Architecture and Evaluation of the MIT Roofnet Mesh Network (DRAFT)

John Bicket Sanjit Biswas Daniel Aguayo Robert Morris

M.I.T. Computer Science and Artificial Intelligence Laboratory
{jbicket, biswas, aguayo, rtm}@csail.mit.edu

ABSTRACT

This paper describes the design and performance of Roofnet, an 802.11b mesh network for Internet access. Roofnet consists of more than 40 nodes spread over an eight square kilometer urban area. Its principle design goal is to allow unmanaged deployment and operation, to better support the growth of community networks with open access policies. This goal has influenced many levels of the design, from the choice of antenna to the routing metric to the IP address allocation plan.

A key challenge for an unmanaged wireless network is that lack of careful planning might render the network's performance unusably low. For example, acceptable performance might require siting relay or wired gateway nodes in optimized locations, or using multiple frequencies to reduce interference, or improving signal strength with directional antennas. This paper evaluates Roofnet to show that it has good performance despite lack of planning: the average inter-node throughput is 88 kilobytes/second, even though the average inter-node route is three hops. The paper answers a number of questions about Roofnet's architecture, such as the performance impacts of node density, gateway placement, and path length; finally, the paper compares Roofnet's performance to the performance of an access-point network using the same set of nodes.

1. INTRODUCTION

This paper describes the architecture and evaluates the performance of Roofnet, a multi-hop wireless Internet access network. Roofnet consists of more than 40 nodes spread over a few square kilometers of a city; each node consists of a PC running special software, an 802.11b radio, and a roof-mounted antenna. The purpose of the research project is to understand how to build better cooperative community wireless networks.

The main principle driving the design of Roofnet is that it should be easy to deploy. Non-technical users should be able to install their own equipment without having to configure it or understand how their node will fit into the larger network. A Roofnet user with a separate wired Internet connection should be able to share that connection with other Roofnet users easily. Roofnet uses automatic address allocation and gateway selection in order to completely avoid configuration for users.

A direct result of the desire for minimum deployment effort is that Roofnet's nodes are placed with essentially no consideration for radio propagation. Each user is a volunteer who mounts an omni-directional antenna on his or her roof and joins the network through whatever neighboring nodes happen to be directly reachable. This approach differs from the more traditional "engineered" network in which nodes are positioned based on propagation surveys or simulations, and in which directional antennas are used to create specific high-quality links.

Reduced efficiency is the price that a non-engineered network must pay in return for ease of deployment. Below a certain point, however, low efficiency might result in an unusable network. The risks are that radio ranges might be too short to connect some nodes; that many links might be low quality; that the Roofnet nodes might interfere with each other and cause persistent packet loss; that standard TCP might interact poorly with low-quality radio links; and that outdoor omni-directional antennas might pick up unacceptable levels of interference from other ISM-band users throughout the city.

This paper describes initial experience with Roofnet. First, it outlines Roofnet's architecture with special attention to aspects that made it easy to deploy. Second, it outlines some of the techniques Roofnet uses to achieve high performance. Third, it measures the user-visible performance of the resulting network. Fourth, the paper explores the effect of varying parameters such as node density on the performance of the network in order to generalize the results. Finally, the paper evaluates how well the most common alternate architecture (base stations) would have performed. This paper complements our previous publications that explain the basis for Roofnet's routing protocol [5] and investigate the details of single-link packet loss patterns [1].

Some of the specific observations made in this paper are as follows. Roofnet works: with few exceptions, users see throughput and delay comparable to DSL links. Throughput decreases with number of hops, but even 8-hop routes get 20 kilobytes/second. Throughput increases with node density, since the routing protocol makes use of short links

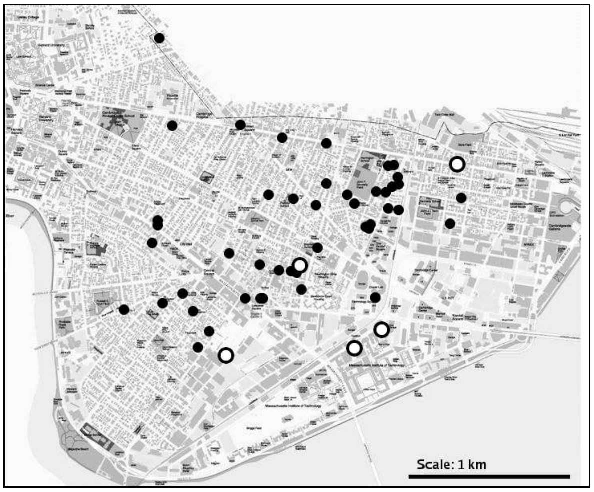

Figure 1: A map of Roofnet. Each dot represents a wireless node, and the five hollow dots show the locations of gateways to the wired Internet. The underlying map data is provided by the city of Cambridge; the empty space at the north end of the map is the neighboring town of Somerville for which no map was available.

and drives them at high bit-rates. While radio ranges are in general short (the median link in Roofnet is about 150 meters), there are a significant number of surprisingly long links. While Roofnet includes a few nodes positioned on the tops of tall buildings, its performance and robustness do not greatly depend on any small set of nodes; it is a true mesh. Finally, while a single-hop base-station architecture would have been possible for Roofnet, for any given number of wired access points, Roofnet's multi-hop forwarding improves coverage and throughput.

2. NETWORK ARCHITECTURE

This section presents an overview of the Roofnet architecture, highlighting the particular techniques that make the network easy to deploy.

Roofnet is deployed over about nine square kilometers in Cambridge, Massachusetts, just north of the Massachusetts Institute of Technology. Figure 1 shows a map of the network and the surrounding area. This area is urban and densely populated. Most buildings are three or four-story homes; all but eight Roofnet nodes are located in such buildings, while the others are in taller buildings. Line-of-sight signal propagation between nodes is often obstructed.

2.1 Hardware

Each Roofnet node consists of a small PC, an 802.11b card, and an 8 dBi omni-directional antenna. The PCs were selected to require the minimum amount of space in the user's home, and to generate as little noise as possible (they have no fans). The PCs are are small enough that an entire Roofnet kit (including PC, antenna, mounting hardware, and cable) can be carried by one person. Each PC has an Ethernet port for the user's network connection, a hard drive for collecting traces, and a CD reader so that a user can upgrade the node if an ordinary over-the-network upgrade fails.

Figure 2: A typical Roofnet antenna installation. The cable leads to a node inside the building.

Each omni-directional antenna has a vertical -3 dB beam angle of 20 degrees. This beam angle is relatively wide, and as a result the antenna need not be perfectly vertical. Higher-gain antennas have narrower beams, and experience shows that users don't install them close enough to vertical and wind tends to tilt them; a tilted narrow-beam antenna radiates mostly into the sky and into the ground. Each antenna is connected to its node with coaxial cable which introduces 6 to 10 dB of attenuation. Figure 2 shows a photograph of a typical installation. The antenna is about 17 inches long. Three nodes have 12 dBi Yagi directional antennas with 45 degree horizontal (and vertical) beam widths.

The wireless network cards are a generic brand based on the Intersil Prism 2.5 chipset. Roofnet uses 802.11b primarily because of the extent of available driver support, and hardware which is known to work. 802.11g might be a desirable alternative, as it should be more resilient to multi-path interference, which is a significant source of packet loss in the network [1]. 802.11a radios operate at 5 GHz; cable loss at this frequency would require the radio (and node) to be on the roof with the antenna. Future versions of Roofnet will probably use this arrangement, with smaller nodes in waterproof boxes and power-over-ethernet.

Roofnet operates in the Prism 2.5 chipset's non-standard "pseudo-IBSS" mode. This is similar to standard 802.11b IBSS mode, in that nodes communicate directly without intervening access points. Pseudo-IBSS omits the BSSID (network ID) mechanism and does not use synchronization beacons. This solves a problem with standard IBSS mode in which partitions form with different BSSIDs despite having the same network ID. These partitions made it impossible to operate Roofnet reliably.

All Roofnet nodes use the same 802.11b channel.

2.2 Software and Auto-Configuration

Each Roofnet node runs identical turn-key software, which consists of Linux, routing software implemented in Click [10], a DHCP server, and a web server so users can monitor the network status. Most users pick up nodes from us at our lab with software pre-installed. We regularly upgrade all the nodes' software over Roofnet, and occasionally by mailing out installation CDs.

From the user's perspective, the node functions similarly to a cable or DSL modem: the user connects the node to its antenna, power, and an Ethernet hub, turns on the node, waits a minute, and then gets Internet access. The Roofnet node acts as a DHCP server, so the user's network connection usually "just works."

In order for Roofnet nodes to be completely self-configuring, the software must automatically solve a number of problems: allocating an address, finding a gateway between Roofnet and the Internet, and choosing a good multi-hop route to that gateway.

2.2.1 Addressing

Roofnet does not use IP routing internally: it has its own routing protocol and packet formats. However, Roofnet nodes still need unique identifiers. Roofnet uses 32-bit identifiers that look like IP addresses, to allow each node to use the same identifier inside Roofnet and at the IP level to talk to the outside world. However, since Roofnet gateways use network address translation on the way to the Internet, the identifiers need not be globally routable IP addresses; they only need to be unique within Roofnet, and use an IP prefix that is unlikely to be in use on the Internet.

A Roofnet node uses the lower 24 bits of its Ethernet address as the low bits of its ID, and sets the high byte equal to an IP class A network number that has been set aside for private networks.

The wired Ethernet port, on which the node acts as a DHCP server, uses a private class C net number. Hosts that are connected to this port and issue DHCP requests are assigned addresses on this subnet, with the Roofnet node as the default gateway.

Each Roofnet node uses NAT between the user's Ethernet and Roofnet, so that the IP packets that Roofnet carries appear to be from the Roofnet node itself.

2.2.2 Gateways and Internet Access

Roofnet is designed assuming that a small fraction of Roofnet users will share their wired Internet access links. This is entirely on a voluntary basis; a Roofnet node must be explicitly configured to share a wired connection. Multiple consumer DSL ISPs in the Cambridge area have AUPs that allow wireless sharing of Internet connections (Speakeasy and Cyberion), so there is no contractual difficulty with this approach. The main challenge have to do with IP addressing.

Each gateway is connected to its Internet connection on its wired Ethernet port, and functions as a NAT for traffic between Roofnet nodes and the Internet. The source address of packets coming from Roofnet is rewritten with the IP address of the gateway's wired interface.

Each gateway periodically floods advertisements through Roofnet indicating that it offers gateway service. When routing traffic for the Internet, each node selects the gateway to which it has the best route metric. The node keeps track of which gateway is being used for each TCP connection, since the gateway's use of NAT means that each TCP connection must continue to use the original gateway even if the routing protocol later decides that a different gateway has the best metric. Each new TCP connection uses the gateway with the best metric when the connection starts. If a Roofnet gateway fails, existing TCP connections through that gateway will fail (because of the NAT), but new connections will use a different gateway.

Roofnet currently has five Internet gateways. Two are located in ordinary residences, one is on the roof of a six-story MIT building, and two are in windows of other MIT buildings.

2.3 Routing Protocol

If Roofnet were based on directional antennas, the radio links could be engineered to act like low-loss wired point-to-point connections, so that existing IP routing protocols would work. Many existing community mesh networks take this approach [16]. Roofnet, in contrast, uses omni-directional antennas for ease of deployment and robustness. Because they are not engineered, Roofnet's radio links have a wide range of packet delivery ratios, complicating the tasks of routing, forwarding and bit-rate selection.

Roofnet uses a routing protocol called RNR (Roofnet Routing). RNR is a pro-active source-routed protocol inspired by DSR [8] and similar in overall structure to MCL [6]. RNR performs its own measurement-based transmit bit-rate selection and chooses bit-rate-aware and loss-rate aware routes using a routing metric derived from ETX [5]. These techniques improve throughputs by a factor of four for typical multi-hop routes.

The rest of this section briefly discusses these techniques. A more thorough presentation and evaluation of RNR is the focus of another paper which is currently in submission. The name of the protocol has been changed here to protect the anonymity of the pending submission.

2.3.1 Bit-rate selection

The Prism 2.5 firmware chooses among the 802.11b transmission bit rates (1, 2, 5.5, and 11 megabits/second) with the following algorithm: after four transmission failures (failures to receive an 802.11 ACK) in a row, reduce the bit-rate; after 10 seconds without a transmission failure, return to 11 megabits/second. This algorithm often results in sub-optimal bit-rate selection; for example, if a link delivers 60% of packets at 11 megabits/second, that rate will delivery higher throughput than even a perfect 5.5 megabits/second, yet the link is likely to drop four packets in a row fairly frequently.

RNR chooses the transmit bit-rate to each neighbor itself, rather than using the Prism 2.5 algorithm. RNR's choice is based on the loss rates of periodic broadcast probe packets sent at each of the 802.11b bit-rates; RNR chooses the bit-rate that will achieve the highest throughput after accounting for the cost of 802.11 re-transmissions.

2.3.2 The ETT Metric

Prior work presented the estimated transmission count (ETX) metric [5], which favors routes with low loss rates since they are likely to provide high throughput.

RNR extends ETX by using routes that minimize the expected transmission time (ETT) required to deliver a packet across the network. ETT takes into account 802.11b transmit bit-rates as well as loss rates. RNR calculates ETT metrics based on loss rates of periodic broadcast probe packets sent at each bit-rate.

3. EVALUATION

This section evaluates whether Roofnet works well, both in absolute terms and relative to its architectural goals and the performance of alternate architectures.

First, basic measurements of throughput and latency over the network are presented. These measurements show that performance on Roofnet is quite good: over all pairs, throughput averages 88 kilobytes/second, and latencies average 34 milliseconds. Users accessing the Internet gateways see on average 190 kilobytes/second, with an average latency of 23.7 milliseconds.

Second is a presentation of radio range measurements and an exploration of the performance of different-size subsets of the Roofnet nodes. These data provide a feel for how well other mesh deployments with different node densities might work.

Third is an assessment of how Roofnet takes advantage of the highly connected mesh provided by omni-directional antennas. This investigation shows that most Roofnet nodes take advantage of many neighbors, but that there are a few nodes in the system which contribute disproportionately to overall throughput.

Fourth is a comparison of Roofnet's multi-hop approach with a single-hop access point architecture. Five well-placed access points would be required to cover all Roofnet nodes, and even more would be required to match the average throughput that Roofnet provides.

Finally, the evaluation presents measurements that suggest that inter-hop collisions are a major limiting factor in multi-hop throughput.

3.1 Methodology

Throughput results were obtained using a 15 second one-way bulk TCP transfer between each pair of nodes; throughput is measured in terms of data bytes delivered to the receiving application. Each transfer is preceded by a 30-second quiet interval during which the sender sends 84-byte pings once per second to establish the route and measure latency. The throughput result and the median round-trip ping time are recorded along with the route in use at the end of the transfer. The measurements were taken with RTS/CTS disabled. (Measurements with RTS/CTS enabled are presented in Section 3.6.)

Roofnet ordinarily provides Internet access to its users, and this service was not disabled during these experiments. Thus the experiments may be affected by Roofnet traffic as well as other uses of the 2.4 gigahertz ISM band.

3.2 Basic Performance

Figure 3 shows the distributions of TCP throughputs among all pairs of Roofnet nodes. The median is about 50 kilobytes/second, though there are pairs with much higher throughputs. Much of the distribution is explained by hop-count, since the successive nodes in multi-hop routes cannot send at the same time. Table 4 arranges the throughput data by hop-count in order to illustrate this point: the routes with low hop-counts have much higher throughputs than those with many hops.

For comparison, Figure 5 shows the theoretical maximum throughput for various bit-rates over multiple lossless hops. These numbers are calculated by accounting for packet transmission times, inter-frame gaps, and link-level acknowledgments. Together, the tables suggest that Roofnet's one-hop routes operate at an average speed consistent with the 5.5 megabit transmission rate, but that the links in longer routes operate at lower efficiency. The reason for the lower efficiency is probably collisions and packet loss due to interference among the hops in a route (see Section 3.6 for more evidence).

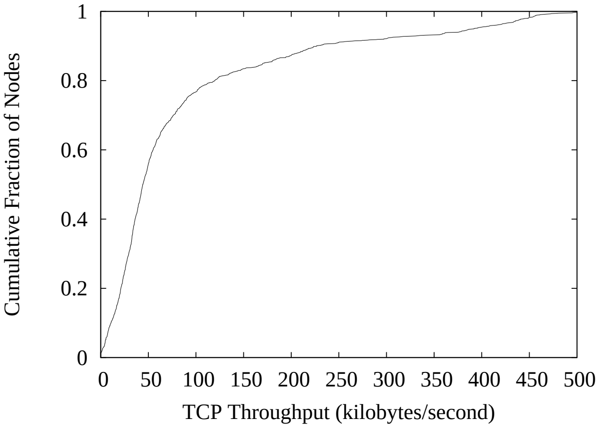

Figure 3: CDF of the TCP throughput between each pair of nodes in the network.

Hops	Pairs	Avg. Throughput (kilobytes/second)	Avg. Latency (milliseconds)
1	179	316.4	12.5
2	354	97.7	22.2
3	354	46.0	38.6
4	256	33.9	44.2
5	127	27.3	60.8
6	54	30.5	80.8
7	38	22.5	72.4
8	17	20.5	97.9
9	6	19.2	121.2

Figure 4: Average TCP throughput and round-trip ping latency between each pair in the network, arranged by the number of hops in a typical path between the pair.

Most Roofnet users talk only to the Internet gateway with the best metric, and thus use routes with fewer hops than the average of the all-pairs routes. Table 6 shows TCP throughputs to each node from its chosen gateway, again arranged by hop count. The maximum hop-count is four instead of nine. The averages for each hop-count are higher than in the all-pairs data because the five Roofnet gateways are slightly better placed than the average Roofnet node. Even at four hops the average of 47 kilobytes/second (or 376 kilobits/second) is comparable to many DSL links. The all-pairs average at nine hops is lower but still comfortably greater than a dial-up modem.

The tables also show round-trip latencies for 84-byte ping packets, to capture the interactive delays a user might see on a relatively idle network. Latency is affected by per-hop processing time as well as 802.11 re-transmissions and back-offs when packets are lost. The bulk of the latency is caused by the packet processing time, though it is likely that this could be reduced substantially by further tuning

	Max Throughput (kilobytes/second)		
Rate	1 Hop	2 Hops	3 Hops
1	76	38	25
2	142	71	47
5.5	324	162	108
11	507	253	169

Figure 5: Theoretical loss-free maximum UDP throughput over one, two, and three hops for each 802.11b transmit bit-rate.

Hops	Nodes	Avg. Throughput (kilobytes/second)	Avg. Latency (milliseconds)
1	18	357.2	9.7
2	10	112.0	17.5
3	9	52.8	43.7
4	7	47.3	43.0

Figure 6: Average TCP throughput and round-trip ping latency to each Roofnet node from its chosen Roofnet gateway, arranged by number of hops in the route from that gateway.

of Click and Roofnet software. The measurements suggest that interactive latency is acceptable over a few hops but would be annoying over nine hops.

3.3 Radio Range and Density

One of the reasons that Roofnet works well is that its nodes are dense enough that they are well connected. This section attempts to provide insight into how well a hypothetical network of lower density might work. Since the relevant density measure is nodes per nominal radio coverage area, the section first presents radio range measurements from Roofnet. Then the section explores the effects of density by presenting performance measurements for varying-size subsets of Roofnet.

Figure 7 shows the relationship between distance and delivery ratio for all Roofnet node pairs, at both 1 and 11 megabit/second transmit rates. The data were obtained with experiments in which each node in turn broadcasts 1500-byte packets. There is only weak correlation between distance and delivery ratio in Figure 7. There are both surprisingly long links with high delivery ratios and short links with low delivery ratios. Figure 8 shows the same data in map form, with darker links indicating higher delivery ratios.

Just because long links exist does not mean that they are useful. Figure 9 shows the distribution of the distances of links actually used in at least one Roofnet route. While the median is just under 500 meters, 10% of the useful links are more than 100 meters long. RNR is often willing to use long-distance links with mediocre delivery ratios, since a one-hop link with 51% delivery ratio can be better than a two-hop route with 100% links. Figure 10 illustrates this with the CDF of the delivery ratios of the links used by RNR: the median delivery ratio is 80%, but 10% of used links have delivery ratios less than 50%.

The main significance of radio range is its effect on node density. A mesh network won't work at all if nodes are too sparse for radio communication. Above a certain density threshold the network will become connected. Increasing density beyond that point may potentially increase performance, since links will become more reliable and support higher transmission rates.

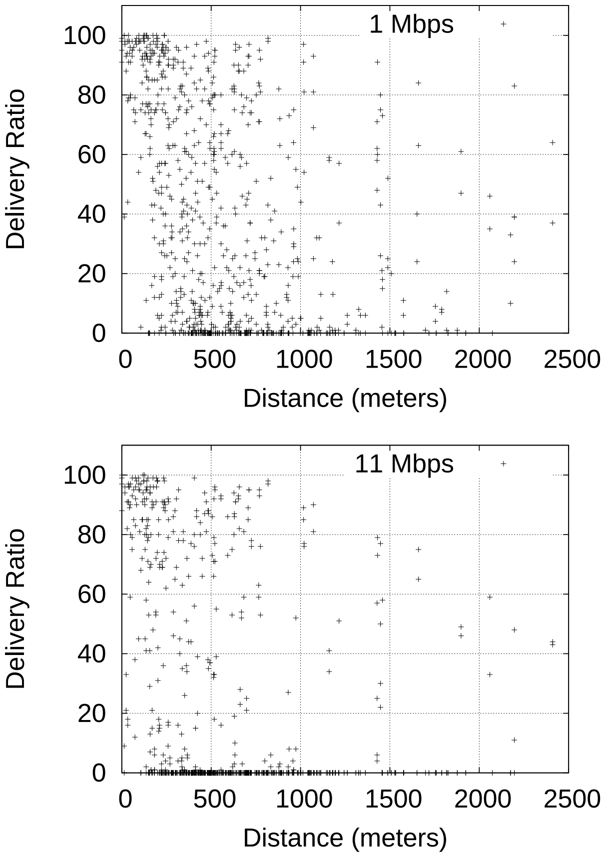

Figure 7: Delivery ratio versus distance for all Roofnet node pairs. The top graph shows delivery ratio for 1 megabit per second and the bottom graph uses 11 megabits per second. These graphs demonstrate there is little useful correlation between distance and delivery rates, and that there are useful links that are surprisingly long.

To explore how less-dense versions of Roofnet would perform, we measure throughputs for various size subsets of Roofnet. Each subset is randomly chosen, and nodes not in a given subset do not participate in Roofnet. In order that throughput measurements from different subsets be comparable, the throughputs are measured between a fixed set of four "edge nodes" shown in Figure 11; these four nodes are members of every subset. No pair among the four edge nodes has usable direct radio connectivity.

The density experiments were performed with the following questions in mind: First, how dense do nodes need to be before the network is connected? Second, if the network is connected, does adding more nodes in the same area affect end-to-end throughput? To explore these questions, we selected a set of nodes on the edge of the network and varied the number of other nodes that participated in routing to explore the effect of node density on throughput and con-

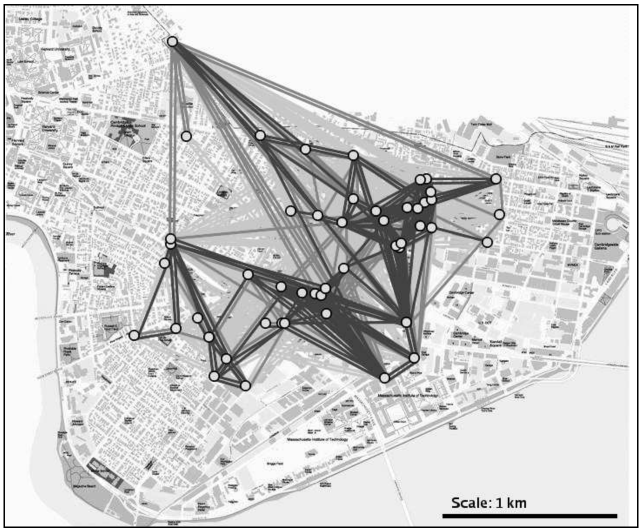

Figure 8: The gray level of each line indicates the delivery ratio at 1 megabit between the indicated node pair. Three gray levels are shown, corresponding to less than 40%, 40% to 70%, and 70% to 100%.

nectivity.

Figure 12 shows how density affects connectivity, throughput, and hop-count among the four edge nodes. The top graph shows that the four nodes almost always become fully connected with 10 other nodes; although Roofnet's node layout is not uniform, this corresponds to a density of about ten nodes per square kilometer. The middle graph shows that adding more nodes beyond that point increases throughput. The reason for this is that the higher density gives RNR more choices, and in particular more short-distance links with lower loss or higher usable transmit bit-rate. The lower graph bears this out: the average hop-count increases with density.

3.4 Mesh Robustness

This section investigates how well Roofnet takes advantage of the routing choices afforded by a mesh architecture and omni-directional antennas.

The most immediate measure of a mesh network's robustness is the number of potentially useful neighbors each node has. Figure 13 shows a histogram of the number of neighbors for each node, where a neighbor is defined as a node to which the delivery ratio is 40% or more. Most nodes have many neighbors, though there are a few isolated nodes.

Of course, if a node never routes through one of its neighbors, the neighbor's value is questionable. For example, if most nodes routed through only one or two neighbors, it might be worth considering building a network with directional antennas pointing just at those neighbors. Figure 14 shows a histogram of the number of unique first hop neighbors used by each node when routing to all the other nodes in the network. While some nodes do indeed route through only one or two neighbors, the majority of nodes use many more neighbors. In this sense Roofnet makes good use of the mesh architecture in ordinary routing.

The fact that there are a few nodes in Figure 13 with large number of neighbors suggests the possibility that Roofnet's performance depends heavily on just a few "super" nodes with good radio connectivity. If such nodes existed, their failure might disproportionately degrade throughput.

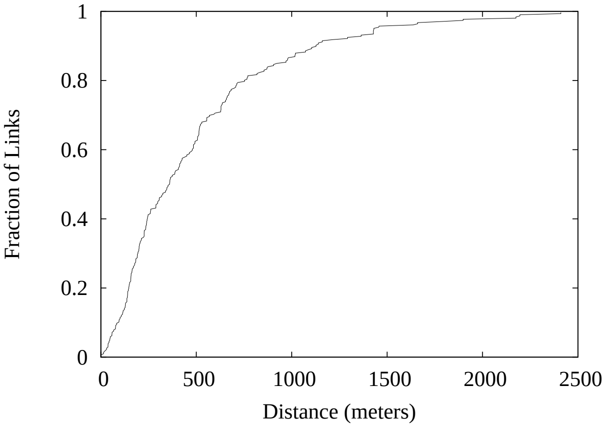

Figure 9: The distribution of distances of the unique links used by the RNR routes between all Roofnet node pairs.

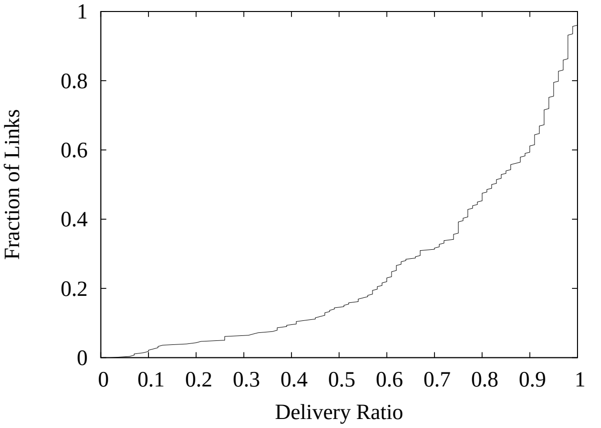

Figure 10: The distribution of delivery ratios of the unique links used by the RNR routes between all Roofnet node pairs, measured with 1500-byte broadcast packets at 1 megabit/second.

Figure 15 shows the effect on throughput of cumulatively eliminating the best-connected nodes in Roofnet. The throughputs plotted are the averages among the four edge nodes in Figure 11. The best-connected nodes are identified by looking for the nodes that appear the most number of times in all-pairs routes.

The data in Figure 15 shows that the best two nodes are important for system performance, since losing both decreases the average throughput by 43%. The decreases after that point are more gradual.

3.5 Architectural Alternatives

One way to evaluate Roofnet's architecture is to compare it to a traditional wireless network based on access points. In such a network, the access points are connected to the wired Internet, and each node connects directly to one access point over a single hop.

Examples of such access-point networks include cellular

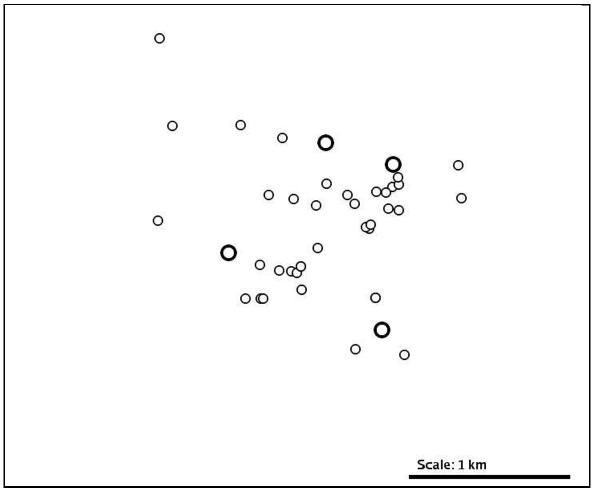

Figure 11: A map of Roofnet indicating the positions of the four edge nodes that participate in the density experiments.

telephone networks, campus-scale 802.11 deployments such as Wireless Andrew [2], and the 802.16 standard for outdoor metropolitan area radio networks [4]. Many community networks, such as Wireless Leiden [16], use a small mesh backbone of point-to-point wireless links, with the backbone nodes providing single-hop access-point service to subscribers.

This section provides an estimate of how Roofnet would compare – in terms of coverage and average throughput – to an architecture in which some subset of existing Roofnet nodes act as wired access points. The analysis was conducted off-line, using the measured TCP throughput between all N^2 pairs in the network presented in Section 3.2. In addition, direct, single-hop measurements were taken between all pairs of nodes in order to simulate the access point architecture; RNR's bit-rate selection technique was still used.

3.5.1 Optimal choice

Figure 16 presents the results for optimal choice of access points and gateways. The columns show the number of nodes which have non-zero throughput to the wired network and the average throughput to the best gateway or access point.

Each line in the table represents the addition of another gateway or access point. Each successive addition is the Roofnet node that would provide single-hop coverage to the most additional nodes. in order to provide single-hop coverage to the most Roofnet nodes. Ties are broken by choosing the gateway that results in the highest average throughput.

The data show that five access points are needed to cover the entire Roofnet network. More would be required to match the average throughput provided by Roofnet's current Internet gateways. Moreover, this table shows that for any given set of gateways or access points, Roofnet's multihop forwarding provides higher average throughput.

The five optimal gateways for the Roofnet network turn out to be nodes located on three-story residences, not the tallest buildings in the network. This is likely due to the fact

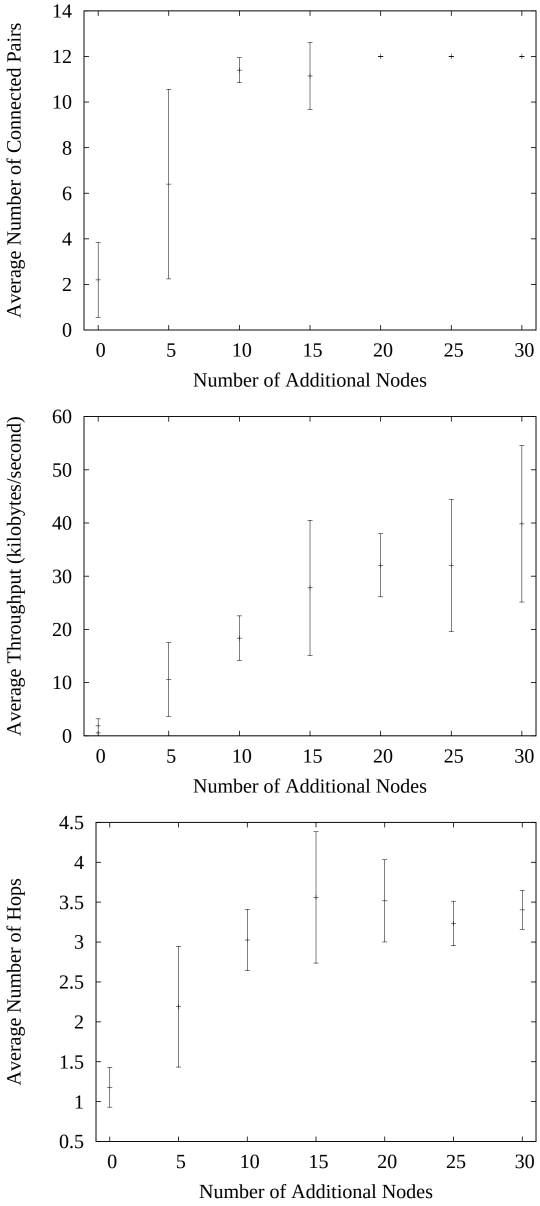

Figure 12: The effects of node density on connectivity and throughput. The x axes show the number of randomly chosen nodes added to the system in addition to the four edge nodes in Figure 11. The y axis in the top graph shows the average number of pairs among the four edge nodes that achieve throughput of more than one kilobyte/second. The y axis in the middle graph shows average throughput among the four edge nodes. The y axis in the bottom graph shows the average hop-count of the routes among the four edge nodes. The error bars show average and standard deviation over five runs with different random choices of nodes added. Increasing density rapidly causes the network to be connected, and more gradually increases throughputs; the reason for the latter is that higher density allows routes to be constructed from more higher-quality hops.

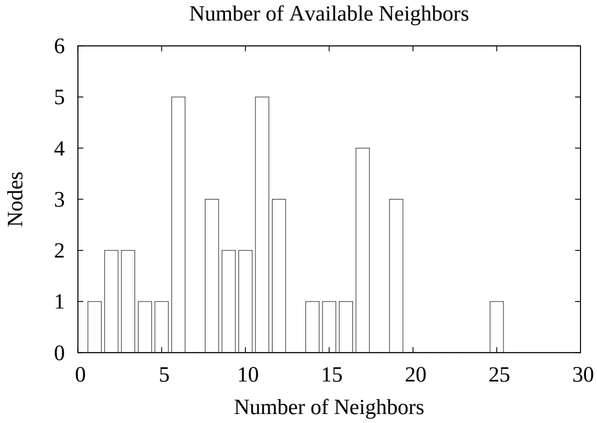

Figure 13: Histogram of the number of neighbors per node. A node counts as a "neighbor" if it has greater than 40% delivery ratio for 1 megabit per second packets.

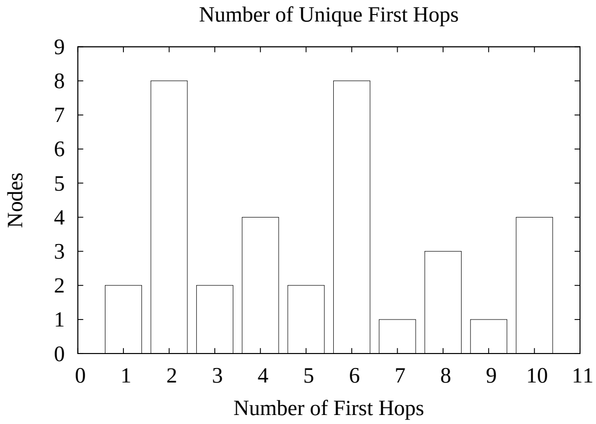

Figure 14: Histogram of the number of different first hops that Roofnet nodes use in all-pairs routes.

that the tallest buildings are located around the perimeter.

3.5.2 Random choice

It is likely that a community network would not be able to choose the best nodes to act as wired access points, but would choose them more or less randomly. Figure 17 shows performance random access point and gateway choice. For each number of access points or gateways, the set was chosen by considering 250 distinct sets of that size and choosing the set with the median average throughput. This was done to select a representative set.

25 access points would be required to cover all Roofnet nodes, many more than with optimal choice. 90% of the nodes are covered with 10-13 access points, but there are a few nodes which are difficult to reach: the histogram in Figure 13 shows these last ten percent of nodes are within the range of three or fewer neighboring nodes.

3.6 Interference

Figure 4 shows that the average throughput in Roofnet falls off much faster with each additional forwarding hop than one might expect. The average single-hop, direct link used in Roofnet provides 316 kilobytes/second. If each link in a two-hop route offered such performance, one would expect the average two-hop throughput to be roughly 158 kilobytes/second. Instead, the average two-hop Roofnet route delivers 97.7 kilobytes/second.

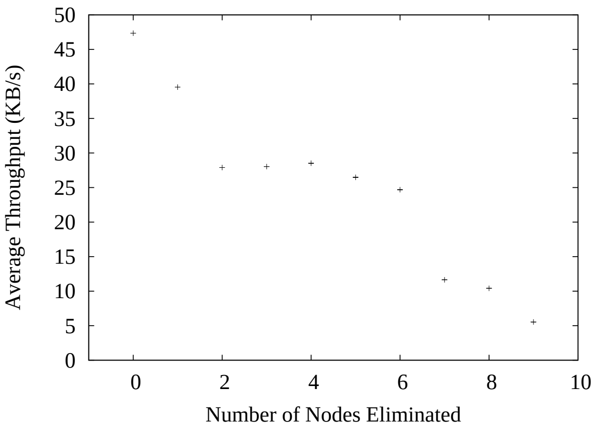

Figure 15: The effect on throughput of eliminating the best-connected Roofnet nodes. The x axis shows the cumulative number of the best nodes eliminated. The y axis shows the average throughput among the four edge nodes in Figure 11.

APs or	Multi-Hop		AccessPoint	
Gateways	Cntd	Throughput	Cntd	Throughput
1	41	119.00 KB/s	25	20.47 KB/s
2	41	202.08 KB/s	34	86.52 KB/s
3	41	235.08 KB/s	38	108.07 KB/s
4	41	261.87 KB/s	40	143.05 KB/s
5	41	255.50 KB/s	41	144.86 KB/s
6	41	273.47 KB/s	41	201.75 KB/s
7	41	287.06 KB/s	41	232.84 KB/s

Figure 16: Comparison of mesh and access-point architectures with optimal choice of gateways and access points. The *Cntd* column shows the number of connected nodes (nodes with non-zero non-zero throughput) to a gateway or access point.

Figure 18 shows this phenomenon more clearly. Every point on the graph represents one node pair. The y-value of the point shows the measured throughput between that pair of nodes. The x-value shows the throughput predicted by inverting the packets-per-second throughput of each link in the route to find the average time per packet, summing these times along the route, inverting again to predict packets-per-second for the whole route, and multiplying by the TCP packet size of 1500 bytes. One would expect every point to fall on or near the $y = x$ line, since the time to deliver a packet over a route should be the sum of the times for each link. Some variation could be expected because the component measurements are taken at different times.

In fact most of the points in Figure 18 fall significantly below the $y = x$ line, meaning that throughputs are lower than predicted. The most obvious explanation for the low performance of multi-hop routes is that concurrent transmissions

APs or Gateways	Multi-Hop Cntd	Multi-Hop Throughput	AccessPoint Cntd	AccessPoint Throughput
1	41	104.03 KB/s	6	90.98 KB/s
2	41	99.94 KB/s	17	65.61 KB/s
3	41	147.06 KB/s	23	105.39 KB/s
4	41	147.06 KB/s	26	105.98 KB/s
5	41	157.48 KB/s	35	117.57 KB/s
6	41	156.06 KB/s	35	117.90 KB/s
7	41	177.52 KB/s	36	132.33 KB/s
8	41	208.43 KB/s	36	163.00 KB/s
9	41	216.60 KB/s	36	174.03 KB/s
10	41	223.92 KB/s	36	184.18 KB/s
⋮	⋮	⋮	⋮	⋮
15	41	260.12 KB/s	39	210.31 KB/s
20	41	291.65 KB/s	40	258.51 KB/s
25	41	304.29 KB/s	41	270.97 KB/s

Figure 17: Simulation of access point infrastructure throughput with random gateway choice. The *Cntd* column shows the number of connected nodes where the node has a non-zero throughput to a gateway or access point.

on different hops collide and cause packet loss.

The RTS/CTS mechanism provided by 802.11b is intended to prevent such collisions. Figure 19 shows the results of Roofnet throughput measurements with and without RTS/CTS, taken between a random subset of node pairs. RTS/CTS does not seem to improve performance. This is consistent with existing observations [17].

As a final test to see whether collisions are reducing throughput, we measured throughputs while intentionally inserting delays between each packet at the sender. The goal was to delay long enough that each packet is forwarded to the end of the route before the next packet starts. This technique applied to two selected two-hop routes increased throughputs from 70 to 107 and 70 to 125 kilobytes/second, respectively. These results suggest that collisions are a likely cause of the lower-than-expected throughputs.

4. RELATED WORK

Many of the basic ideas in wireless mesh networking were first developed for the DARPA Packet Radio Network [9]. Roofnet uses derivatives of routing protocols from prior research in mobile ad-hoc networking; RNR is loosely based on DSR [8]. A number of research groups maintain wireless testbeds which to evaluate real-world performance of MANET protocols [13, 12, 14, 11, 3].

A number of community wireless mesh network efforts exist, such as Seattle Wireless, the San Francisco BAWUG, Madrid Wireless, the Southampton Open Wireless Network, and Wireless Leiden [16]. Many of these mesh nets use directional antennas and the OSPF routing protocol. This approach works well for relatively sparse networks, whereas Roofnet is targeting future networks where networks are dense. Wireless Leiden [16] discusses various awkward aspects of directional antennas, including the effort required to plan, install and periodically adjust the antennas.

Commercial mesh Internet access services and technologies exist, such as Ricochet [15], Tropos Networks, and MeshNetworks Inc.

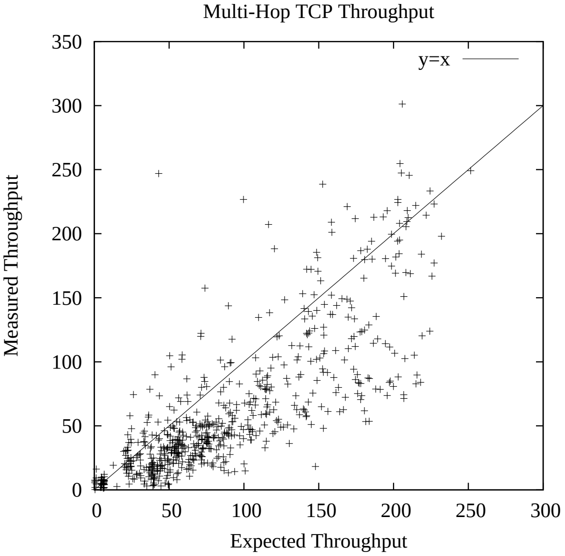

Figure 18: Comparison of measured multi-hop throughput and throughput predicted from component single-hop throughputs. Each point represents one node pair. The y value shows the measured TCP throughput between that pair in kilobytes/second. The x value shows the predicted throughput obtained by combining the single-hop throughputs of the links in the route between the pair. Only node pairs with multi-hop routes are included.

Hops	No RTS/CTS count	No RTS/CTS tput	With RTS/CTS count	With RTS/CTS tput
1	6	228.18	4	166.37
2	9	81.85	11	75.67
3	16	40.91	14	42.28
4	4	40.01	4	36.07
5	3	20.68	4	25.08

Figure 19: TCP throughputs with and without RTS/CTS.

Sensor networks use multi-hop wireless networks to collect data, and thus face problems similar to Roofnet. For example, Yarvis et al. [18] observe that hop-count performs poorly as a routing metric for a sensor network, and present the results of using a loss-aware metric instead. Ganesan et al. [7] describe the behavior of a wireless sensor net, including some unexpected phenomena similar to those seen in Roofnet.

5. CONCLUSIONS

This paper presents the design and evaluates the performance of an urban rooftop 802.11b mesh network. The architecture of the network is centered around design choices that make it easy to deploy: omnidirectional antennas, self-configuring software, and link-quality-aware multi-hop routing. Through the participation of volunteers, the network has grown to more than 40 nodes in the course of a year, with a minimum of administrative or installation effort on the part of the researchers.

An evaluation of network performance shows that Roofnet works well: average throughput between nodes is around 88 kilobytes/second, and the entire network is well served by just a few Internet gateways whose position is determined by convenience. Besides high performance, the high degree of connectedness provided by omni-directional antennas makes Roofnet more robust.

This paper also shows the effect of density on throughput: as the network becomes more dense, throughput increases. Finally, a comparison with hypothetical single-hop, access point networks shows that the promised benefits of mesh networking really do exist: multi-hop forwarding broadens coverage and increases throughput.

6. REFERENCES

[1] Daniel Aguayo, John Bicket, Sanjit Biswas, Glenn Judd, and Robert Morris. A measurement study of a rooftop 802.11b mesh network. In *Proc. ACM SIGCOMM Conference (SIGCOMM 2004)*, September 2004.

[2] Bernard J. Bennington and Charles R. Bartel. Wireless Andrew: Experience building a high speed, campus-wide wireless data network. In *Proc. ACM/IEEE MobiCom*, pages 55–65, September 1997.

[3] Kwan-Wu Chin, John Judge, Aidan Williams, and Roger Kermode. Implementation experience with MANET routing protocols. *ACM SIGCOMM Computer Communications Review*, 32(5), November 2002.

[4] IEEE Computer Society LAN MAN Standards Committee. *Air Interface for Fixed Broadband Wireless Access Systems*. New York, New York, 2002. IEEE Std. 802.16–2001.

[5] Douglas S. J. De Couto, Daniel Aguayo, John Bicket, and Robert Morris. A high-throughput path metric for multi-hop wireless routing. In *Proceedings of the 9th ACM International Conference on Mobile Computing and Networking (MobiCom '03)*, San Diego, California, September 2003.

[6] R. Draves, J. Padhye, and B. Zill. Comparison of routing metrics for static multi-hop wireless networks. In *Proc. ACM SIGCOMM Conference (SIGCOMM 2004)*, September 2004.

[7] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical report UCLA/CSD-TR 02-0013, UCLA CS Department, 2002.

[8] David B. Johnson. Routing in ad hoc networks of mobile hosts. In *Proc. of the IEEE Workshop on Mobile Computing Systems and Applications*, pages 158–163, December 1994.

[9] John Jubin and Janet D. Tornow. The DARPA packet radio network protocols. *Proceedings of the IEEE*, 75(1), January 1987.

[10] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M. Frans Kaashoek. The Click modular router. *ACM Transactions on Computer Systems*, 18(4), November 2000.

[11] Henrik Lundgren, Erik Nordstrom, and Christian Tschudin. Coping with communication gray zones in IEEE 802.11b based ad hoc networks. In *ACM WoWMoM Workshop*, September 2002.

[12] David A. Maltz, Josh Broch, and David B. Johnson. Experiences designing and building a multi-hop wireless ad hoc network testbed. CMU-CS-99-116, Carnegie Mellon University, School of Computer Science, March 1999.

[13] David A. Maltz, Josh Broch, and David B. Johnson. Quantitative lessons from a full-scale multi-hop wireless ad hoc network testbed. In *Proceedings of the IEEE Wireless Communications and Networking Conference*, September 2000.

[14] Eric Nordstrom. APE - a large scale ad hoc network testbed for reproducible performance tests. Master's thesis, Uppsala University, June 2002.

[15] M. Ritter, R. Friday, R. Garces, W. San Filippo, and C-T. Nguyen. Mobile connectivity protocols and throughput measurements in the Ricochet microcellular data network (MCDN) system. In *Proc. ACM/IEEE MobiCom*, July 2001.

[16] Rudi van Drunen, Jasper Koolhaas, Huub Schuurmans, and Marten Vijn. Building a wireless community network in the Netherlands. In *USENIX/Freenix Conference*, June 2003.

[17] Kaixin Xu, Mario Gerla, and Sang Bae. Effectiveness of RTS/CTS handshake in IEEE 802.11 based ad hoc networks. *Ad Hoc Network Journal*, 1(1).

[18] Mark Yarvis, W. Steven Conner, Lakshman Krishnamurthy, Jasmeet Chhabra, Brent Elliott, and Alan Mainwaring. Real-world experiences with an interactive ad hoc sensor network. In *Proceedings of the International Workshop on Ad Hoc Networking*, August 2002.