ESD342 Roofnet Report

Derek Rayside, Jennifer Underwood, Yingxia Yang

May 2006

Annotated Table of Contents

- RoofNet: Introduction to the RoofNet Wireless Mesh Network

  Provides a general overview of the RoofNet network and describes the system in terms of: the stimulus, main actors, and stakeholders; the sources of needs and requirements; the system boundary; the class project mission statements; and a brief discussion of the historical background and evolution of the RoofNet network.

- RawData: Detailed Overview of the RoofNet Data

  We analyzed the RoofNet network architecture using data from their ● SIGCOMM'04 paper. This section discensses the data available.

- GeographicalMaps: Maps of RoofNet and Identification of Gateways

  The geography of the RootNet network is an interesting aspect of the system since RoofNet is an un-planned network. We considered the possibility that the geographical locations of the nodes and Gateways were important to the performance of the network. We visually identify the location of the nodes and Gateways here.

- Effect of Increasing Attempted Data Rates

  The RoofNet data is separated into 4 distinct experiments, representing four different attempted data rates. We undertook analysis to understand the effect of increasing the attempted bit rates on the network topology. We find that varying the attempted data rates seems to change the network topology in largely predictable ways.

- ContrastingTopologies

  RootNet is a ' mesh' network. We compare it with a random graph model, and with two models corresponding to conventional wired computer networks. Based on the numerical metrics results, we found the architecture of RoofNet exists somewhere between a random graph and the benchmarking models. We hypothesize some explanations for these observations based on the mechanisms behind mesb networking.

- OperatorDiagnostics

  The network administrator in charge of a RoofNet installation will need to make incremental ' repairs" and "improvements" to the network. These diagnostics guide the operator towards where problems may be, and which fixes may have the best return on investment.

- CambridgePublicInternet

  The city of Cambridge is going to deploy a beta-test wireless mesh network based on the RoofNet routing protocols in the summer of 2006.

- ReflectionsAndComparisons

last edited 2006-05-16 19:45:53 by vpn-eighty-six

RoofNet

1. General Overview

This is an exploratory project that reflects the diversity of its objectives:

- To study routing protocols per Jennifer's original proposal for the class project (greatly limited by the data available).
- To apply the tools/methods of the class to the project and to analyze the system from an architectural point of view.
- To solve problems identified as concerns by the RoofNet team.
- To find out what's going on with the continued deployment of RoofNet in Cambridge.

These objectives converge and diverge in different ways.

2. System Description

We are studying "RoofNet". RoofNet is a somewhat ambiguous term, with several meanings.

The Many Meanings of RoofNet:

- A research group at MIT http://pdos.csail.mit.edu/roofnet/doku.php.
- The research produced by that group (eg, routing protocols, analysis of link quality, etc).
- A test deployment of that research by the research group in the area around MIT.
- A test deployment of that research by the research group in Tent City in Boston.
- A test deployment of that research by the city of Cambridge.

Wireless Mesh Networking:

The RoofNet deployment is a wireless mesh network. A conventional wireless network is bipartite: there are access points (gateways) and clients, and clients only connect to access points. A mesh network has two classes of nodes, but is not bipartite: the clients can all connect to each other, and route each other's packets towards their ultimate destination (which is probably a gateway). This has a number of practical advantages:

1. fewer gateways are required
2. the clients can be further from the gateways (as long as there are other clients in between to relay their packets towards the gateway)
3. the network is more robust, because each client probably has multiple viable paths to route its data on (in a conventional network, each client has only one path: a single link to its nearest gateway)

Brief Contrast with other Mesh Networks

There are a number of other wireless mesh networks. Tropos is the dominant company in the area. RoofNet is distinguished from the most other wirless mesh technologies in a number of ways:

1. Commercial mesh technologies use more conventional routing algorithms. The purpose of the RoofNet project is to explore novel routing algorithms.
2. RoofNet client nodes are entirely self-configuring. Many commercial mesh technologies require a technician to configure nodes when they are deployed.
3. Getting signal inside buildings is a major challenge for wireless mesh networks. Most commercial implementations suggest hanging more nodes on more telephone poles. The first incarnation of RoofNet technology solved this problem by running an ethernet cable from the roof into one's apartment. Landlords do not like this. The next incarnation will have two kinds of client nodes: solar powered rooftop repeaters, and small inside window-ledge repeaters. The RoofNet team is going on sabbatical to develop these window-ledge devices, and the MuniMesh team (see below) is developing the roof-top repeaters.

2.1. Stimulus, Main Actors, Stakeholders

2.1.1. RoofNet Research Team

The RoofNet research team is led by Professor Robert Morris of MIT's Computer Science and Artificial Intelligence Laboratory. Students who have worked on RoofNet include: Dan Aguayo, John Bicket, Sanjit Biswas, Ben Chambers, and Douglas De Couto.

The primary goal of the RoofNet research project is to explore novel routing algorithms for wireless mesh networks.

Most members of the RoofNet research team are going on sabbatical to a startup company named Meraki Networks, where they are working on new hardware for RoofNet nodes.

Image removed for copyright reasons.
Photo of someone holding a meraki device.

Image removed for copyright reasons.
Photo of the inside of an electrical device with the size of a playing card.

2.1.2. MuniMesh

"MuniMesh" is Kurt Keville and Bob Keyes: two "volunteer reseearchers" who are working on the technology transfer aspect of things. They work on assembling and deploying hardware for the city of Cambridge, and also on practical software engineering aspects of deploying RoofNet technology, but not on the core routing protocols. For example, they are working on making economical, physically robust, and solar-powered rooftop repeaters. The original roof-top equipment used by the RoofNet research team were regular PCs (hence relatively expensive) and required running both power and ethernet wires to the roof, which many landlords do not appreciate.

They are also writing a book on municipal wifi.

2.1.3. City of Cambridge

The main stakeholders and actors within the City of Cambridge:

- Councillor Henrietta Davis, current chair of the Cable TV, Telecommunications, and Public Utilities Committee
- Mary Hart, CIO
- Linda Turner, project manager
- Bob Coe, technical lead

Other involved parties:

- Cambridge Housing Authority
- Cambridge Health Alliance
- Museum of Science
- Harvard

2.2. Sources of Needs and Requirements

Robert Morris, the Professor leading the RoofNet project, is particularly concerned about the problem of nodes "on the periphery". Our research attempted to tackle this question in addition to analyzing the topology of the RoofNet network architecture.

2.3. System Extent (Boundary and Quantities)

The system is the RoofNet network as it existed in 2004. The boundary ends at the Gateways; there is no consideration of transition or interaction with the external WWW/internet, only the interactions internal to the wireless mesh RoofNet network. Furthermore, the aspects of the system under study are limited by the data available (please see the RoofNet Data (RawData)).

2.4. Mission Statements

Our project goals are the following:

  - Analyze the effects of increasing the attempted data rate.
  - Analyze and benchmark the network topological properties for the aggregate data.
  - Analyze the robustness of the RoofNet architecture.
  - Analyze the performance of the periphery nodes.
  - Understand the current political situation in Cambridge involving RoofNet deployment.

3. System Historical Background and Evolution

RoofNet has been deployed for several years in the Central Square area of Cambridge, MA. In the course of the experiment, the RoofNet network has grown in size. For example, in 2004, the network consisted of 38 nodes and 3 Gateways. In 2005, it had growu to 50 nodes. Please see GeographicalMaps to see this evolution of the network.

The version of RoofNet deployed in the Central Square area of Cambridge, MA, consists of PCs and roof-mounted antennas. The deployment of current and future systems is moving away from this rooftop deployment and toward "small and many", similar to the concept of sensor networks. The current implementation of this morphing strategy is encompassed in the "Tent City" project described here.

The system architectural structure does not seem to be changing; deployment remains de-centralized, each uode still functions as both a client and a router/repeater, and the only access to the external www/internet is through specified Gateway nodes.

The future of the RoofNet deployment in Cambridge is discussed in the CambridgePublicInternet section.

4. Assessment of System Effectiveness

Please refer to the following sections (also described in the Annotated Table of Contents section) that assess the system effectiveness:

- Effect of Increasing Attempted Data Rates
- ContrastingTopologies
- OperatorDiagnostics

The analyses performed in the above sections were conducted using the data discussed in the RoofNet Data Section (RawData).

5. Reflections and Comparisons

Please refer to the ReflectionsandComparisons page.

RawData

1. Sources of RoofNet Data

There are two main sources of RoofNet data publicly available. The trace data for both types of data can be found here. This section dicusses the data available and motivates the use of the 2004 SIGCOMM trace data for our project. Data from 2005 was not available.

1.1. 2004 SIGCOMM Paper

The 2004 SIGCOMM paper, *Link-level Measurements from an 802.11b Mesh Network*, can be found here. The paper focuses on analyzing the patterns and causes of packet loss in the Roofnet network. This section discusses the structure and content of the trace data used in the 2004 SIGCOMM paper.

1.1.1. Nodes

The 2004 SIGCOMM data contains specific information on each of the RoofNet nodes. This information is provided in text files that identify the IP Address (or Node ID) and geographical coordinates for each node.

- **Coordinates:** The coordinates are provided in terms of latitude and longitude.
- **IP Addresses and RoofNet IDs:** The RoofNet IDs are the unique RoofNet-specific identifiers assigned to each node. The node IDs can be found by the two low-bytes of the IP Address. For example, the building NE43 Gateway IP addresses were 5.4.102.110 and 5.5.92.100. Thus, the node IDs are 26222 and 23652.

Excerpt from the Coordinates file bottom of the page here:

IP Address	Latitude	Longitude
5.3.173.178	42.363546	-071.099826
5.4.160.160	42.360150	-071.088829
5.4.160.150	42.362881	-071.110256
5.4.168.216	42.363532	-071.099663

Interestingly, there is a separate coordinates file contained within the SIGCOMM trace file. This coordinates file lists the node ID (instead of the IP Address) and geographical coordinates.

Excerpt from the Coordinates file contained within the SIGCOMM trace data:

RoofNet Node ID	Latitude	Longitude
26206	42.365494	-71.096788
23652	42.363601	-71.09108
44466	42.361125	-71.092605

1.1.2. Traffic data

The RoofNet 2004 SIGCOMM traffic data was collected in the space of a few hours over a single night. The network was separated from the Internet to ensure that no outside traffic would contaminate the data.

The data is relatively clean in the sense that it is self-contained within the RoofNet network. For this reason, it is also relatively contrived: it does not necessarily give an accurate sense of what nominal traffic levels are like. However, it does give a sense of the connectivity and topology of the network. All of our analyses in this report are based on this topological information (ie, they are not based on studies of actual traffic patterns: we are looking at a map of the territory, not video of cars on the road).

Experiments

The data is separated into 4 distinct experiments. In a given experiment, each node takes a turn sending a series of 1500-byte broadcast packets at a specified attempted data rate. All of the other nodes listen (including the Gateway nodes). Each experiment represents a different attempted data rate (1, 2, 5.5, and 11 Mbps).

Structure of the Traffic Data

The traffic data is provided in three pieces within the 2004 SIGCOMM trace data.

- **Sent Packets:** The raw data containing information on all of the packets that were sent. The data file tracks the experiment ID (corresponding to the attempted data rate), the source Node ID (who sent the packet), a unique sequence number for each packet (locally assigned), and a time stamp for when the packet was sent.
- **Received Packets:** The raw data containing information on all of the packets that were received. The data file tracks the experiment ID, the destination Node ID (who received the packet), the unique sequence number for the packet (as assigned by the sender), a time stamp for when the packet was received, and signal and noise values as measured by the 802.11 card.
- **Summaries:** The raw data containing information on each link. This data file combines information on each source and destination node pair, the experiment ID, and provides the delivery ratio as defined by the fraction of packets sent by the source that were received by the destination node. The file also notes the signal and noise average values for all received packets.

1.2. Other Data

The trace data here includes traffic data collected over the course of several months. These traces probe traffic including packets that are passing through the Gateway nodes to the rest of the internet. Although this data represents sampling of nominal traffic patterns, it is not self-contained to the RoofNet network itself, though it would provide information as to congestion patterns. Because there is much uncontrolled data, the RoofNet team suggested we not use this data.

2. Data Inconsistencies and Issues

Although the raw data provided by the RoofNet team was carefully organized and archived, we stumbled across a few challenges to implementing our desired analysis of the network architecture.

2.1. 2004 SIGCOMM Data Inconsistencies

The first challenge involved inconsistencies within the 2004 SIGCOMM data.

2.1.1. Traffic Data Inconsistent with Coordinate Data

There were eight instances in which nodes were referenced in the traffic data as source and/or destination nodes but did not appear in the Node-ID coordinate file.

2.1.2. Coordinate Data Inconsistent with Map

There were three instances in which nodes were referenced in the coordinate file but whose latitude-longitude coordinates were not consistent with the map provided in the paper.

2.1.3. Gateway Identification

The paper only makes vague references to the location of the Gateway nodes. We felt it was important to know their location in order to understand their potential impact on the architecture.

2.2. 2004 SIGCOMM Traffic Data Issues

The manner in which the traffic data was collected limited the amount and type of analysis we could meaningfully perform. We had great interest in considering the relationship between the RoofNet architecture and how it performed in terms of congestion and routing.

2.2.1. No Global Clock Synchronization

The traffic data was not synchronized. Each RoofNet node locally estimated the time a packet was sent and received. Since the clocks at each node were not synchronized, there were multiple instances of packets arriving before they were sent if a global time were assumed.

2.2.2. No Global Unique Packet Identifiers

The packet numbers were not globally assigned. Each node locally assigned unique packet identifiers. This made tracking the route packets took through the network impossible.

3. Resolution of Inconsistencies and Issues

To resolve the inconsistencies and issues discussed in Section 2, we met with members of the Roofnet team.

3.1. Resolution of Inconsistency 2.1.1.

The RoofNet team provided us with the coordinate data for 6 of the 8 inconsistent nodes.

As for the other two nodes: At the time of the experiments, node 36879 did not have a separate roof-mounted antenna, but did share an apartment with 26206. They lost track of node 43220, but based on its local connections and an approximate idea of the geographical layout of the network at that time, I guessed its location.

3.2. Resolution of Inconsistency 2.1.2.

We were told that the origins of the map used in the 2004 SIGCOMM paper are lost to the mists of time. They told us to rely on the resolved data.

3.3. Resolution of Inconsistency 2.1.3.

We were given more specific Gateway information. The 2004 RoofNet map with Gateways highlighted is shown below:

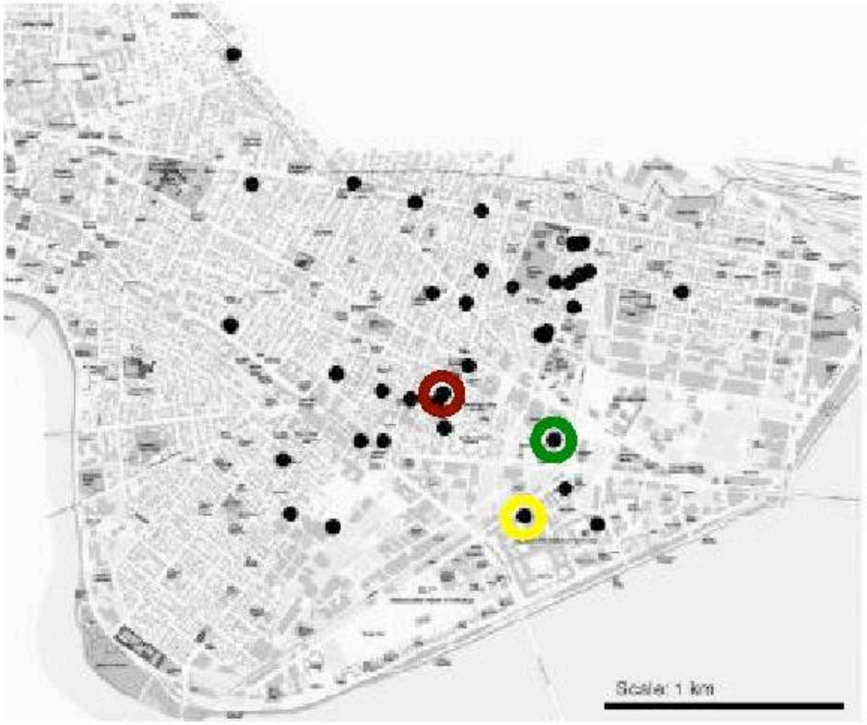

LEGEND:

- **(Green) Building NE43:** Gateway nodes 26222 and 23652
- **(Yellow) Building 36:** Gateway nodes 44466/3370
- **(Red) Cherry Street:** Gateway node 26206

3.4. Resolution of Issues 2.2.1 and 2.2.2

The packet/traffic data issues meant we had no real way of modeling congestion or routing performance. Any kind of traffic flow analysis would require some global knowledge of time. Thus, we could not perform congestion analysis using the 2004 SIGCOMM data. The non-unique packet identifiers was not an issue with the 2004 SIGCOMM data because of the manner in which the experiments were conducted.

last edited 2006-05-16 23:37:04 by 128

GeographicalMaps

The RoofNet network has grown since its inception. This makes it interesting as a dynamic system. The geography itself is also quite interesting since RoofNet is an unplanned network with de-centralized deployment.

Consider the evolution of the RoofNet network over the course of a year (2004 to 2005) in the maps below (courtesy of the 2004 SIGCOMM paper and the RoofNet webpage). It should be noted when viewing these maps that RoofNet was deployed to study and was not a commercial venture.

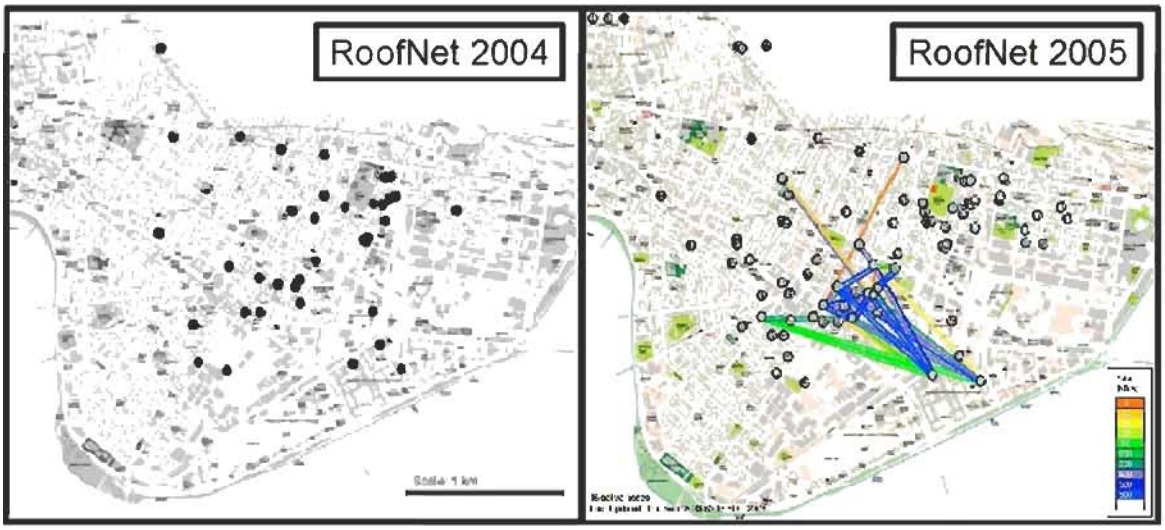

Below is a map of the 2004 Roofnet network (left, courtesy of the SIGCOMM2004 paper) as well as a map of the RoofNet coordinate data plotted in OPNET (right). The location of the Gateway nodes are highlighted in both maps. Note the differences in the map from the paper and the OPNET map using the coordinate data from the 2004 SIGCOMM raw trace data. These differences are discussed in the RoofNet Data (RawData) section of the report.

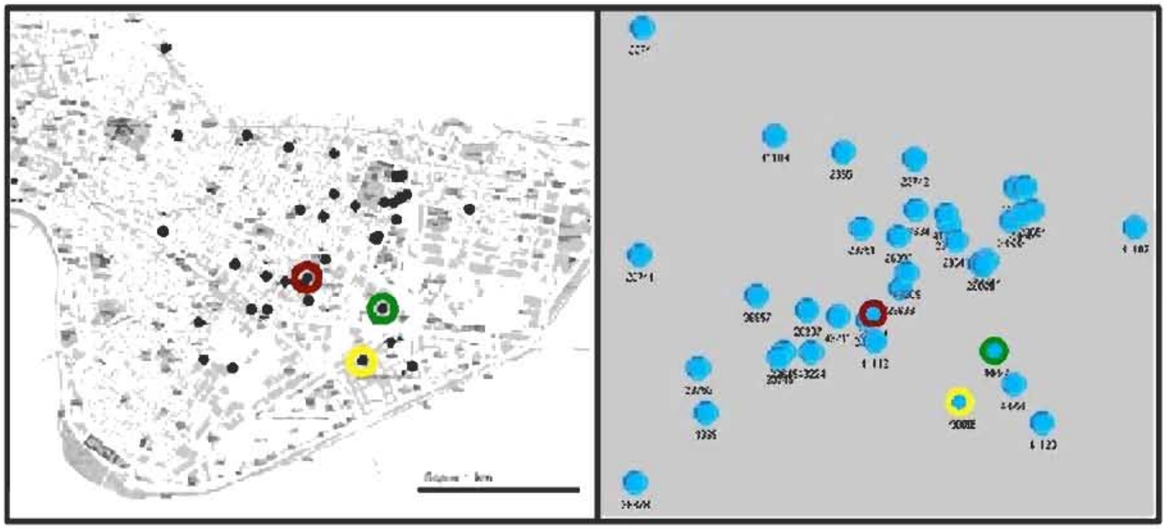

Legend:

- **(Green) Building NE43:** Gateway Nodes 26222 and 23652
- **(Yellow) Building 36:** Gateway Nodes 44466 and 3370
- **(Red) Cherry Street:** Gateway Node 26206

last edited 2006-05-16 20:10:00 by vpn-eighty-six

Effect of Increasing Attempted Data Rates

1. Effect of Increasing Attempted Data Rates

As mentioned in the RoofNet Data (RawData) section of this report, the RoofNet SIGCOMM2004 data is broken up into 4 separate experiments. In each experiment, each node attempts to send data at a specified bit rate. This protocol is in contrast to the TCP/IP protocols that adjust the bit rate real-time to compensate for congestion and poor link quality. This section discusses the analysis undertaken to understand the effect of increasing the attempted bit rates on the network topology.

1.1. Data

By Experiment: Please refer to the Roofnet Data (RawData) section of this report.

Aggregate Data: The aggregate data is a dataset constructed from the distinct experiment data for the purposes of our project. If a link between any two nodes exists at any point in time in any of the experiments, the link exists in the aggregate data. Link quality measurements are taken to be the average over all instances of the link.

1.2. Connectivity

In the class, we discussed connectivity as being a metric capturing the fraction of nodes connected in a network (lecture 6). In this section, we focus our analysis of connectivity in terms of the number of edges in the network, average degree per node, and the Maximal In-degree and Out-degree. We can gain insight into the connectedness of the network topology as a whole by comparing the connectedness as a function of attempted data rates. Later sections will explore other metrics for describing connectivity.

Not unexpectedly, we found that the connectivity of the RoofNet network varies as the attempted bit rates are increased. The connectivity maps for each experiment are shown below. The maps were generated by importing the 2004 SIGCOMM traffic data into OPNET. The reason for the "thinning out" of connectivity between the 1, 2, and 5.5 Mbps experiments is straightforward. Higher data rates require more energy to be successfully transmitted from one node to another. Obstacles, multi-path fade, distance, and atmostpheric phenomenon all affect the effective received energy of a signal. Thus, we expect fewer links as the data rate increases.

Strangely, there are two links that suddenly appear in the 11 Mbps experiment that weren't in the other experiments. These two links are circled in the 11 Mbps connectivity map below. This result is contrary to expectation given the above reasoning. However, the data was collected in a matter of a few short hours over one night. It is entirely possible that some kind of obstruction existed during the first 3 experiments that did not exist in the fourth experiment. This obstruction could be something as simple as a tree moving in the wind, a large truck temporarily parked in between the two nodes, it stopped raining, etc.

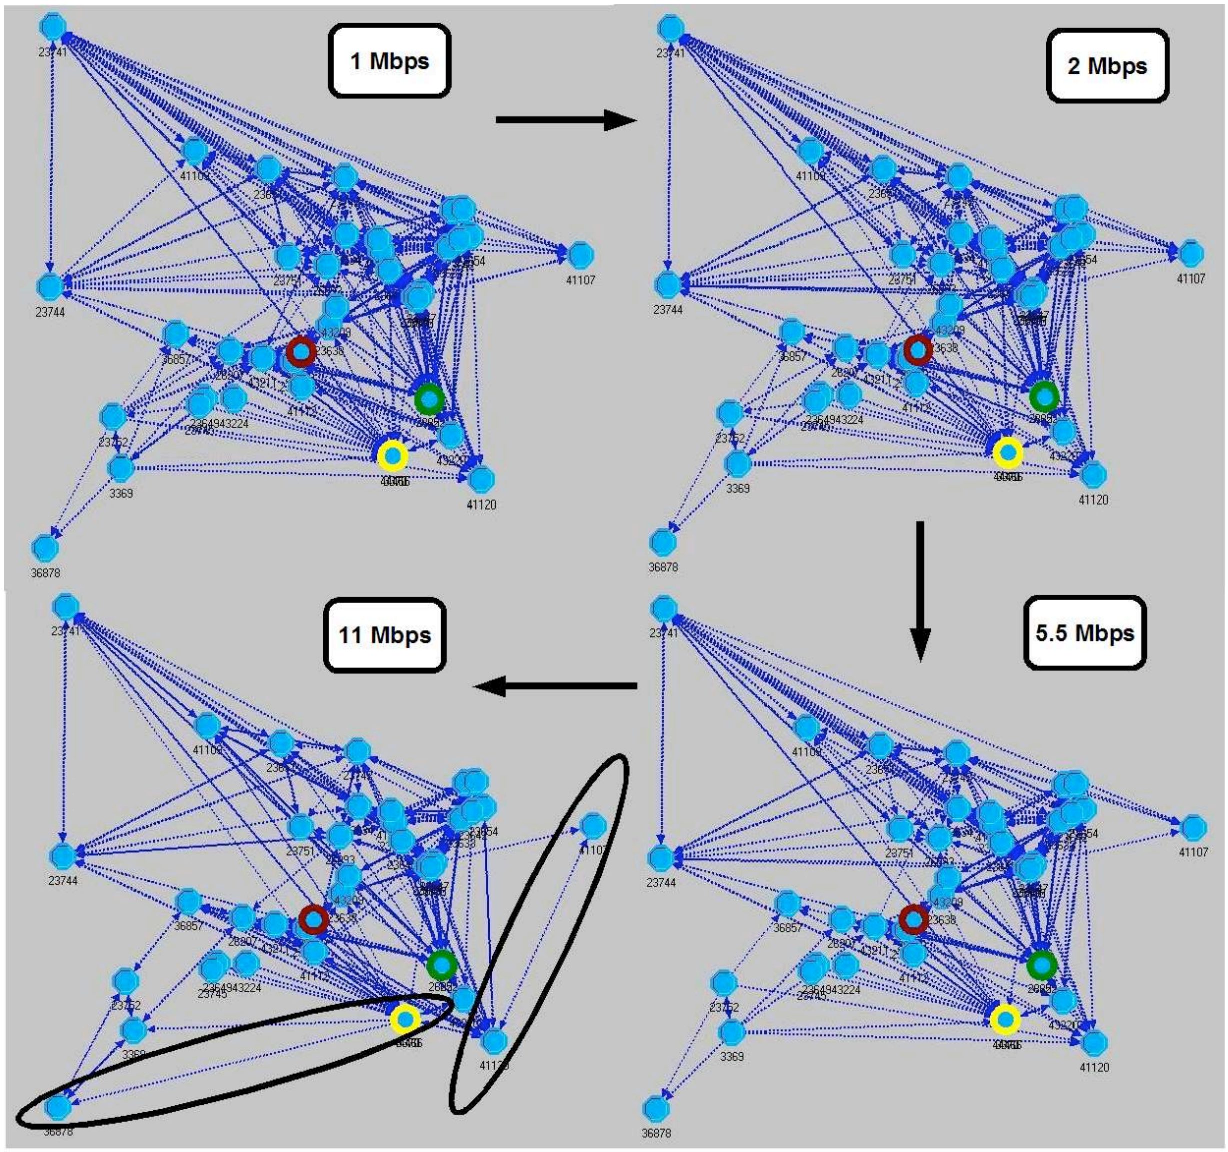

Still, the expectation that connectivity will "thin out" as data rate is increased is confirmed in the graphs below. We can see that the number of edges in the network steadily decreases as the attempted data rate is increased. Likewise, the average degree per node also steadily decreases, implying that the average number of links into and out of a given node "thins out". The differences in the Maximal In-degree and Out-degree plots imply the asymmetry of the links that is known to exist for the RoofNet network.

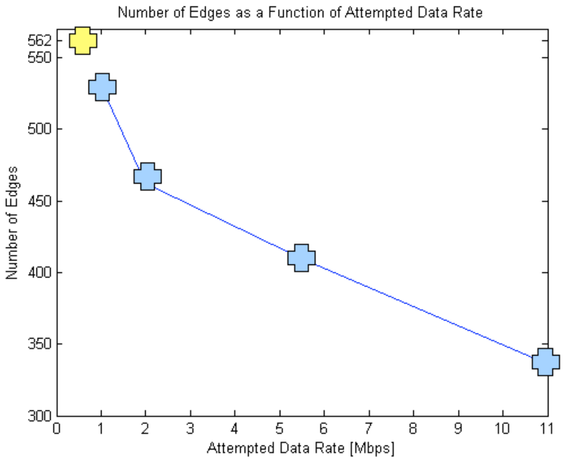

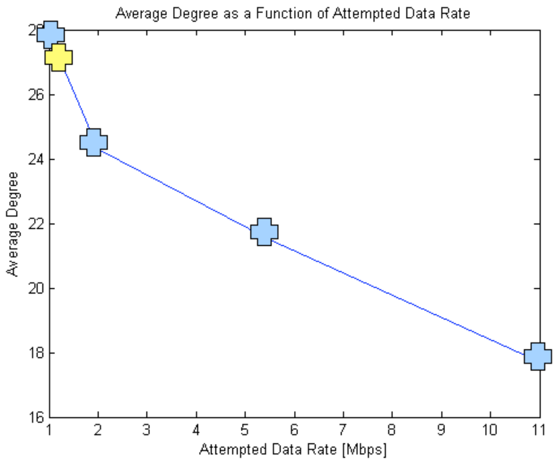

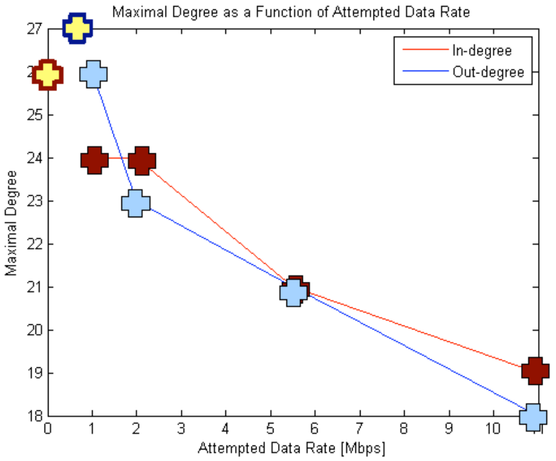

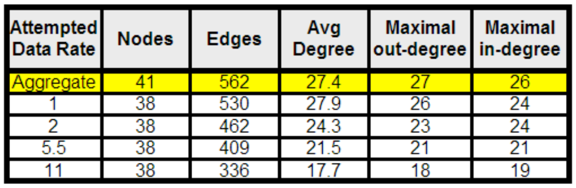

Attempted Data Rate	Nodes	Edges	Avg Degree	Maximal out-degree	Maximal in-degree
Aggregate	41	562	27.4	27	26
1	38	530	27.9	26	24
2	38	462	24.3	23	24
5.5	38	409	21.5	21	21
11	38	336	17.7	18	19

LEGEND:

- **Blue plus sign:** symbolizes the results for each of the experiments, and Maximal Out-degree in the bottom graph.
- **Red plus sign:** symbolizes the Maximal In-degree in the bottom graph.
- **Yellow plus sign:** attempts to locate the aggregate result assuming the apparent trend continues.
- **Yellow plus sign with blue trim:** Same as Yellow plus sign but for the Maximal Out-degree.
- **Yellow plus sign with red trim:** Same as Yellow plus sign but for the Maximal In-degree.

1.3. Clustering and Path Length

Related to connectivity are the ideas of clustering and path length. The clustering coefficient captures some knowledge about clusters of connectivity by evaluating the degree to which nodes linked to a common node are likely to have direct connectivity. Path length likewise captures some aspects of connectivity by measuring how far (in terms of number of hops, for example) a packet must travel between a source node and a destination node. The more connected the network, the shorter one would expect the path length to be.

The set of graphs below demonstrate the effect that increasing the attempted data rate seems to have on the clustering coefficient and the weighted and unweighted harmonic path lengths. The weighted path lengths account for the weight of each link on the basis of its delivery probability. Unweighted assumes that any link that exists has a weight of 1, thus making it analogous to weighting based on the number of hops to traverse the network.

The clustering coefficient drops significantly between 1 Mbps and 2 Mbps, and steadily decreases to 11 Mbps. Thus, as the attempted data rate increases, it becomes more and more unlikely that nodes linked to a common node have direct connectivity between themselves. The sudden drop between 1 Mbps and 2 Mbps could imply some kind of phase transition (the links dropped happened to be important ones, for example), though more targeted studies would have to be done to confirm this hypothesis. One would thus expect the average path length in terms of the number of hops (unweighted) to increase just as rapidly between 1 and 2 Mbps and start to level off after that (though steadily increasing). Sure enough, this is exactly what happens in the unweighted case.

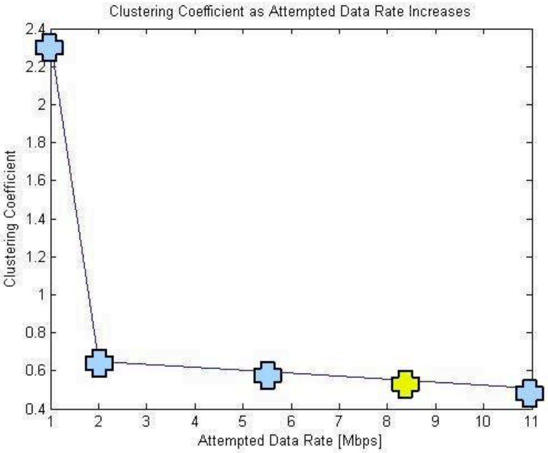

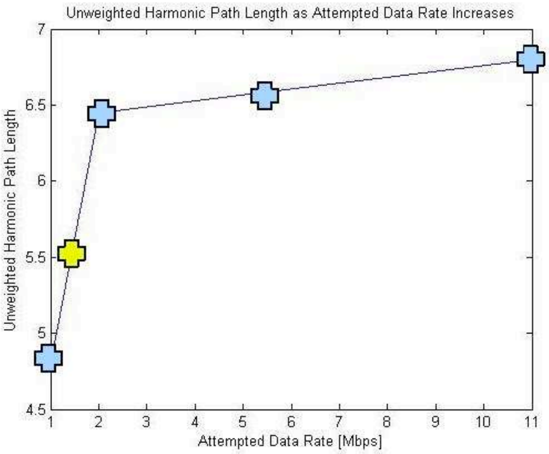

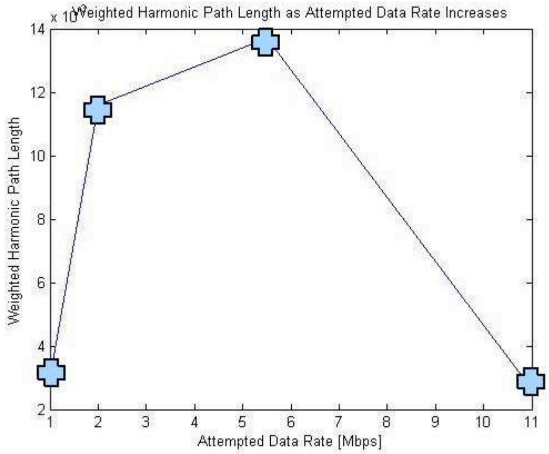

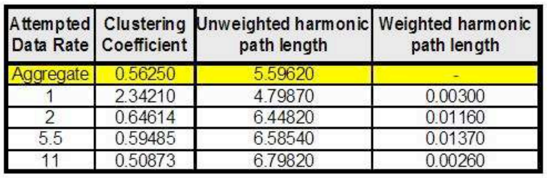

Attempted Data Rate	Clustering Coefficient	Unweighted harmonic path length	Weighted harmonic path length
Aggregate	0.56250	5.59620	-
1	2.34210	4.79870	0.00300
2	0.64614	6.44820	0.01160
5.5	0.59485	6.58540	0.01370
11	0.50873	6.79820	0.00260

The weighted harmonic path length follows this trend until the transition between 5.5 Mbps and 11 Mbps when there is a sharp drop in the path length. The only effective difference between the weighted and unweighted case is that the weighted case applies more weight to links with higher delivery probabilities. Thus, the greater the path length, the greater the probability of service should be. This would imply that there is a sharp drop in the delivery probability between 5.5 Mbps and 11 Mbps. This expectation seems to be confirmed by the data in the 2004 SIGCOMM paper (see Figure 4 below).

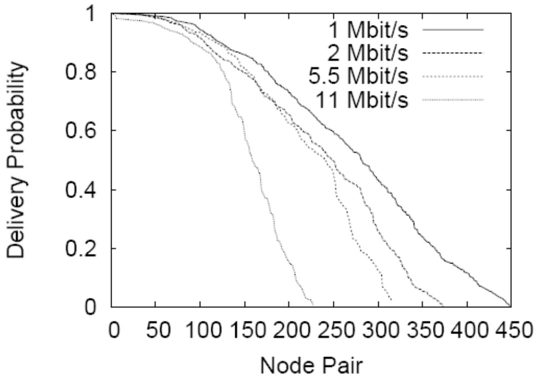

Figure 4: The distribution of link delivery probabilities for 1500-byte broadcast packets. Each point corresponds to one sender/receiver pair at a particular bit-rate. Points were restricted to pairs that managed to deliver at least one packet during the experiment. Most pairs have intermediate delivery probabilities.

1.4. Centrality

The centrality metric attempts to capture information about the amount of centralization in the network. The Degree Centrality metric defines the node that is most central as the node with the most links (Lecture 6). The Network Centralization Index (care of UCINET) measures the overall degree of centrality in the network. Ie, how much the network is controlled by nodes that are more important.

From the graphs below, it appears that the greater the attempted data rate, the more the network is controlled by more important nodes. Meanwhile, the degree centrality (both in terms of In-degree and Out-degree) decreases. This result makes sense because the more links that are dropped in the network as it "thins out" due to the increased attempted data rate, the more critical for performance certain critical paths through the network become.

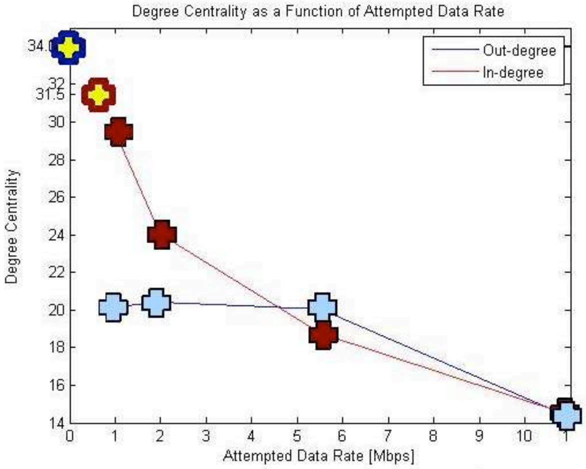

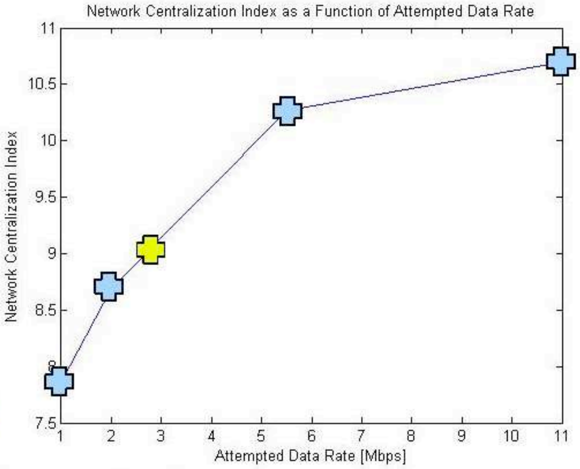

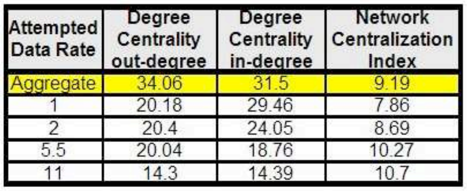

Attempted Data Rate	Degree Centrality out-degree	Degree Centrality in-degree	Network Centralization Index
Aggregate	34.06	31.5	9.19
1	20.18	29.46	7.86
2	20.4	24.05	8.69
5.5	20.04	18.76	10.27
11	14.3	14.39	10.7

1.5. Degree Distribution

The degree distribution is a histogram of the degrees of the nodes in the network. From the graphs below, it is interesting to note that the *shape* of the cumulative degree distribution hardly changes at all as the attempted data rate increases, nor are these shapes very different from cumulative degree distribution for the aggregate data. What does happen: the graph seems to shift to the left slightly and contract ("bunch" up). Could this imply some inherent structure in the RoofNet architecture? It is difficult to say given the limited data available, but it is a curiosity since so much else seems to change significantly as the attempted data rate is increased.

There seems to be a more noticeable change in the histograms themselves. As the data rate increases, the peaks of the histogram shift left, seemingly corresponding with the shifting and contracting in the cumulative distribution.

1 Mbps:

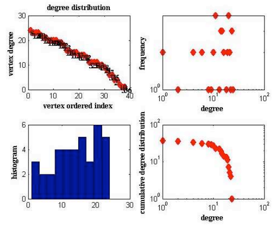

2 Mbps:

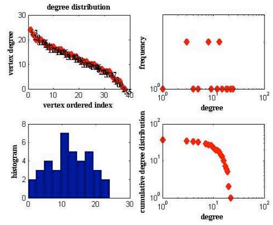

5.5 Mbps:

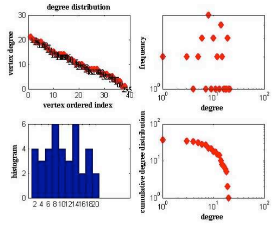

11 Mbps

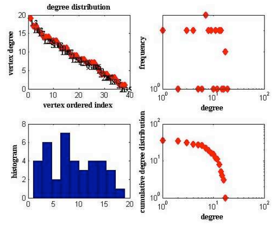

RoofNet Asymmetrical Aggregate:

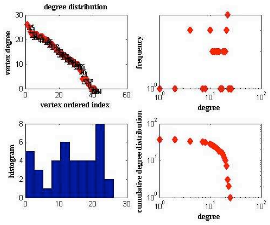

1.6. Summary

This analysis demonstrates that changing the attempted data rate in wireless mesh networks has the effect of changing the network topology. Furthermore, it seems to change the topology in largely predictable ways. Determining the extent of how this effect might be reproducible would require further analysis.

last edited 2006-05-16 20:14:45 by vpn-eighty-six

ContrastingTopologies

The purpose of the analysis in this chapter is to compare Roofnet topology with other networks with different topologies so that we can examine some -ilities of mesh networks. Two benchmarking models and 1 random graph were generated to serve this purpose. The RoofNet architecture was found to exist between the random graph and the benchmarking models. We hypothesize explanations for these observations based on the mechanisms behind mesh networking technology, which is used in Roofnet.

1. Data Process

As described in Effect of Increasing Attempted Data Rates, aggregate data was generated by aggregating the data from the four experiments contained in the 2004 SIGCOMM data. This process ignores the delivery probability differences and the differences in the bit rate; therefore, it collects every possible connection between nodes. This is elaborated as follows:

1. It is noted that the real data is very asymmetric in the sense that:
    a. The connections between two nodes are directed instead of undirected. For example, there exists only the link from node A to node B, but there is no link from B to A. The worst case example involves 3 nodes in the network that send many packets to other nodes but never receive any packets. The reasons for this phenomena are so far unclear.
    b. As stated before, delivery probability is taken to be the criteria for connectivity. It was found that some nodes can only send (or receive) with very high delivery probability, while receiving very low quality signals. This can be illustrated in the Chart below:

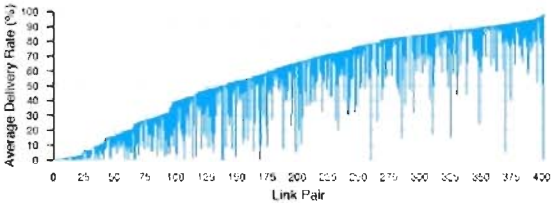

Because the purpose of this analysis for aggregate data is to study the topological differences with other networks, we symmetrized the links when we perform the numerical analysis. Therefore, the network to be compared with the random graph and the benchmarking models is a symmetric network. The network with asymmetric links is also evaluated for comparison. The graph on the left below is the directed, asymmetrical RoofNet network as plotted in UCINET. The graph on the right is the undirected, symmetric RoofNet network.

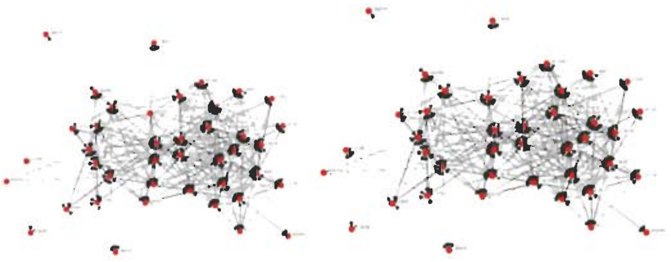

2. Because of the method to aggregate the data, there can be many links between A and B. These links incorporate topological information for the different bit rates (from the four experiments) and the different delivery probabilities over time due to weather, multi-path fade, and other signal disturbances. To serve our purpose of topological analysis, regardless of the delivery probability, if a link exists at any point in any of the experiments, the aggregate data assigns a link (value 1), otherwise it doesn't (value 0). Therefore, the network under analysis is also unweighted network.

2. Model Generation

Two models are generated for benchmarking purposes. Each model has the same number of nodes as the actual Roofnet network.

As shown below, MODEL 1 is analogous to a LAN (Local Area Network) system. There are 3 hubs in this network. Each hub is associated with a LAN and is the Gateway between the users in the LAN and the rest of the internet. The users in this network do not talk with each other. They only talk with the hub. The hubs can talk to each other.

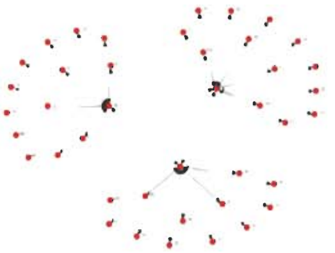

Model 2 is analogous to a WAN (Wide Area Network) system in the sense that it is a collection of LANs. Thus, each user has access to a LAN hub which has access to a WAN hub and connects to other LANs and WANs through these interconnections.

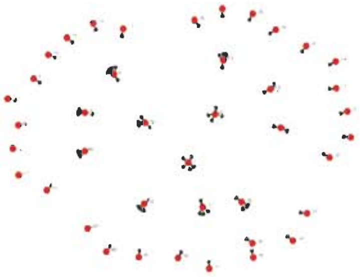

3. Contrasting with LAN/WAN Benchmarks

By using some of the analysis functions in UCINET and Gergana's MATLAB routines, the following network topological parameters were calculated:

System	n	m	k	c	L1	L2	r	Cb	Cd,
LAN (Model 1)	41	82	2	0.0007	0.6039	9.8306	-0.8623	52.34%	34.17%
WAN (Model 2)	41	82	2	0.025	0.9048	13.4575	-0.355	46.13%	7.88%
Roofnet(sym)	41	638	15.6	0.6986	0.4123	6.2269	0.0117	10.15%	32.69%
Roofnet(asym)	41	562	13.7	0.5625	0.367	5.5962	0.0633	9.19%	32.69%

- n: number of nodes
- m: number of links
- k: average degree
- c: clustering co efficiency
- L1: average path length
- L2: Harmonic path length
- r: degree correlation
- Cb: Betweenness Centrality (Network Centrality Index)
- Cd: Degree Centrality

These numerical results indicate the following:

1. The RoofNet network is a highly clustered network. The clustering coefficient, 0.6986, is much higher than for MODEL 1 and MODEL 2. This reflects the routing rules of Roofnet because every node can talk with a number of nodes nearby (each node is simultaneously a client and a router/repeater). In the traditional network, users can only talk with hubs but with each other.
2. It is not surprising that the degree correlations in MODEL 1 and MODEL 2 have negative degree correlations while Roofnet has a positive degree correlation. This again reflects the fact that the routing protocol of the wireless mesh network

doesn’t limit the users to talking only with servers/hubs. Instead, a user can be a user as well as an intermediate to transfer a packet. If a user can’t directly talk with a gateway, it can take a multi-hop path through other users to finally connect with a gateway.

3. The betweenness centralities in MODEL 1 and MODEL2 are much higher than in the Roofnet network. This can be explained by the importance of hubs in the two models. The following 3 charts about degree distribution, prestige and acquaintance can illustrate this point. All of these analytical results consistently show that Roofnet is very decentralized.

4. Contrasting with Random Graph

By using the MATLAB routine that Gergana wrote for generating random graphs, a random graph (Erdos-Renyi graph) was generated with the parameters: n=41, p=0.35, E=638 (p=0.35 is because when all nodes are connected to each other, the links would be 41*41; now there aer 638 links, so 638/41*41 = 0.35). The same numerical analysis as before was done for the random graph as well as the Roofnet network after the data was processed. The results are shown below:

System	**n**	**m**	**k**	**c**	**L1**	**L2**	**r**	**Cb**	**Cd**,
Roofnet(sym)	41	638	15.6	0.6986	0.4123	6.2269	0.0117	10.15%	32.69%
Random(sym)	41	638	15.6	0.779	0.2909	4.7243	-0.0445	0.25%	11.83%

We find that in terms of properties such as the clustering coefficient, degree correlation and degree distribution, the Roofnet network is very similar to the random network. It has no preferential attachment.

However, the betweenness centrality and degree centrality metrics of the RoofNet network are very different from the random graph. It seems that there are some important nodes with high betweenness, which makes the betweenness centrality much higher than for the random graph. Linking with the mechanism of how Roofnet works, this could be explained as follows: in Roofnet, nodes can't talk with just any of the nodes in the network (like in the random graph) because of being geographically too far from each other, so they have to link through some nodes geographically in-between. This implies that the real Roofnet network would indeed have a lot higher betweenness centrality.

degree distribution for random graph (symmetrical)

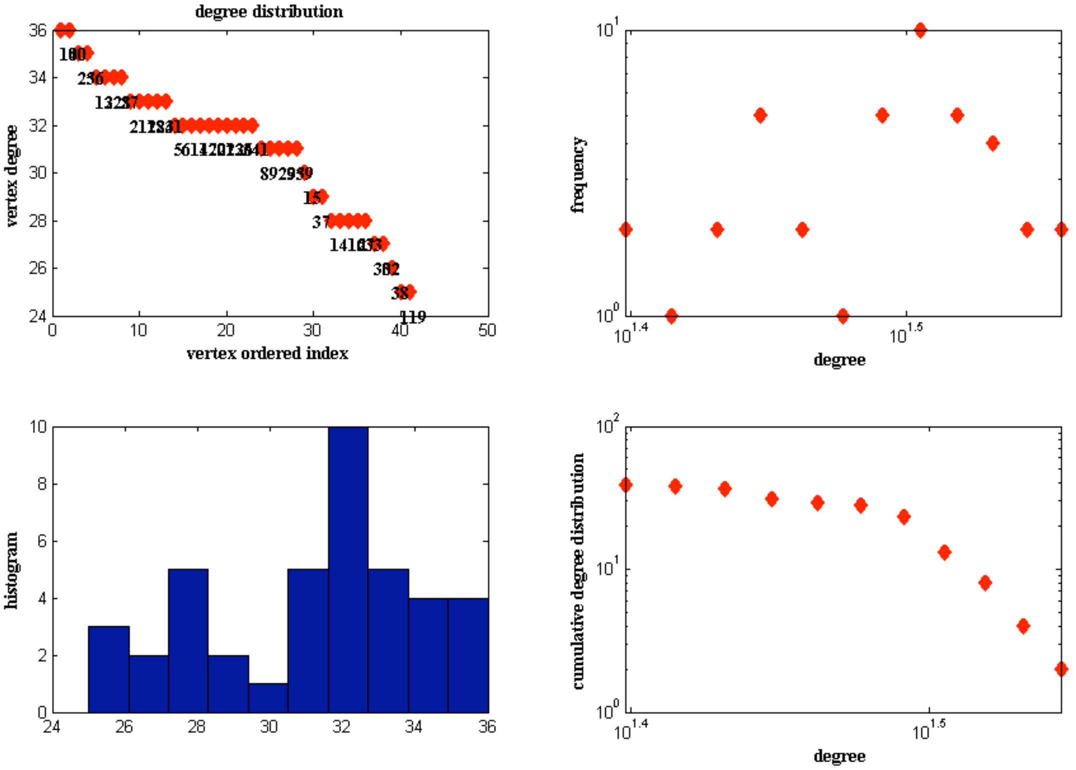

degree distribution for Roofnet (symmetrical)

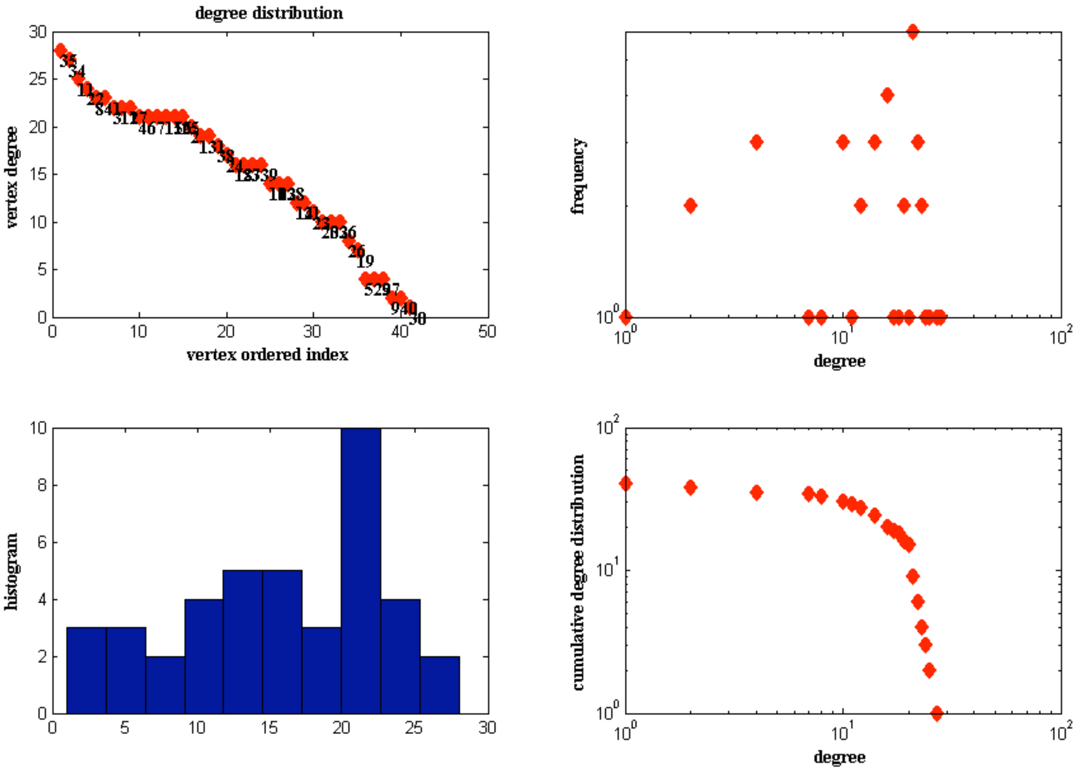

5. Robustness Analysis

Essentially, a mesh network has decentralized infrastructure, is relatively inexpensive, and is very reliable and resilient since each node need only transmit as far as the next node. Nodes act as repeaters to transmit data from nearby nodes to peers that are too far away to reach, resulting in a network that can span large distances, especially over rough or difficult terrain. * referred to Wikipedia. It would be a necessary aspect to analyze the robustness of this network by using the network analysis tools.

One way to examine the robustness of a network is by removing nodes in the network to see how resilient the whole network is. One of the criteria is the remaining number of nodes after removing nodes one by one. In ● Doyle's paper, they mentioned that the internet is a RYF (Robust Yet Fragile) system meaning that it is unaffected by random component failures but vulnerable to targeted attacks on its key components. By having two models which simulate the structure of internet, we show that the architecture of Roofnet is different from the internet: that is, robust but not fragile.

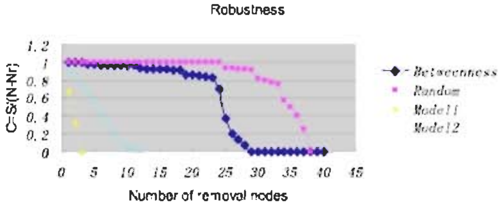

C=S/(N-Nr). S is the number of remaining nodes; N is the total number of nodes; Nr is the number of nodes that are removed.

The yellow and blue line represents removing nodes from Model 1 and Model 2 respectively by targeting the nodes with high betweenness value. We can see that the remaining nodes drop dramatically. The purple line represents removing nodes from Roofnet (asym) randomly and the blue line represents removing nodes from Roofnet(asym) in the order of betweenness. The chart clearly shows that there is no big difference between removing nodes randomly and by betweenness until the number of removed nodes is greater than or equal to 23. Even with removing nodes with the high betweenness value, it doesn't make the whole network connectivity drop suddenly like in Model 1 and Model 2. Therefore, Roofnet as a mesh network is robust but not fragile.

6. Other analysis results

6.1. Prestige

<table>
<tr><td>Model 1:</td><td>Model 2:</td></tr>
<tr><td>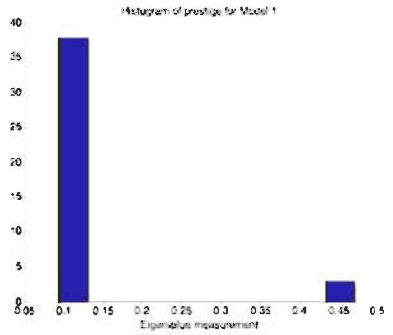
</td><td>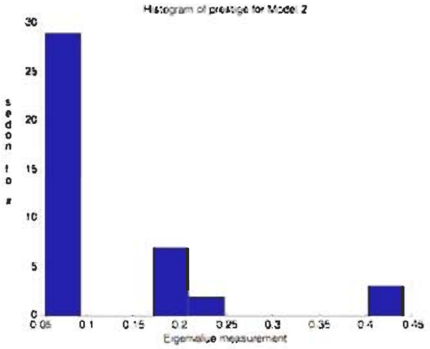
</td></tr>
<tr><td>RoofNet (asymetrical)</td><td>RoofNet (symmetrical)</td></tr>
</table>

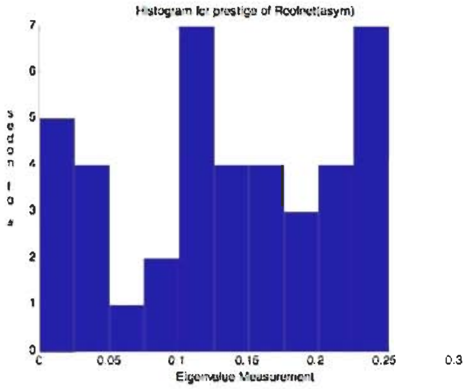

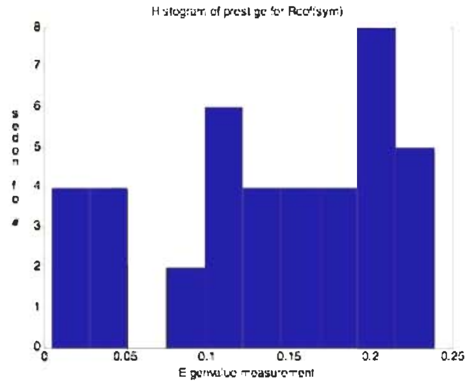

Random

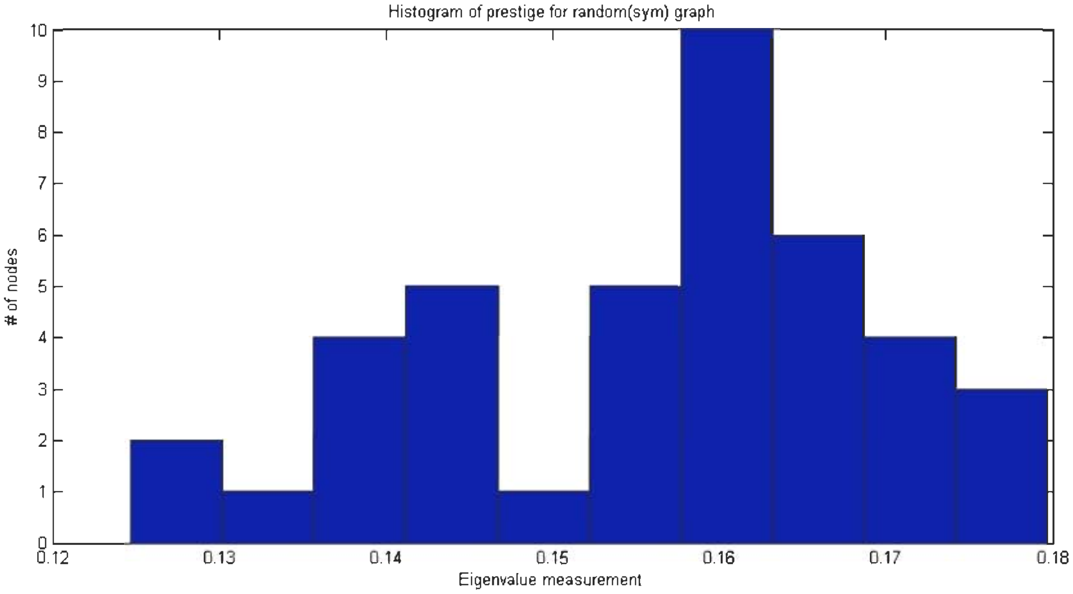

6.2. Acquaintance

Model 1:	Model 2

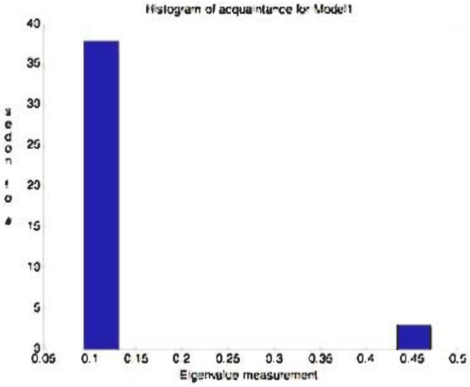

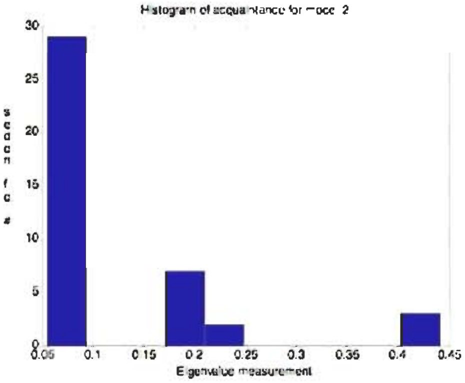

RoofNet (asymetrical)	RoofNet (symmetrical)

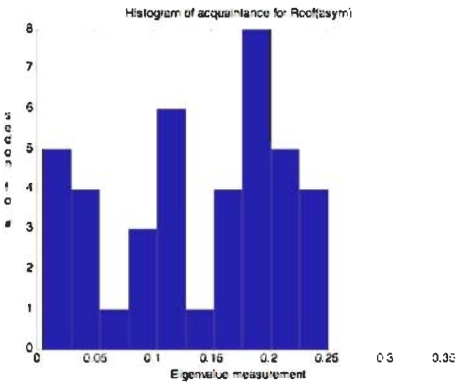

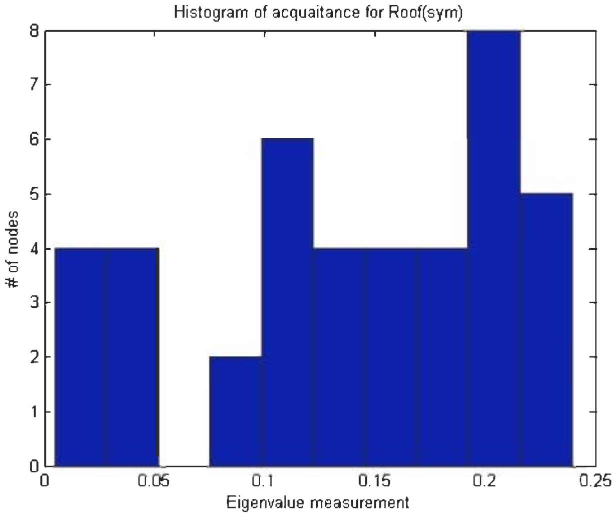

Random (symmetrical)

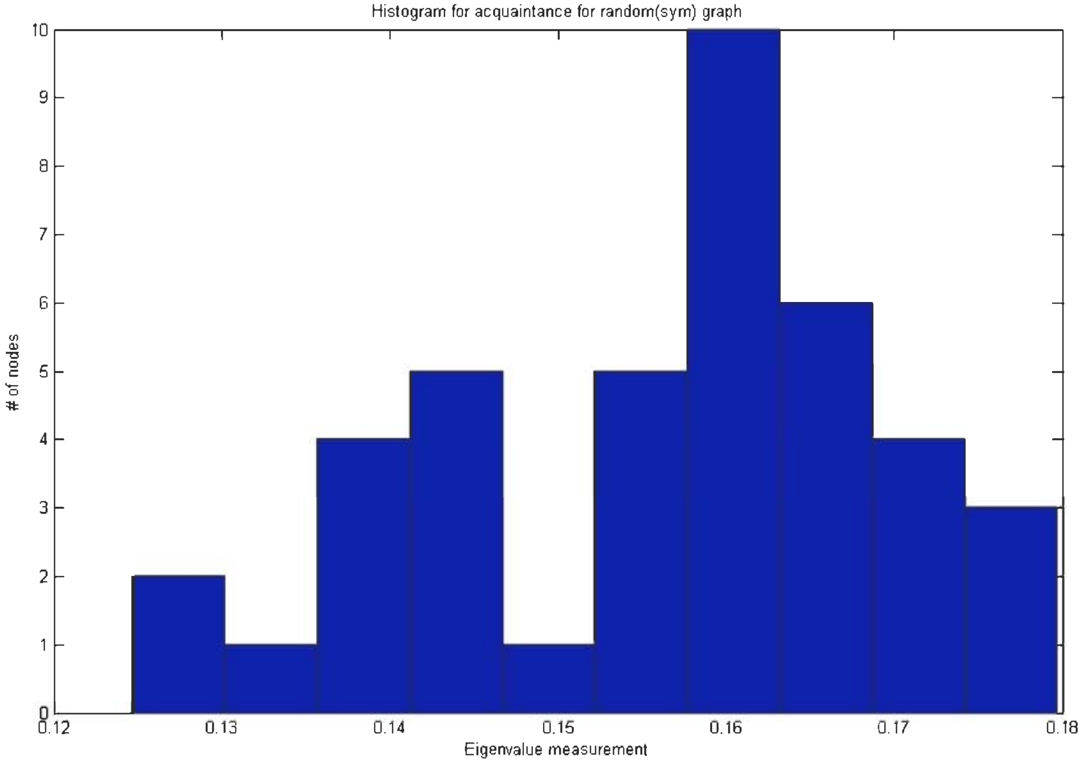

last edited 2006-05-16 20:53:04 by vpn-eighty-six

OperatorDiagnostics

In a real-world wireless mesh network deployment, some nodes may be better connected than others. The person operating the network might like to know which nodes are not well connected, and how to improve their connections.

1. Diagnostic Objectives

1. **Classify nodes according to connectivity:** gateway, isolated, single-hop, multi-hop, mid-point, periphery.
2. **Connectivity characteristics of each node:** describe the connectivity of each node at different delivery probability rates.
3. **Which link to improve?** which link should the operator focus improvement efforts on to have the greatest overall benefit for the network?
4. **Anti-Redundant Nodes:** n1 and n2 are redundant if they act as mid-points for the same (similar) set of nodes, and connect to the same (similar) set of gateways. A node is *anti-redundant* if there is no other node in the network that is structurally similar to it. More redundant nodes will increase the robustness of the network to node failure.

2. A Variation of Betweenness Centrality

Wireless mesh networks can be deployed for a variety of purposes. For example, one use is to connect police, ambulance, and fire vehicles to each other. This kind of usage is like the kinds of social networks that are traditionally studied in the networks literature: the assumption is that everyone wants to talk to everyone else, if they had the opportunity.

However, the city of Cambridge is planning to deploy RoofNet technology to provide residential internet access. In this usage, the individual nodes are not really interested in talking to each other: they just want to get to the gateway (and get information back from the gateway). Because it's a mesh network, they will talk to each other as a means to achieving the end of talking to the gateway, but talking to each other is not their objective.

To more directly model the concerns of the Cambridge deployment, we have modified some of the metrics used in class. In particular, we compute 'gateway betweenness' as a variation of 'betweenness centrality'. Gateway betweenness is the number of paths that go through a node to a gateway (the semantics of what's considered a valid path are discussed below). We also compute 'betweenness in-degree' and 'betweenness out-degree', which have natural interpretations in this domain: the number of nodes that may take a path through N to get to a gateway, and the number of gateways that N can reach, respectively.

3. Path Semantics

We consider that a path has the following characteristics:

1. from a node to a gateway
2. maximum four hops
3. each hop must have at least 5% chance of delivery success
4. path delivery probability = product of hop probabilities

Paths in the opposite direction (from gateway to node) are also a legitimate input to these analyses. However, here we just explore the analyses in the direction to the gateway. Recall as well that links in RoofNet are often asymmetrical.

This notion of path is more conservative than the one used by the RoofNet ExOR routing algorithm. ExOR explores multiple paths simultaneously, and at each hop evaluates which path is working best. The notion of path presented here corresponds more to a conventional routing algorithm that attempts to select the single best path. If this conventional notion of path identifies a good route, ExOR will also find that route. ExOR may have a higher delivery success rate in situations where no conventional good path exists. So this conventional notion of path is a conservative approximation of the expected ExOR performance.

4. A summary graph

Shows proportion of gateway, middle, periphery, and isolated nodes when paths of different quality are considered. The far left considers all existant paths, even of very low quality, and we see that there are no isolated nodes: there are 4 gateways, mostly middles, and two or three periphery. Periphery here means a node that is connected (not isolated), but which no other node is using as a mid-point.

At the far right we see that only approximately 60% of the nodes have high quality paths to a gateway (>90% delivery probability).

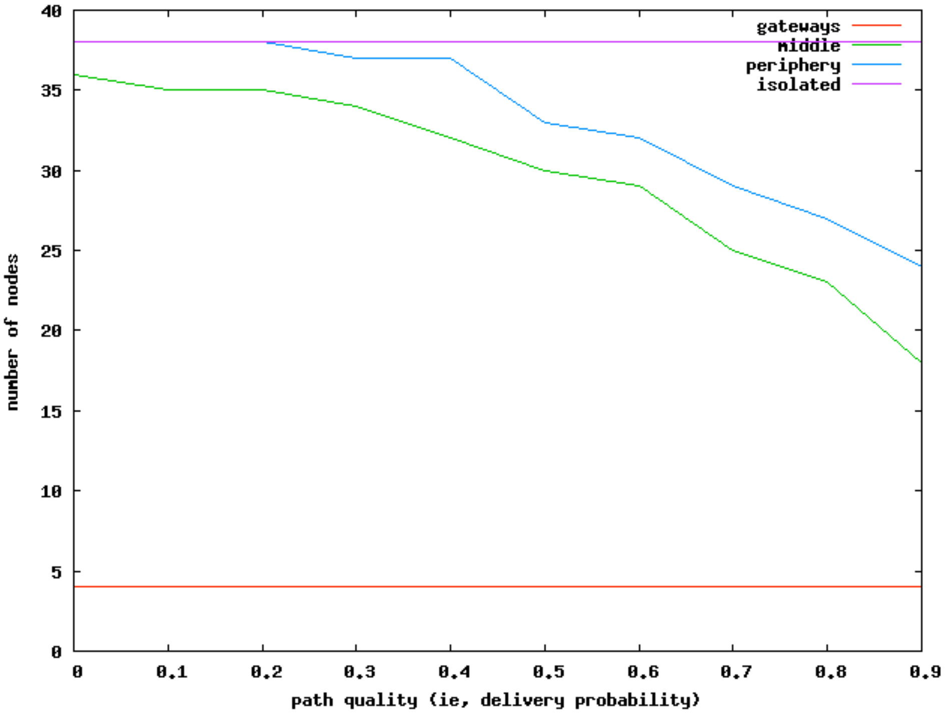

5. Classifying Nodes

Here we classify nodes according to their connectivity to a gateway:

- **Isolated** nodes are those that have difficulty finding a reasonable quality path to a gateway. "Reasonable" presently means "better than 60%".
- **Single-hop** nodes are those adjacent to the gateway
- **Multi-hop** nodes are those that can reach a gateway via some other node
- **Mid-point** nodes are those that relay packets towards a gateway on behalf of others
- **Periphery** nodes are those that do not relay packets for others, which could be for a few reasons:
  - the periphery node itself does not have a good path to the gateway
  - the periphery node is already at the maximum hop-distance to the gateway
  - other nodes do not have good paths to the periphery node

The single/multi-hop criteria and mid-point/periphery criteria are orthogonal. In other words, all four boxes in the following table are possible:

	Single-hop	Multi-hop
Mid-point		
Periphery		

The classification is performed automatically based on the charts below (the entire chart section of the report is generated by a script, and the results are pasted in here).

5.1. Legend

- x-axis: path quality (ie, delivery probability)
- green line: number of reachable gateways ("betweeness out-degree" in social network lingo)
- red line: number of nodes who can reach a gateway through this node ("betweeness in-degree" in social network lingo)
- blue line: number of immediately adjacent gateways
- magenta line: number of paths that this node is a midpoint on ("betweeness centrality" in social network lingo)
- turquoise line: number of paths this node has to a gateway

5.2. Patterns and Interpretations

- isolation: look at the green lines (how many gateways can it see?)
- importance: look at red line (how many nodes is it a midpoint for?)
- redundancy: not yet implemented.
- only a red line: gateway node
- vertical lines: stuff only works for lower path quality (to the left of the vertical line)
- no visible lines: all values are zero or one (one doesn't show up because it's a log scale; it's a log scale because the number of paths is typically much larger than the number of nodes, and these plots have both kinds of lines)

5.3. The Charts

Click on a chart to see it full size.

5.3.1. Gateways

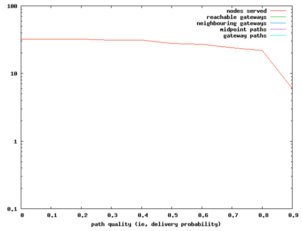

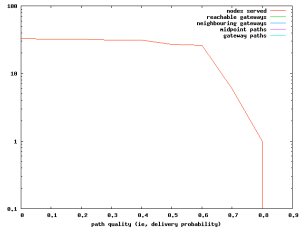

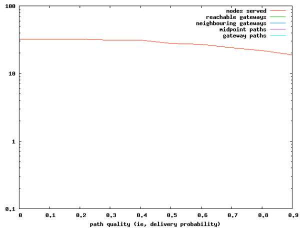

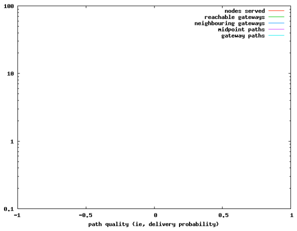

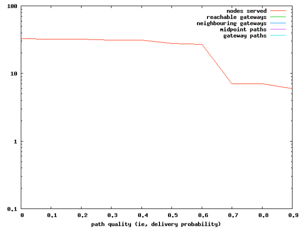

5.3.2. Isolated Nodes

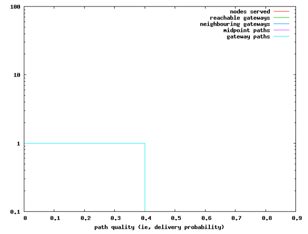

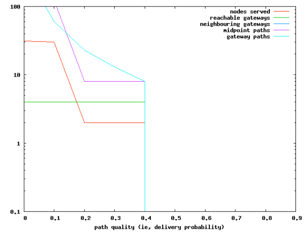

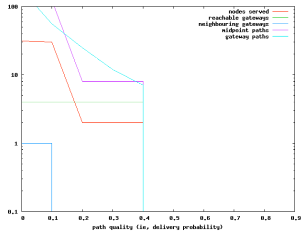

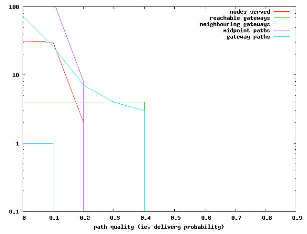

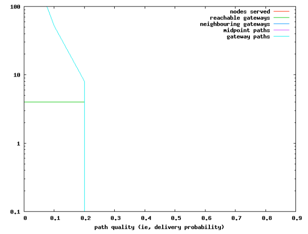

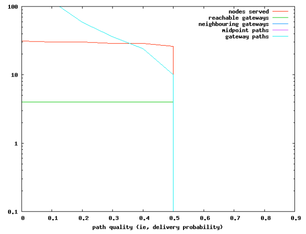

5.3.3. Single-hop Mid-point Nodes

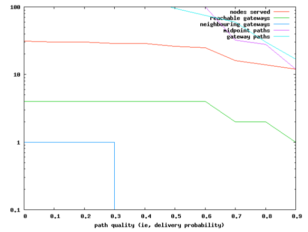

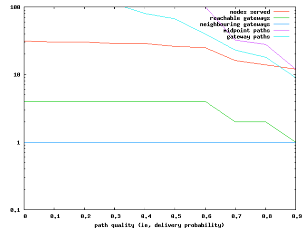

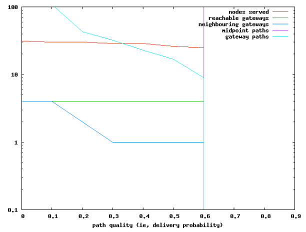

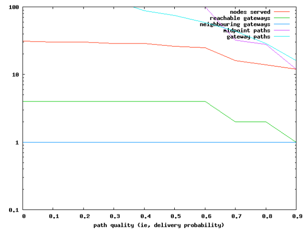

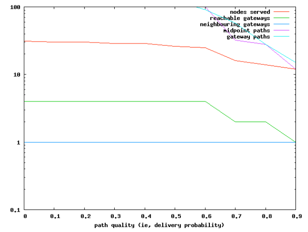

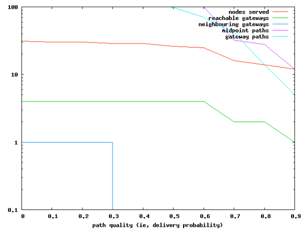

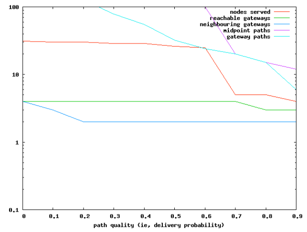

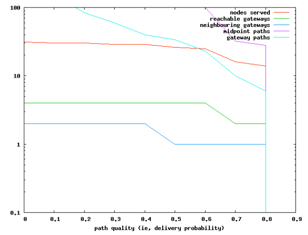

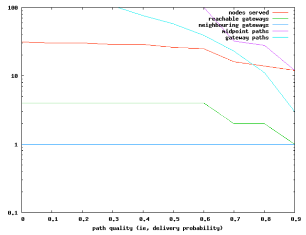

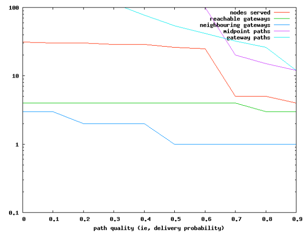

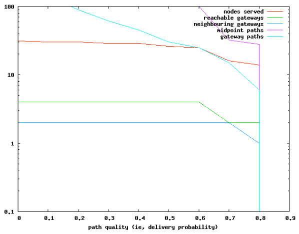

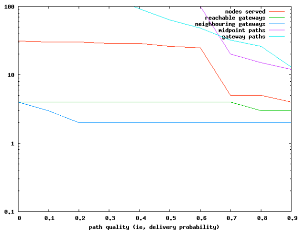

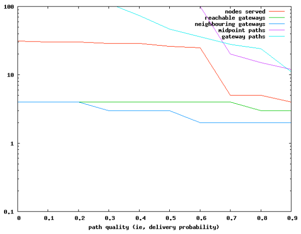

5.3.4. Single-hop Periphery Nodes

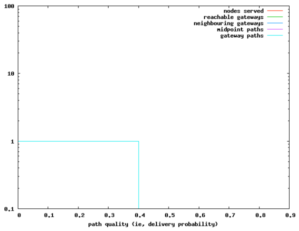

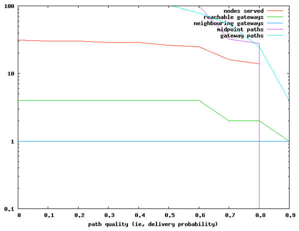

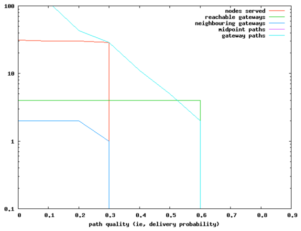

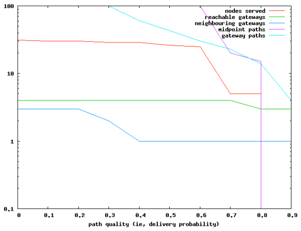

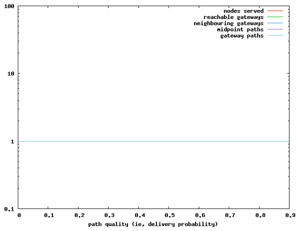

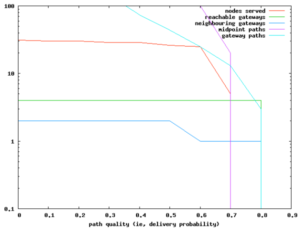

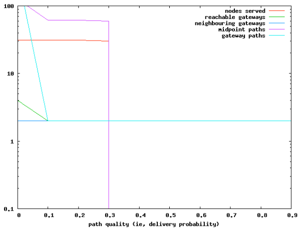

5.3.5. Multi-hop Mid-point Nodes

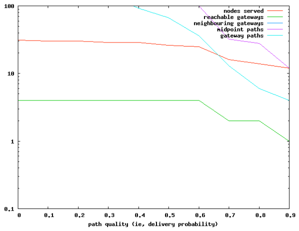

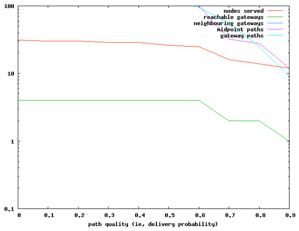

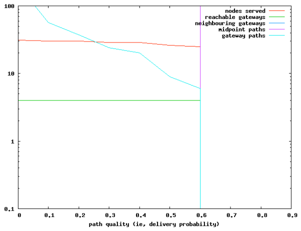

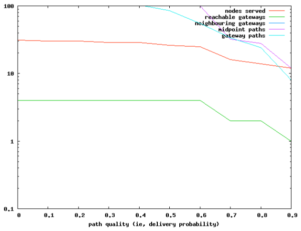

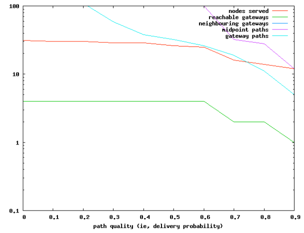

5.3.6. Multi-hop Periphery Nodes

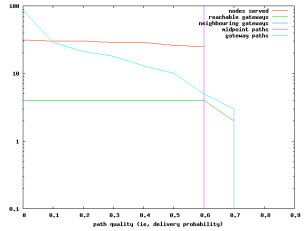

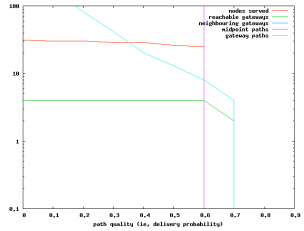

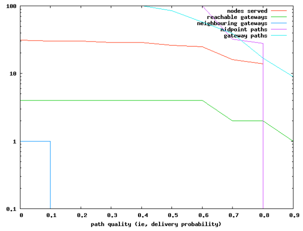

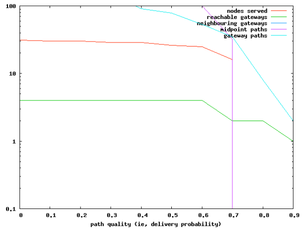

6. Improving the Mesh by Strengthening an Edge

Some parts of the network may be isolated from the gateways (ie, have only low quality paths, or no paths, to the gateways). The network connectivity may be improved by 'strengthening' an edge. The question is, which edge should be strengthened? Our analysis is: for each edge in the network with delivery probability greater than 5%, hypothetically increase its delivery probability to 99%, re-analyze the network, and determine how many previously isolated nodes have become connected. We re-analyze the network at the 90% success rate (ie, a node will become re-connected if it gains a new path of >= 90% delivery success).

We assume that brand new edges cannot be added to the network. If a faint edge already exists, then we know it is possible to communicate between that pair of nodes. Nodes that are not presently able to communicate may be too far apart, or divided by obstacles, etc. The strength of an existing edge could be improved through a number of practical strategies, such as: adding an intermediate node, directional antennae, moving physical or electro-magnetic obstacles.

For the SIGCOMM'04 data we find that 14 nodes are considered isolated at the 90% level, and that there 342 edges with delivery probability >= 5% (there are 220 edges below 5%). Improving any of the following four edges re-connects three nodes (and it happens to be the same three nodes in each case):

Edge	Strength	Reconnected Nodes
3369 -> 26207	0.45	23752 3369 36878
3369 -> 44466	0.12	23752 3369 36878
36878 -> 44466	0.13	23752 3369 36878
23752 -> 44466	0.06	23752 3369 36878

The figure below shows the geographical map of RoofNet with the re-connected nodes in the large yellow bubble in the bottom left corner. The four black lines indicate the new edges in the table above. Gateways are highlighted with red bubbles. Other isolated nodes are highlighted with smaller yellow bubbles. The figure shows that isolated nodes are not necessarily geographical outliers, while geographical outliers tend to be isolated.

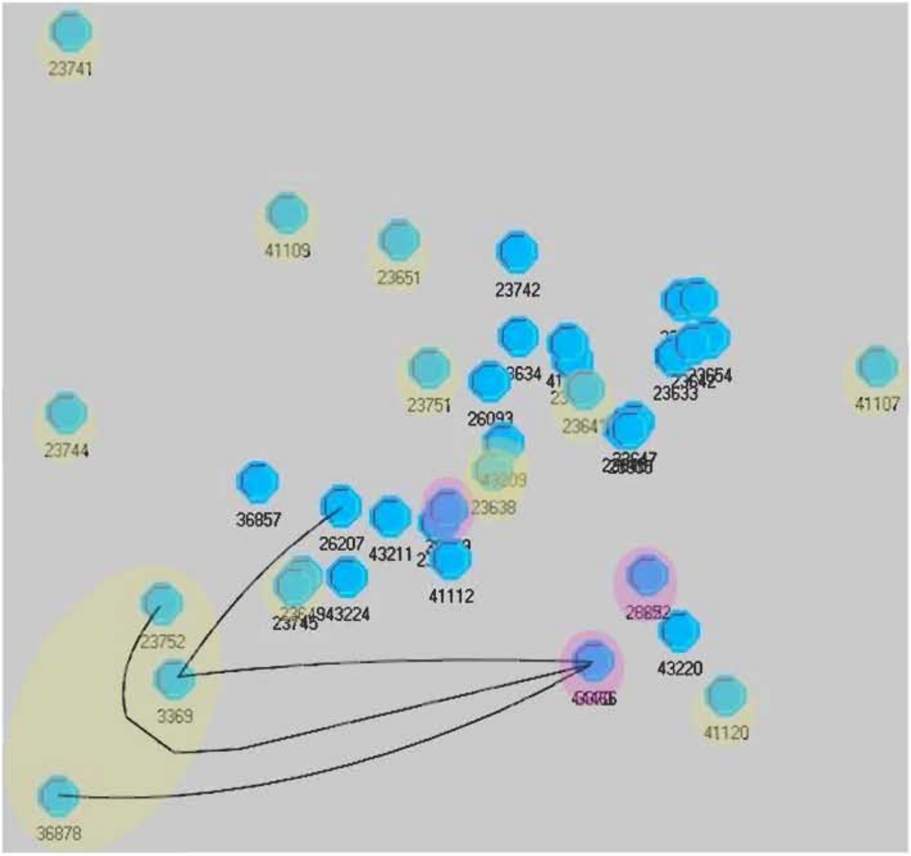

The fact that each of these four edges would re-connect the same three nodes suggests that a community-finding/clustering algorithm may also be able to identify this group.

7. Anti-Redundant Nodes

Two nodes n1 and n2 are *structurally equivalent*, or redundant, if they connect to the same set of other nodes. In a mesh network, redundancy may be viewed as a heuristic for robustness. If every node in the network has a structurally equivalent partner, then the network should be robust to node (and edge) failures. The network operator's objective, then, is to find the nodes that are the least redundant, and then add nodes to the network to make them more redundant.

There are two common ways to measure how structurally equivalent a pair of nodes are: by measuring the Euclidean distance of their relation to other nodes, or by measuring the Pearson correlation of their relation to other nodes. We have computed both, and then rank the nodes from least redundant to most redundant:

Pearson Correlation: 23649 41107 23652 23742 26093 41120 23634 23751 44466 23638 36879 43220 23741 26206 41123 43209 23641 23651 3370 23739 41112 36857 26207 43211 23642 23654 23633 23744 41109 23635 23645 23752 23647 23740 3369 36878 23734 41105 23745 26222 43224

Euclidean Distance: 23652 26093 41120 23742 3370 23741 26206 23634 41112 44466 23642 23739 43220 23654 23633 36857 26207 43211 41123 41109 23647 23740 23638 36879 23635 23645 23651 43209 23641 23744 23751 23734 41105 23752 3369 36878 23649 41107 23745 26222 43224

These two measures agree that nodes 23745, 26222, and 43224 are structurally equivalent with each other, which we have verified by hand in the original edge list. However, the position of other nodes in the list can vary substantially: for example, node 41107 is the second most anti-redundant node by the Pearson measure, yet is considered one of the most redundant nodes by the Euclidean measure. Some other nodes appear at similar places in both lists: 23652 is the most anti-redundant node by the Euclidean measure, and the third most anti-redundant node by the Pearson measure. At present it is unclear why these two measures sometimes give very different results.

It is still possible for the network to be robust without this kind of redundancy, but reducing anti-redundant nodes is one way to target the network operator's efforts to make the network more robust.

It might be more profitable to compute the minimum cut-set between each node and the gateway.

last edited 2006-05-17 12:10:21 by c-65-96-163-242

CambridgePublicInternet

1. Overview

The ● City of Cambridge is planning to deploy RoofNet wireless mesh technology to select neighbourhoods, starting in the summer of 2006. The first selected neighbourhood is the part of ● Area 4 on the other side of Portland Street from Tech Square.

The city seems to view the deployment of wireless mesh technology as a "digital divide" issue, and has a complementary digital divide project to provide free computer equipment to select residents. Considering the project this way substantially reduces the city's financial and technical risk. For example, Philadelphia also conceptualizes their wireless network as a way for municipal services, such as libraries, to connect to the internet. Supporting this kind of use requires broader and more consistent coverage than Cambridge initially envisions.

Cambridge has the luxury of focusing its wireless mesh efforts on the digital divide because it already has a fibre-optic network for municipal services. Around 50 years ago when the phone company put its wires underground, the city required them to also install empty conduits for future city use. In recent years the city has been threading fibre through these conduits. Mesh access points will be established in select neighbourhoods in Cambridge by connecting to this fibre network.

The ● February 1st article in ● The Tech claiming that Cambridge plans to deploy for the entire city this summer seems to be mistaken. None of the people we have spoken to who are actually involved in the project, especially those who work for the city, display this kind of irrational exuberance. It is possible that they have drastically scaled back their plans since February, but we doubt that is the case. The ● February 2nd article in the Boston Globe mostly echoes the claims of The Tech article.

2. History

In early 2005 the city decided to explore the possibility of a free municipal wireless mesh network. This decision was made in part because of a refusal of local broadband providers to adjust pricing to residents income-level. Councillor Henrietta Davis is currently the chair of the Cable TV, Telecommunications and Public Utilities Committee, and has been an organizer of this project from the beginning.

In mid-2005 MuniMesh (Kurt Keville and Bob Keyes) approached the city to discus the RoofNet technology. An official committe was formed in November 2005, and the first beta deployment is planned for summer 2006.

The beta deployment will be based on three primary gateways: the Lombardi building beside city hall, a tall apartment building owned by the city on the Cambridgeport side of Mass Ave, and MIT.

last edited 2006-05-16 23:17:57 by 128

ReflectionsAndComparisons

1. Reflections and Comparisons

1.1. Analogies to Other Systems

During the analysis, we compared the RoofNet wireless mesh network with two models simulating LAN (local area network) and WAN (wide area network) systems. We developed a random graph to identify the differences between Roofnet and a random graph. By examining the network architectural metrics, we have two hypotheses about the -ilities of Roofnet: one is that Roofnet is a very decentralized network relative to the internet; thus, it is robust and not a fragile network architecture (rather than "robust yet fragile" architecture of the internet network). The other is that Roofnet is very similar to a random graph in terms of the clustering coefficient and degree correlation properties; however, we find that RoofNet is centralized relative to the random graph because of its geographical and technical constraints.

1.2. Learning from this Project

We learned how to use UCINET and MATLAB to perform network analyses. The application of tools and methods helped us to appreciate the numerical metrics and link the topology to the properties of the network architecture. These metrics can provide some measure the network and help us to understand the network, especially when the network is extremely complex.

We also learned to think about complex system architectures in different ways. Without any expertise about Roofnet, we worked on this project from an architectural perspective. The analysis gave us insights into the properties of this network vs. other networks. There certainly is a lot more work that can be done in this area! The two models we considered are quite simple. If we had access to more (clean) data about Roofnet and the Internet, we might have been able to compare these two networks directly to see the differences between the -ilities of each network. It would still be interesting to examine the relationship between decision (routing, congestion, etc) protocols on the topological properties of these types of networks.

On the other hand, it seems that it would be very difficult to gain any insight about a complex system using these metrics without specific knowledge of the system. One has to know some technical aspect of the system to be able to link those metrics with the actual properties of the system to get anything useful and meaningful. The meaning of the metrics seem to be subject to a great deal of interpretation based on the system under study.

1.3. Comments on System Architecture Analysis and Description

It seems as though the metrics themselves need a lot of work. Either they are not enlightening because they tell us something we already know about the network or they don't seem to say anything meaningful about the structure of the network (sometimes with or without knowledge of the system under study). Perhaps it would be useful to focus on finding ways to parameterize the network structure. One of the most interesting and telling studies seemed to be the parameterization of organization structure in the Dodds, Watts, and Sable paper.