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:

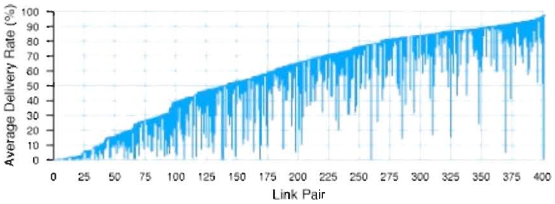

Because the purpose of this analysis for aggregate data is to study the topological differences with other networks, we symmetrized the links when we perform the numerical analysis. Therefore, the network to be compared with the random graph and the benchmarking models is a symmetric network. The network 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.

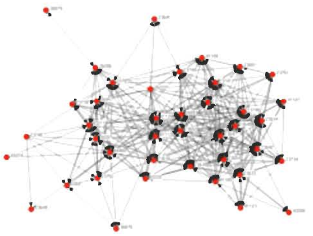

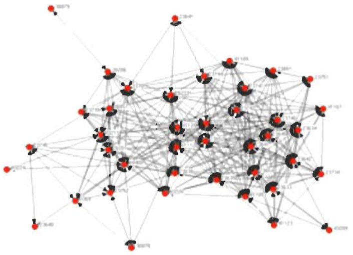

2. Because of the method to aggregate the data, there can be many links between A and B. These links incorporate topological information for the different bit rates (from the four experiments) and the different delivery probabilities over time due to weather, multi-path fade, and other signal disturbances. 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.

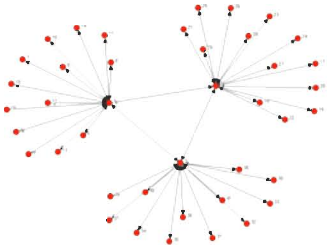

Model 2 is analogous to a WAN (Wide Area Network) system in the sense that it is a collection of LANs. Thus, each user has access to a LAN hub which has access to a WAN hub and connects to other LANs and WANs through these interconnections.

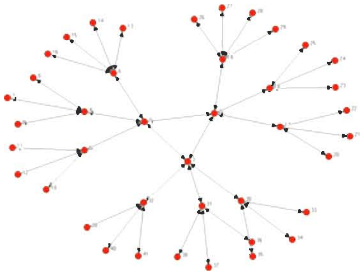

3. Contrasting with LAN/WAN Benchmarks

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)

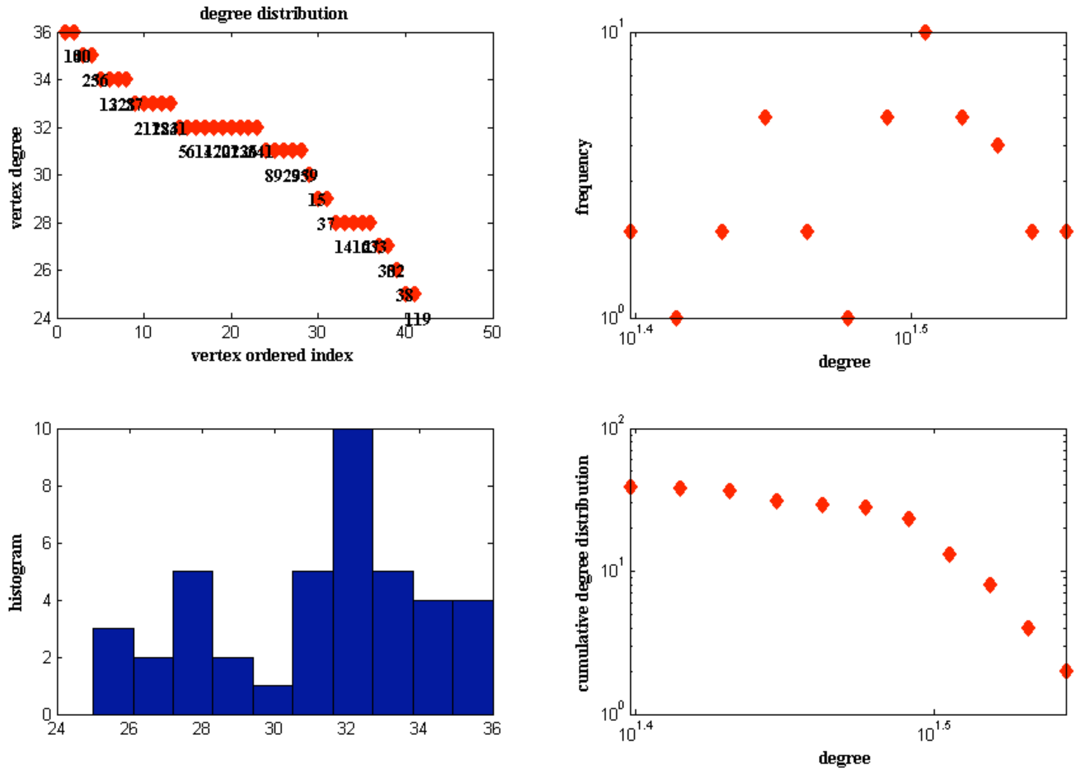

degree distribution for Roofnet (symmetrical)

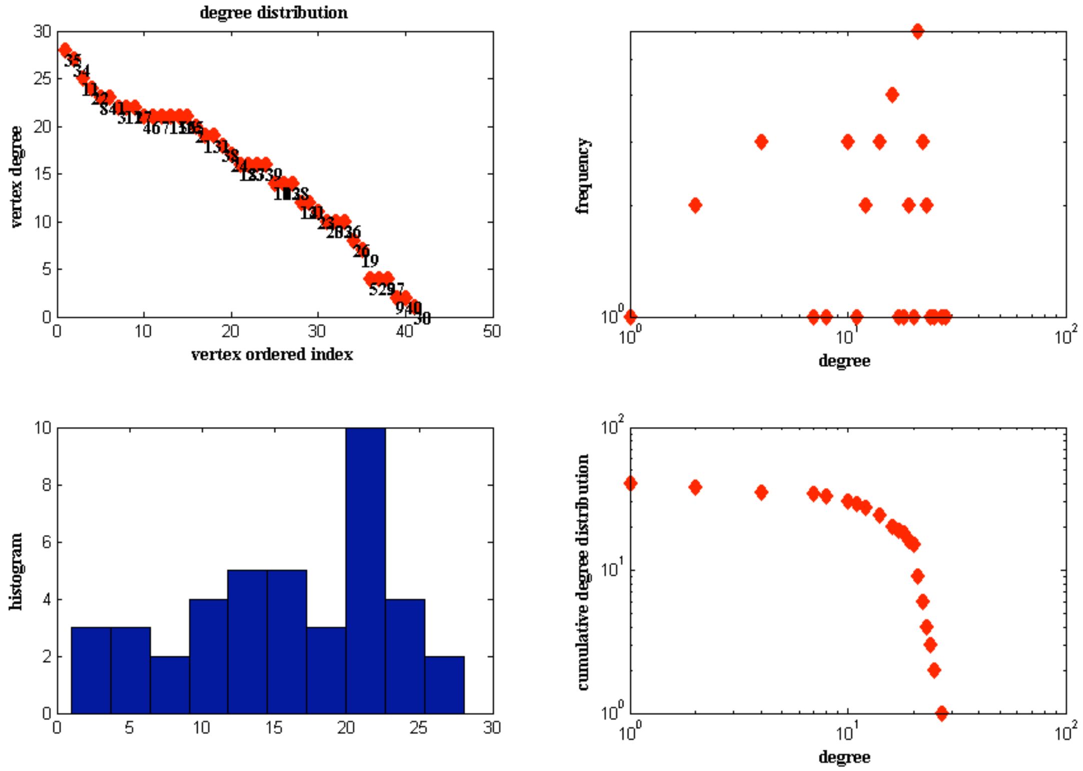

5. Robustness Analysis

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.

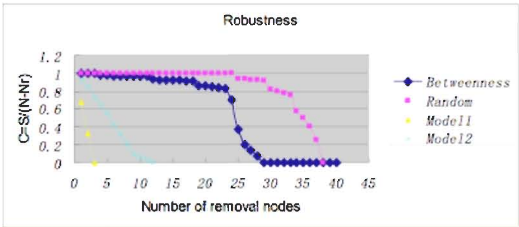

C=S/(N-Nr). S is the number of remaining nodes; N is the total number of nodes; Nr is the number of nodes that are removed.

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

Model 1:	Model 2:
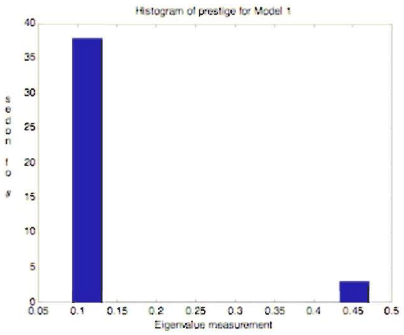 	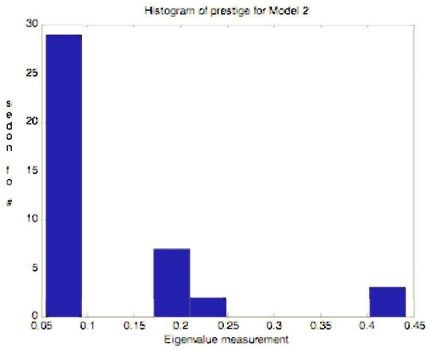
RoofNet (asymetrical)	RoofNet (symmetrical)

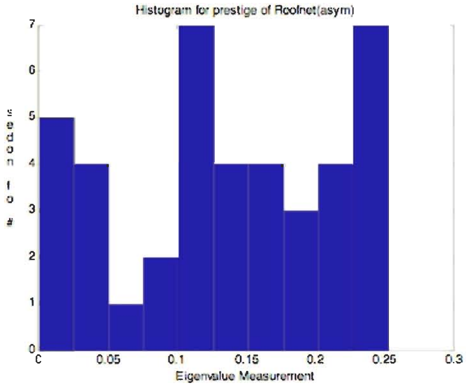	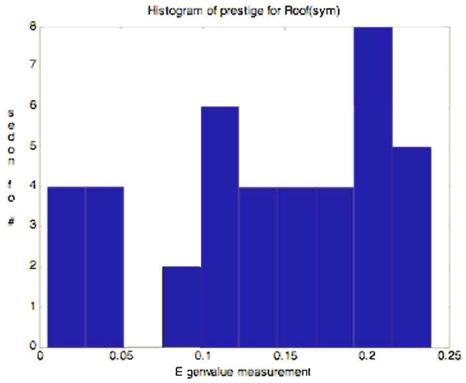
Random	
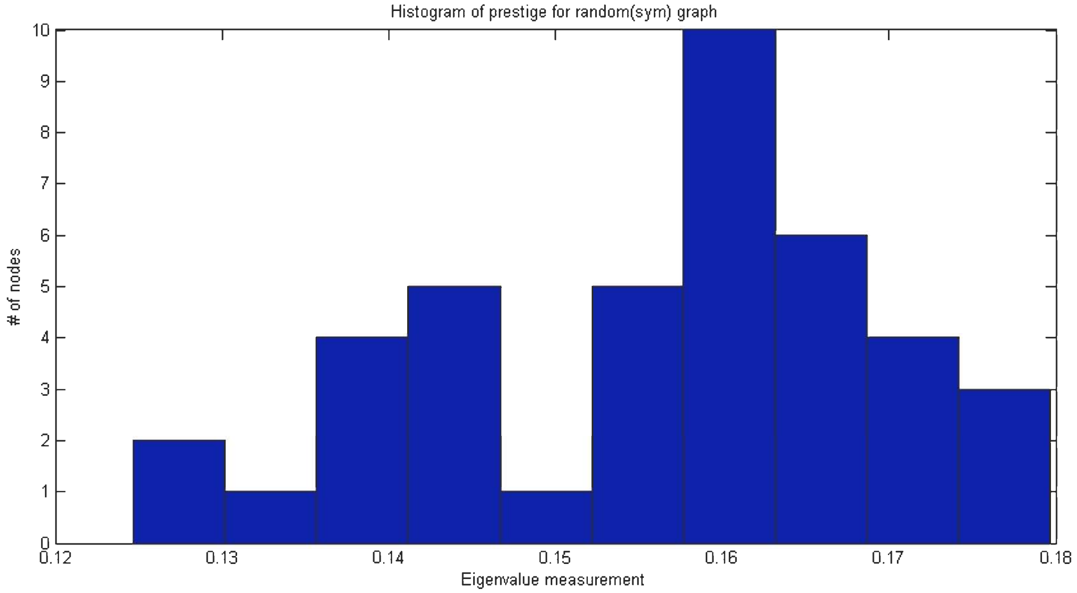	

6.2. Acquaintance

Model 1:	Model 2
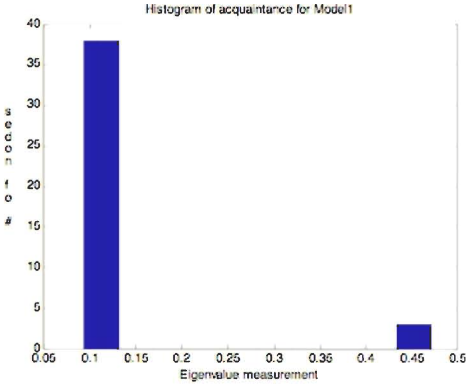	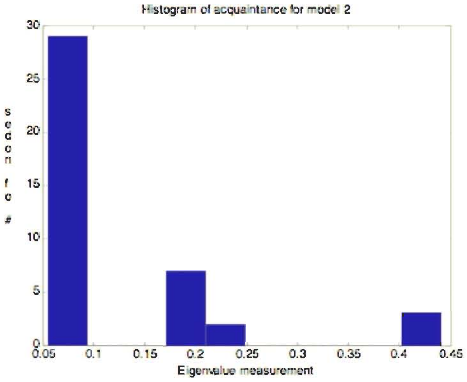
RoofNet (asymetrical)	RoofNet (symmetrical)

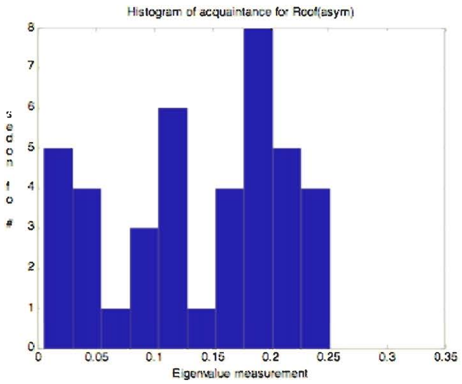

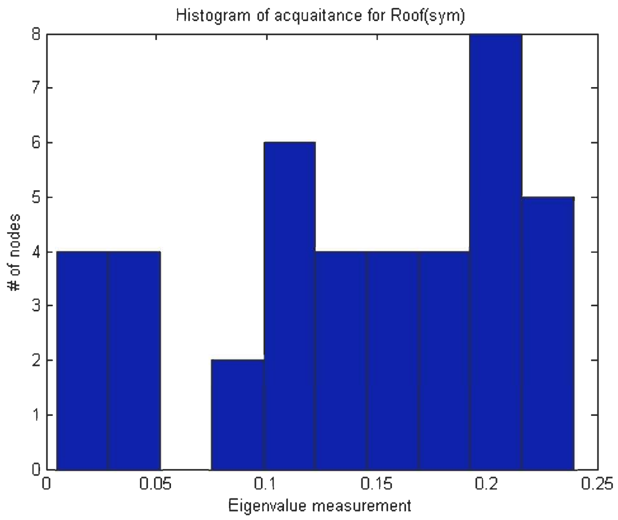

Random (symmetrical)

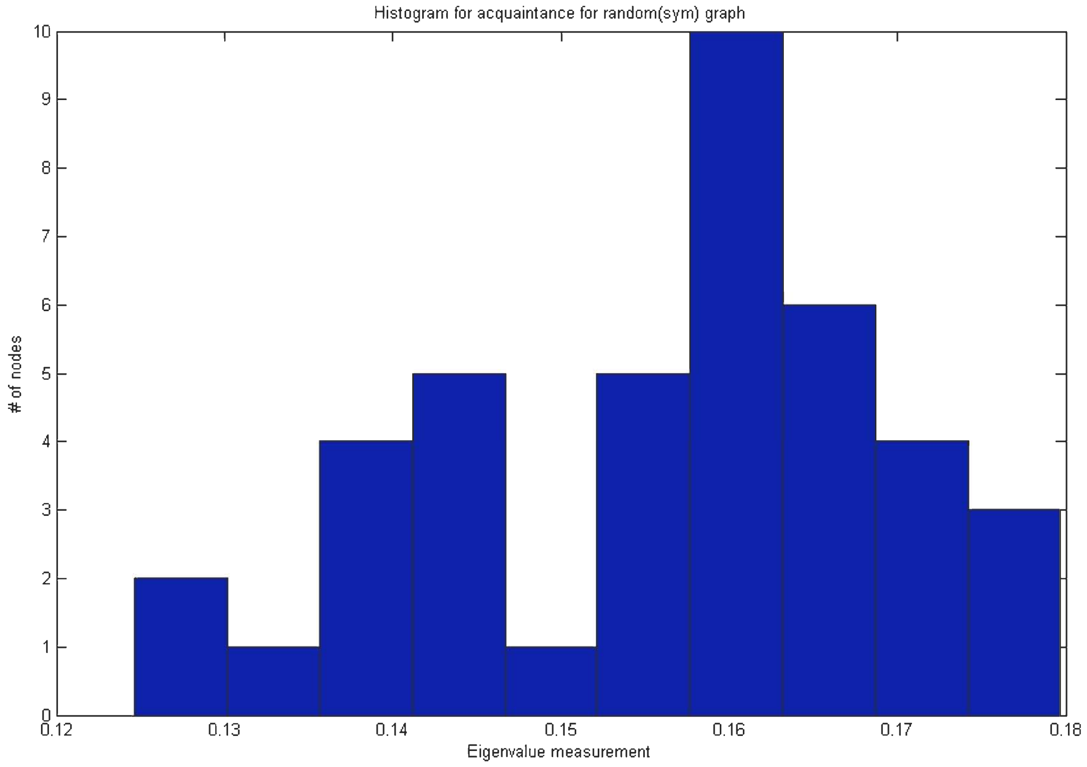

last edited 2006-05-16 20:53:04 by vpn-eighty-six

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).

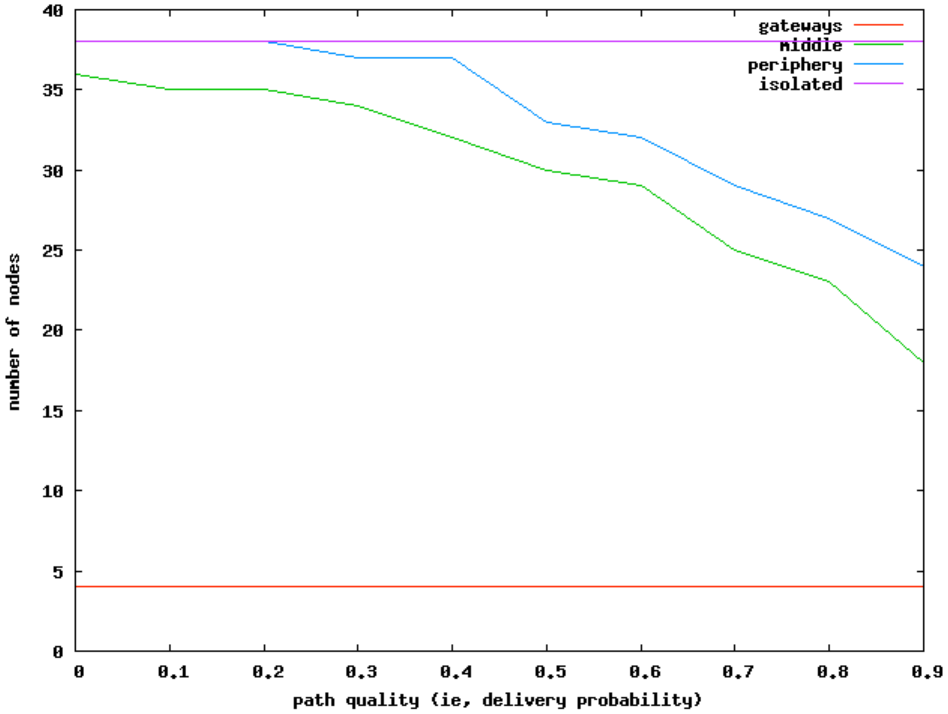

5. Classifying Nodes

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

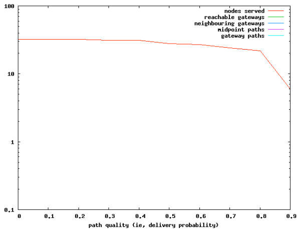

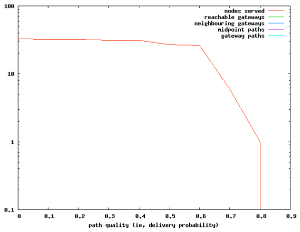

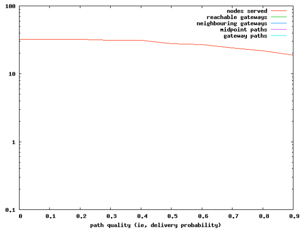

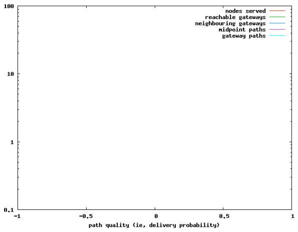

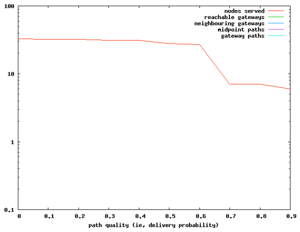

5.3.2. Isolated Nodes

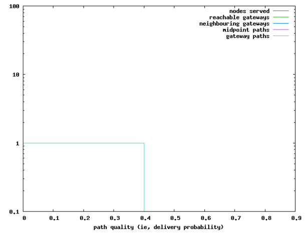

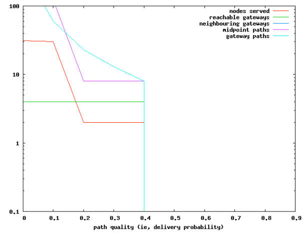

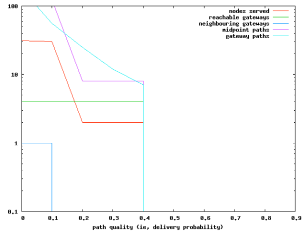

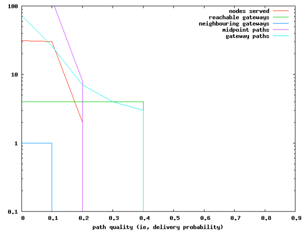

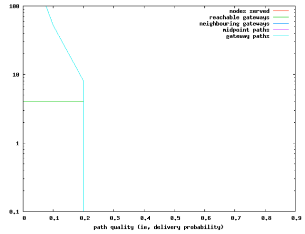

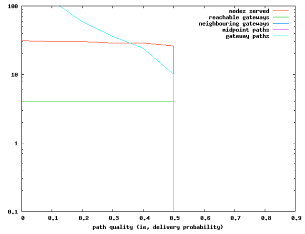

5.3.3. Single-hop Mid-point Nodes

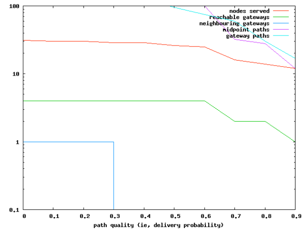

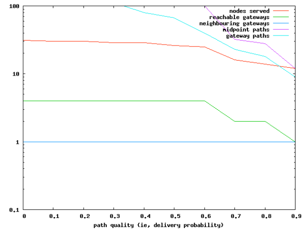

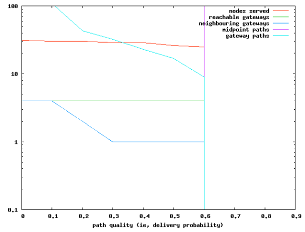

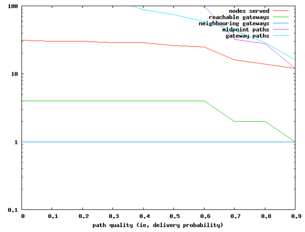

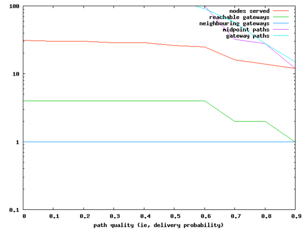

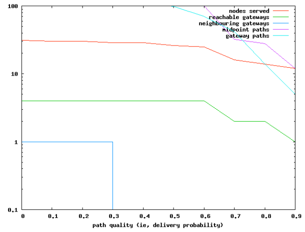

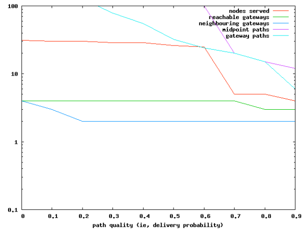

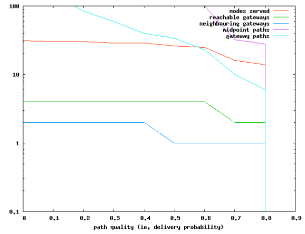

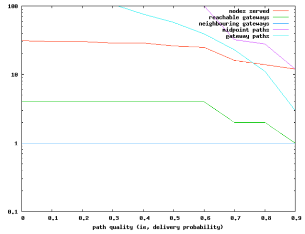

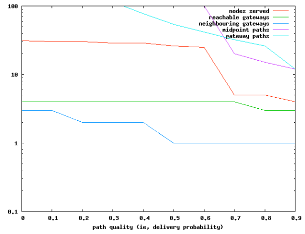

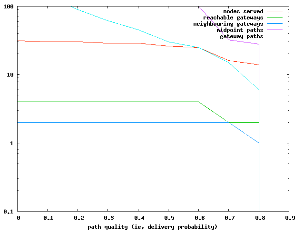

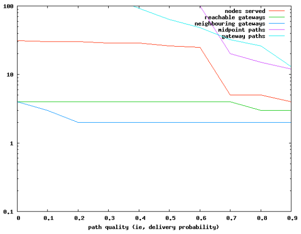

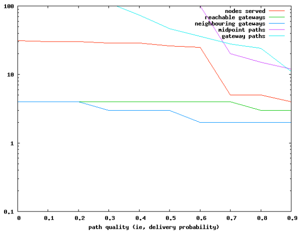

5.3.4. Single-hop Periphery Nodes

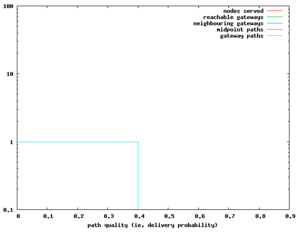

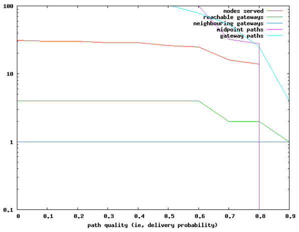

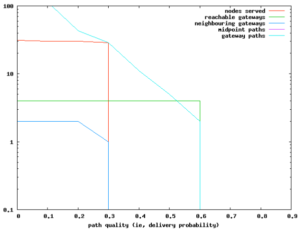

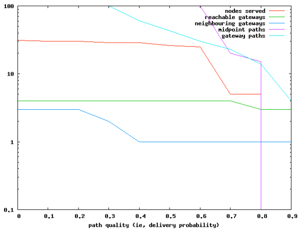

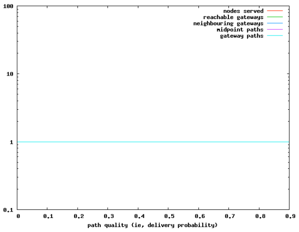

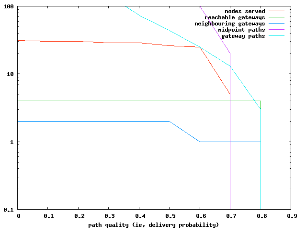

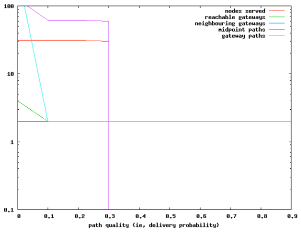

5.3.5. Multi-hop Mid-point Nodes

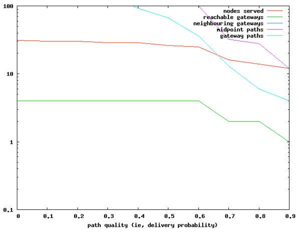

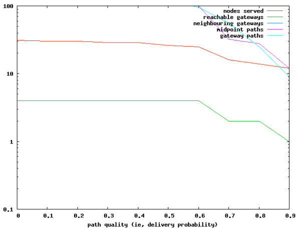

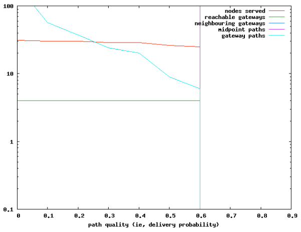

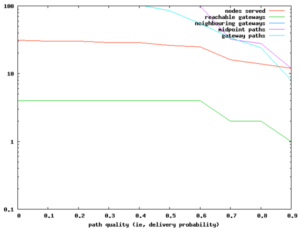

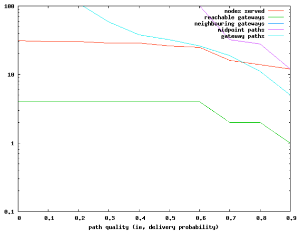

5.3.6. Multi-hop Periphery Nodes

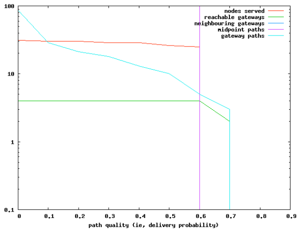

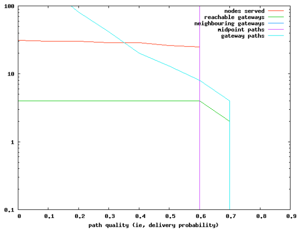

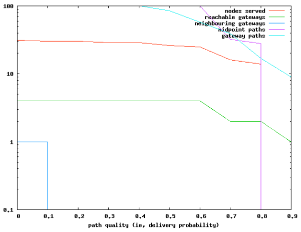

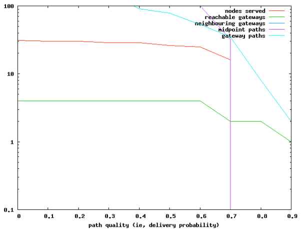

6. Improving the Mesh by Strengthening an Edge

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.

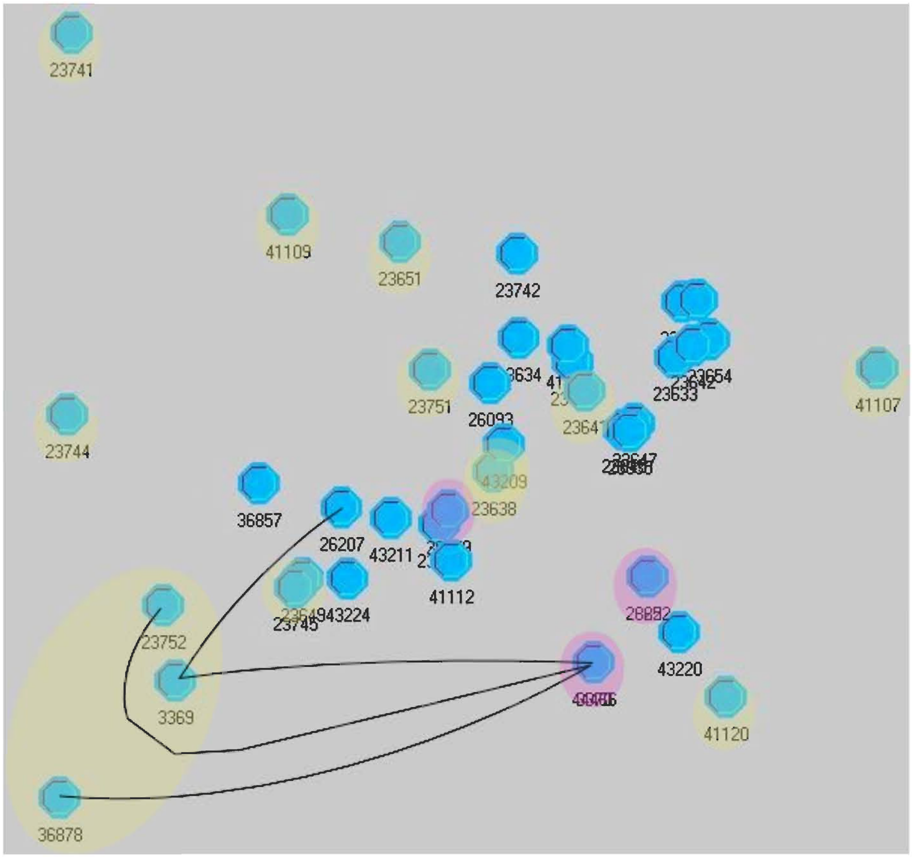

The fact that each of these four edges would re-connect the same three nodes suggests that a community-finding/clustering algorithm may also be able to identify this group.

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.

last edited 2006-05-16 19:43:11 by vpn-eighty-six