest throughput. Srcr nodes maintain a database with some of the ETT metrics between links on the Roofnet and use Dijkstra’s algorithm[4] on that database to find the best route.

To estimate the end-to-end throughput of a route, Srcr uses the following relation:

$$t = \frac{1}{\sum_i \frac{1}{t_i}}$$

Equation 1: Approximation used by Srcr to estimate the end-to-end throughput

where t_i indicates the throughput of the hops between origin and destination. This relation is reasonably accurate in routes with few hops but is not so accurate with large amounts of hops, because of the effect of several throughputs at the same time where

[4] Algorithm used in computer networks to find the shortest path between nodes.

the total throughput must be shared.This leads to collisions and lost packets (which will be seen in chapter three).

To learn the ETT metric of the links, the nodes have three different ways:

- When a node forwards a packet, it includes a metric of the current link on it, in order to inform the other nodes of the route.
- If a node has to send a packet to another one, but cannot find a route, it sends a query to all its neighbors and learns the metric of the links through the responses that it receives.
- Nodes can overhear queries and responses from other nodes and add the metric in those packets to their own databases.

In addition to these methods, the gateways periodically flood (send to each node) a dummy query to allow all the nodes to find a route to them. The gateways learn the route to reach each node from the info that the packets sent by these nodes contain.

A problem is that the flooded queries do not always follow the best route. Therefore the node has to compare the new received info with its old links database in order to find the best route with the lowest metric.

If a link fails, the upstream node sends a notification to each packet source saying that this link is non-operative. If a link stops being the best route, the nodes learn the new metric of the forwarded data packets. If a link changes to have a better metric, the nodes learn it from the dummy queries from the gateways or from the information included in forwarded packets. Then they recalculate the best link with Dijstra’s algorithm.

2.2.5. Bit-rate Selection

Srcr have to decide between all the possible bit-rates to find the best relation between bit-rate and delivery probability. The common algorithms usually choose the bit-rates with less loss-rates, but in a mesh network that is not the best solution. In this network a high bit-rate with up to 50% loss is preferable to the next-lowest bit-rate.

To find the optimal bit-rate, the Roofnet uses SampleRate algorithm. This algorithm, in contrast to the Srcr algorithm, uses real time measurements to decide which bit-rate is the best. Like Srcr, SampleRate calculates the throughput with each bit-rate taking into acount the delivery probability, but these calculations are based on actual data transmissions. Therefore SampleRate has a quicker behavior in changing scenarios than Srcr. Periodically SampleRate sends a data packet of each bit-rate to estimate when a new throughput is better than the one they used before.

3. Evaluation and conclusions of the measured mesh network

The deployment of a mesh network is easier than the deployment of a planned network but it has the risk of a worse performance. Some problems that can appear are:

- Radio ranges might be too short to connect some nodes.
- Many links might be of low quality.
- Nodes can interfere with each other and cause persistent packet loss [3].
- Standard TCP might interact poorly with low-quality radio links.
- The outdoor omnidirectional antennas might pick up unacceptable levels of interference from other users in the same frequency band [5].

All the elements of the mesh network have been studied individually by other research groups, but not as a complete network.

The evaluation is focused on the following characteristics:

- The effect of node density on connectivity and throughput.
- The characteristics of the links that the routing protocol elects to use.
- The usefulness of the highly connected mesh afforded by omnidirectional antennas for robustness and throughput.
- The potential performance of a single-hop network using the same nodes.

Some aspects of the mesh network will not be measured, like the effect of multiple current flows, scalability of the network and protocol, change in network performance over time and the effect of long-term evolution of the network. The measurements are obtained from the received information of the final application, not in the low level.

The results are obtained from four sets of measurements:

- The “multi-hop TCP” data-set: TCP throughput measured with 15 seconds of one-way bulk TCP transfer between each pair of Roofnet nodes.
- The “single-hop TCP” data-set: TCP throughput on the direct radio link between each pair of Roofnet nodes.
- The “loss matrix” data-set: loss rate between each pair of Roofnet nodes at each bit-rate.
- The “multi-hop density” data-set: measures TCP throughput between four choosen nodes, while varying the number of Roofnet nodes that are participating in routing.

With these measurements we obtain some important conclusions: the mesh network works, all the nodes reach good throughput level and the average throughput between nodes is 627 kbits/second. It is comparable with a typical DSL link in terms of throughput and delay and the roundtrip latency averages 39 milliseconds. The behavior of the net is not based on a few number of nodes, it is really mesh.

3.1. *Basic Performance*

This figure shows the distribution of TCP throughput among all pairs of nodes of the Roofnet (cumulative fraction of pairs). The median is 400 kbits/second and the average is 627 kbits/second.

Table 1 shows how many pairs are connected with each number of hops and which throughput and latency is obtained in these connections (in multi-hop TCP).

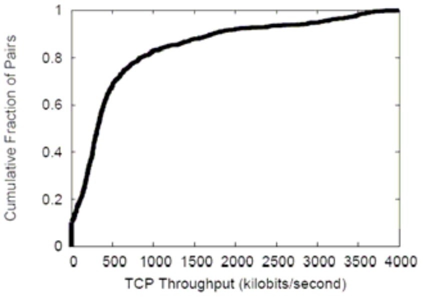

Figure 4: Distribution of TCP throughput

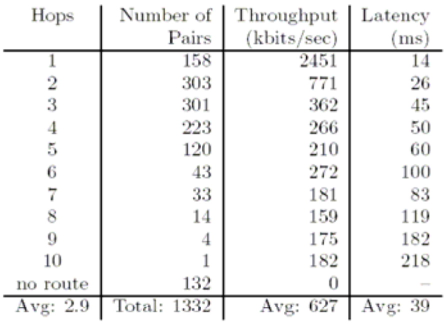

Hops	Number of Pairs	Throughput (kbits/sec)	Latency (ms)
1	158	2451	14
2	303	771	26
3	301	362	45
4	223	266	50
5	120	210	60
6	43	272	100
7	33	181	83
8	14	159	119
9	4	175	182
10	1	182	218
no route	132	0	–
Avg: 2.9	Total: 1332	Avg: 627	Avg: 39

Table 1: Hops needed between two nodes to establish each link

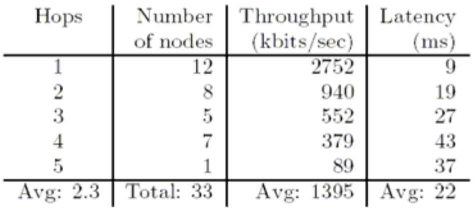

Hops	Number of nodes	Throughput (kbits/sec)	Latency (ms)
1	12	2752	9
2	8	940	19
3	5	552	27
4	7	379	43
5	1	89	37
Avg: 2.3	Total: 33	Avg: 1395	Avg: 22

Table 2: Hops needed to establish a link between each node and a gateway

Table 2 shows how many hops are necessary to connect the nodes to a gateway, and the corresponding average throughput and latency.

The conclusion at this point is that Srcr protocol prefers multi-hops links to the larger-hops links (the number of hops of a link is normally more than one, and there are more nodes that need more than one hop to reach a gateway than the nodes that need only one). Most of the information travels by these links, in both cases, to connect nodes with each other or to connect nodes with a gateway.

3.2. *Link Quality and Distance*

Most of the links are between 500 and 1300 meters long and can transfer about 500 kbits/second, but Srcr uses preferably the shortest links with highest throughput, even if that means that a link needs several hops. We can see it in these figures:

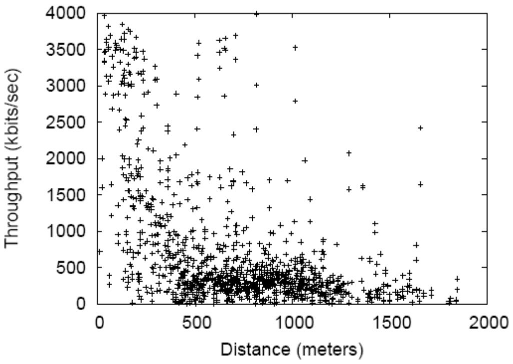

Figure 5: Throughput and distance of each link

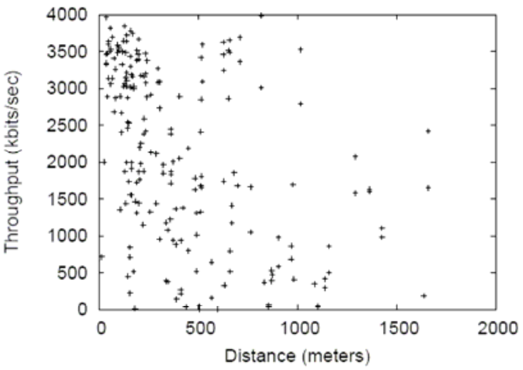

Figure 5: Throughput and distance of each link choosen by Srcr protocol

In the figures (x-distance between nodes of the link, y-throuthput of the link), each point represents a link between two nodes, the first one has all the links and the second one has only the links that Srcr chooses to work. We can see that the links choosen are typically the shortest links with high throughput (second figure, concentration in the upper left corner).

3.3. Effect of density

To measure the effect of density in the mesh network, the behaviour of all nodes of the Roofnet is simulated as follows. For each number of nodes *n*, a random set of *n* nodes are selected. Then an estimated throughput between each pair of nodes is calculated.

The conclusions are that the network starts to approach all the pairs' connectivity characteristics when the number of nodes are 20 or more (about five nodes per square kilometer), that is to say, all the links between the nodes selected have more or less the same properties. With lower number of nodes the characteristics of the links between nodes oscillate between no-connectivity and total connectivity. This is due to the fact that with more nodes there are more posibilities to find a path with short high-quality links.

3.4. Mesh Robustness

One of the aspects of robustness in a mesh network is through how many neighbors a node routes the info. A neighbor of a node is another node to which the delivery probability is 40% or more. If most of the nodes route always through the same neighbor, then it is better to use a directional antenna pointing to this neighbor.

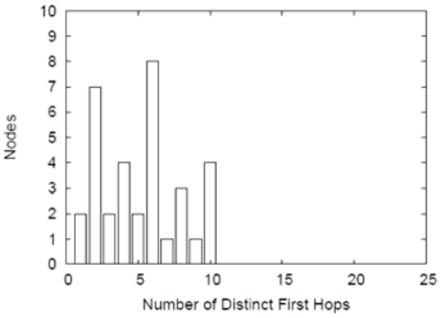

Figure 6: Number of nodes that use every number of distint first hops

The conclusion at this point is that most of the nodes use many different neighbors as their first hop, so the mesh architecture is useful. This fact is shown in figure six, which shows how many distinct first hops a node uses. Only 2 nodes use always the same first hop.

Another aspect is how many pairs are disconnected and how the average throughput decreases if some of the links are eliminated. In the simulation there are four different orders: random elimination, long x fast elimination (to delete the link with highest product of distance and throughput), fastest elimination (to delete the link with highest throughput) or most effect elimination (to delete the link most important to keep the average throughput). The effect can be observed in figures seven and eight.

The fastest links are more important than the long/fast links for throughput (the average throughput decreases faster when we eliminate the fastest links). Also, a lot of links have to be eliminated to reduce the throughput by half. The mesh network is a very robust network.

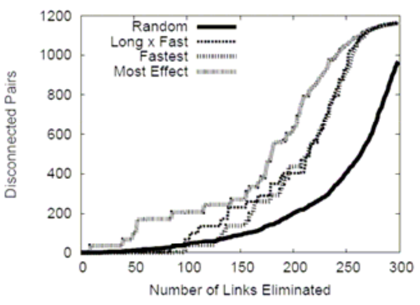

Figure 7: Evolution of the number of disconnected pairs according to the number of links eliminated

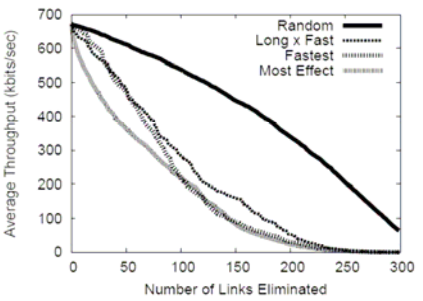

Figure 8: Evolution of the average throughput according to the number of links eliminated

3.5. *Architectural Alternatives*

In order to evaluate the mesh architecture, it is compared with a single-hop network with the same topology and direct radio links between each node and a gateway. That is an access-point network. In these networks the placement of the gateways is crucial. With the optimal choice of the gateways' placements, five gateways are needed to cover all the nodes of the single-hop network (only one gateway in the multi-hop network). 25 gateways are needed with random choice of placement (only 8 in multi-hop network, and with one gateway 34 of 37 nodes are covered). In both cases the multi-hop network increases throughput.

The conclusion is that the Mesh multi-hop Network improves the connectivity and throughput.

3.6. *Inter-hop Interference*

The figure nine shows the expected throughput estimated by Srcr (with equation 1) versus the measured throughput for each link. The measured throughput of single-hop routes (most of the crosses above 2000 kbits/second) differs from the expected throughput but the error is not biased. However the measured throughput of multi-hop routes (lower throughputs) is often lower than the estimated.

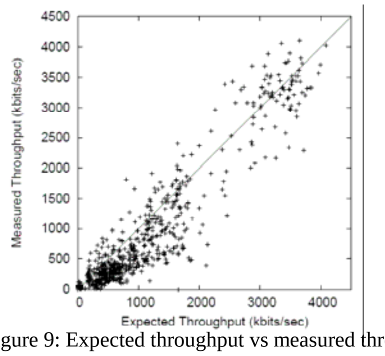

Figure 9: Expected throughput vs measured throughput

It can be due to the collisions between packets of different nodes in the most used links. In some other experiments the packets are delayed until the previous connection is finished, and the results support this hypothesis.

The mechanism RTS/CTS tries to solve this problem but does not improve performance. It is not effective at avoiding collisions.

The conclusion at this point is that the inter-hop interference can be a limiting factor in the effective throughput of the mesh networks.

4. Conclusions

The mesh network architecture provides reasonably good throughput and connectivity, with very robust behaviour. It improves the results obtained with a single-hop network. Furthermore, the mesh network is easy to deploy and self-configurated, so it is an optimal solution for average users, who do not need to know how the network works. The mesh network does not need many gateways to offer Internet connection with an average throughput, like the throughput offered by DSL links.

All these advantages are the reason why the mesh network architecture is being used more and more all over the world. Here in Berlin there are a lot of Non-Governmental Organizations or commercial associations that group all the mesh network's users to exchange news, problems and solutions [3, 4].

5. References

J. Bicket, D. Aguayo, S. Biswas & R. Morris. Architecture and Evaluation of an Unplanned 802.11b Mesh network, MobiCom 2005

[2] The Click Modular Router
Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti and M. Frans Kaashoek.
Laboratory for Computer Science, MIT.

[3] Humbolt-Universität zu Berlin. Computer Science Departament.
http://sarwww.informatik.hu-berlin.de/research/wireless_mesh/wireless_mesh.htm

[4] Sombrutzki R., Zubow A., Kurth M., Redlich JP. Self-Organization in Community mesh networks The Berlin RoofNet. http://www.citeulike.org/article/1378877

[5] Mathias Kurth, Anatolij Zubow, Jens-Peter Redlich. Multi-channel link-level measurements in 802.11 mesh networks. http://portal.acm.org/citation.cfm?id=1143736

[6] Wikipedia. http://es.wikipedia.org/wiki/Portada