

Empirical Analysis of the Single-Radio Mesh Architecture Performance Limitations and Challenges

Andres Arjona
Nokia Siemens Networks
Espoo, Finland
andres.arjona@nsn.com

Cedric Westphal
Nokia Siemens Networks
Mountain View, California
cedric.westphal@nsn.com

Sami Takala
Nokia Siemens Networks
Espoo, Finland
sami.j.takala@nsn.com

ABSTRACT

Wireless mesh networks are being deployed to provide broadband wireless connectivity to city-wide hotspots. The most successful commercial architecture in these deployments thus far is a single-radio architecture: mesh nodes carry only one radio, which is used both to receive the traffic from the WiFi clients and to relay this traffic through the mesh to the wired internet gateway.

In this paper, we study the performance of a representative single radio mesh network and evaluate several key quantities: the fairness in the bandwidth allocation, the impact of the hidden node terminal, and the number of concurrent VoIP supported by the network.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement Techniques and Performance Attributes.

General Terms

Measurement, Performance, Experimentation.

Keywords

mesh, 802.11, WiFi, multihop, VoIP, performance, measurements

1. INTRODUCTION

Over the last few years, wireless mesh networks that provide WiFi services over large areas, including entire municipalities, have been successfully deployed. Many vendors have proposed products and solutions to provide wireless broadband connectivity over extended areas [3][12][16][17][21].

Cedric Westphal's research was partially supported by a grant from the Japan Society for the Promotion of Science (JSPS).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IWCMC'07, August 12–16, 2007, Honolulu, Hawaii.
Copyright 2007 ACM 978-1-59593-695-0/07/0008...\$5.00.

The mesh network is composed of three basic elements: an 802.11 WiFi client, such as a laptop or a VoIP phone, for the end-user to connect to the network; an access point (AP) for the client to attach to; and some mesh routers to forward the traffic over potentially multiple wireless hops in between the AP and the wired internet gateway. The mesh routers form the backhaul of the mesh network (see Figure 1).

Metropolitan wireless mesh networks are seen as a potential disruption to cellular operators, by providing VoIP services to mobile handsets ubiquitously in a city-wide area. As such, the performance of VoIP connections is critical to the commercial success of these mesh networks.

The commercial solution most commonly deployed is a single radio mesh node, where both AP and backhaul functionality are merged on the same platform. The same node functions both as an access point to the end-user clients and as a relay for the traffic to the gateway.




Figure 1. WiFi Mesh in an Outdoor Scenario

Most vendors offer a single radio solution, and the market leader, Tropos [21] offered exclusively single radio products until recently. However, single radio architectures were chosen mostly due to their lower price point, and to their implementation simplicity in an emerging market. A systematic study of the network performance did not enter in the selection of the architecture.

This paper sets out to study the performance of single-radio mesh networks by conducting a methodic performance analysis of a

mesh network composed of single-radio platforms of the latest current generation.

The key contributions of the paper are to evaluate and quantify the performance of a single-radio mesh network for fairness of the bandwidth allocation over several users, for resiliency to the hidden terminal problem, and for support for VoIP connections.

2. SINGLE-RADIO WIFI MESH ARCHITECTURES

The single-radio architecture uses the same radio frequency for wireless access and backhaul. This technical choice represents the lowest cost in the deployment of a wireless mesh network. Many metropolitan wireless mesh deployments are currently based on this architecture, such as those from Earthlink and Tropo in Philadelphia and San Francisco or Google's network in Mountain View [2]. The single-radio architecture is substantially cheaper than dual or multi-radio architectures. On the other hand, multi-radio architectures offer better performance, as the backhaul traffic typically is in a different frequency band as the access traffic, and does not interfere with it. For a performance benchmark of multi-radio mesh equipment please refer to [11].

Further, as depicted in Figure 1, nodes might be several hops away from the wired gateway, and might need to use other nodes as relays. Therefore, not only does the backhaul traffic interfere with the access traffic, it also interferes with itself. As more mesh access points are added, a higher percentage of the wireless bandwidth is dedicated to packet forwarding. Thus, less capacity is available to for the users' traffic [4].

We are interested in assessing the upper bound on the mesh network capacity using single radio nodes. As such, we make some simplifying assumptions:

- We consider a line topology. Actual mesh networks follow tree topologies, where the different branches of the tree interact with each other. Considering a line network gives us an upper bound on performance versus the number of hops.
- We consider nodes which interfere only with their neighbors. Again, different topologies and node placement will have different results in terms of achieving a transmission schedule. However, removing interference from nodes more than one hop away gives an easily reproducible upper bound.
- We consider networks with size limited to three hops. While this could seem small, it is actually the deployment rule followed in real life networks: nodes are never further than three hops away from a gateway. Typically, the gateway is not a wired gateway, but a point-to-multipoint link with a higher capacity, such as an Alvarion [1] or Motorola Canopy [13] link.

Under these assumptions, only one wireless link can be up at any given time, and consequently, the expected throughput of a single-radio mesh network should coincide with its theoretic behavior and be proportional to the inverse of the number of hops.

While we are considering single-radio mesh networks in general, there are differences between the different possible implementations. For instance, the use of proprietary Medium Access Control (MAC) mechanisms in the backhaul, and

specialized routing algorithms can result in differences in performance. [18]

We evaluate in this document a single-radio node from the latest generation, the Meraki mini node [12]. The Meraki mini exists in indoor and outdoor configuration, and was introduced commercially in August 2006 as a fully functional beta version. The testing took place in November 2006, at which point the firmware had been upgraded to the latest version.

Meraki is a start-up founded by students from the RoofNet project [15] at MIT, and the product includes the state-of-the-art result of this research group.

3. METHODOLOGY AND TEST ENVIRONMENT

The evaluation methodology consists of a variety of test cases with multiple combinations of number of hops, number of clients and network traffic. With such variables we are able to evaluate the single-radio solution from the following test objectives:

- a) Tradeoff between client bandwidth and number of hops from the gateway.
- b) Tradeoff between bandwidth share for clients near to the gateway over clients far from the gateway.
- c) Effect and capacity for simultaneous clients connected to different network nodes.
- d) Effect of rate limiting the users at 1Mbps, which is a typical value for subscriber service in commercial mesh networks.
- e) Hidden node problem effect in single-radio equipment.
- f) VoIP quality with multiple clients and number of hops from the gateway.

The test bed configuration consisted of single-radio Meraki mesh access points located individually in RF enclosures. The access points had default configuration and maximum modulation enabled. Each of the RF enclosures was connected to each other in such a way that multiple hops are possible. No attenuators were placed between the access points in order to allow the backhaul to operate at a maximum rate. One of the access points at the edge of the wireless network was connected to a switch in the wired network, where the management consoles and a web server were located (see Figure 2). Laptops were connected directly to the RF enclosure that contained the respective access point to which the client was to associate. The number of laptop clients varied between one and three. The RF enclosures assured that the laptop clients were only able to receive a high WiFi signal from the access point located in the same RF enclosure. The signal to noise ratio for each of the laptop clients in the enclosures was large enough to support the highest transmission rates (54Mbps).

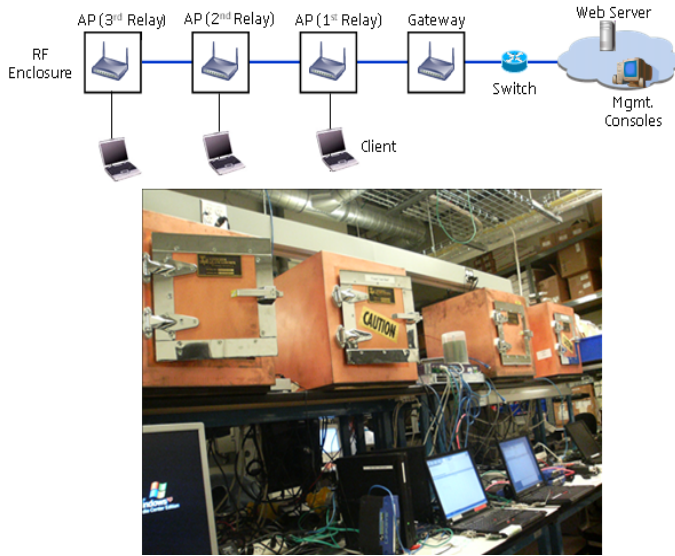


Figure 2. Single-Radio Mesh Testing Configuration

The tests executed consisted of multiple iterations of test cases for each of the mentioned test objectives. Throughput was measured by means of 25MB FTP file downloads and uploads. Delay was characterized with PING. HTTP Browsing was measured as download times for webpages with different sizes and layouts located in a web server. Background traffic was generated with FTP continuous file downloads and uploads. Finally VoIP quality and capacity is measured with IxChariot testing software. Chariot's mean opinion score values are based on ITU-T E-Model. [6][9][10]

4. MEASUREMENT RESULTS

4.1 Available Bandwidth vs. Number of Hops

The performance tests carried out show that the maximum available bandwidth per cell is determined at the gateway; 5.2Mbps and 4Mbps for downlink and uplink respectively when there is no background traffic. Regardless of supporting 54Mbps transmission rates, the Meraki nodes limit client bandwidth to 5Mbps because they are also mesh nodes and dedicate a fraction of their time listening to potential traffic from other mesh routers. The same throughput is achieved regardless of whether there was only one client associated to the access point or two. The maximum available bandwidth is shared among the two clients associated to the access point.

However, Figure 3 shows that the bandwidth share between the users is not equal. One of the users is able to get a considerably higher throughput, especially when both users are associated to the gateway or the 1st relay. In contrast, Figure 4 shows that the share for the uplink direction is almost equal, yet the maximum throughput achievable is lower than for downlink.




Figure 3. Bandwidth vs. Number of Hops [DL]




Figure 4. Bandwidth vs. Number of Hops [UL]

The capacity decrease in our test bed approached the expected (1/number of hops) rate. However, we must note that in our test environment, access points and clients are isolated via the use of RF enclosures (see Figure 2). Therefore, access points can only hear adjacent neighbors. In a real outdoor configuration, the gateway node is likely to be located in a middle location to reduce the number of hops to its neighbor nodes (see Figure 5). As a result, the gateway will probably hear more than one or two relay nodes. Since all clients and access points must operate in the same frequency channel, and contend for access based on the 802.11 standard, the capacity will possibly be lower in such scenarios, and our measurement give an upper bound on the achievable throughput.




Figure 5. Single-Radio Mesh Cluster

In regards to latency, the average delay was very low (10 to 15 ms), despite of the number of hops from the gateway. The browsing experience was good in average, with download times between 2 and 6 seconds for 275KB pages.

4.2 Near Client vs. Far Client

The bandwidth allocation already exhibits unfairness when clients are connected to the same access point. The following test results

show that connecting the clients to different APs only exacerbate the unfairness. To investigate this, the setup consisted of one client associated with the gateway and a second client associated to the 1st, 2nd and 3rd relaying AP respectively (see Figure 6 and Figure 7).




Figure 6. Near Client vs. Far Client [DL]




Figure 7. Near Client vs. Far Client [UL]

The results show that the maximum throughput of the furthest client, that is, the one that is far from the gateway, is directly affected by the clients utilizing resources from the gateway. When clients are located at different nodes, the total bandwidth available for all nodes is reduced to roughly 3.6Mbps. Additionally, the client in the gateway increases its share of the bandwidth allocation the further away from the gateway the other client connects to the network. The approximate bandwidth gain is around 14%. Moreover, browsing experience is affected very much when a client is more than one hop away (2nd relay or further) and simultaneous traffic from the other client exists. In such cases, the webpage download times increase drastically to 10-30 seconds. Therefore, in a real deployment, the number of hops should be limited to one or two at the maximum in order not to compromise the browsing application performance.

4.3 Client Traffic Mixes

The following tests aimed at understanding the effect of multiple users in different nodes and number of hops from the gateway. The test setup consisted of associating a client and running FTP file download and uploads on the gateway and the 1st relay, and subsequently increasing the number of clients by adding another client to the 2nd relay or further. Figure 8 shows the test configuration for three clients, one in the gateway, one in the 1st Relay and one in the 3rd Relay.




Figure 8. Test Example: Simultaneous Clients in Gateway, 1st Relay and 3rd Relay

The results show that the limit for simultaneous clients in different nodes appears to be three. Already with three clients, most of the time, the client that is furthest from the gateway does not receive any service at all in the downlink direction, and sometimes it does not either in the uplink. This issue occurs both when connecting to the 2nd relay and to the 3rd relay.

Even when the third client does not receive any service, as long as it tries to access the medium it will affect the performance of clients in the gateway and 1st relay considerably. The available bandwidth can drop as much as 75% as long as the third client tries to gain access to the medium. Figure 9 shows the bandwidth performance view for the client connected to the gateway when a third simultaneous client further apart intends to access the medium.




Figure 9. Third Simultaneous Effect on Performance

4.4 Effect of Maximum Bandwidth Limit of 1Mbps per User

In order to investigate the problem with multiple users further, we limited the maximum available bandwidth per user to 1Mbps, using a management feature of the Meraki nodes. Several of the metropolitan wireless mesh networks with single-radio architecture enforce this limit, including the Google network in Mountain View [2][8]. With this limitation, the furthest client is able to receive service. However, latency increased to an average of 200ms. This illustrates the need for rate limiting, traffic monitoring and potentially connection admission control in the mesh network.




Figure 10. Bandwidth Share for Three Simultaneous Users in Different Nodes

4.5 Susceptibility to Hidden Node Problems

In order to increase the throughput of the mesh network, single radio nodes customarily disable the RTS/CTS exchange. This exposes the nodes to the hidden terminal problem [14][20].

In order to test the susceptibility of the nodes to the hidden node, clients placed in different RF enclosures were forced to be unaware of each other while still being able to transmit data to a common access point. The results show that the existence of even one hidden client is very harmful to the network performance. This problem affects uplink transmissions much more than downlink. The reason for this is that even though clients are not aware of each other, the access point is aware of them and synchronizes transmissions to each of them (see Figure 11).

Furthermore, while the scenario is symmetric for both clients, the bandwidth allocation is actually unfair. The throughput for one client is very low, while the throughput for the other client fluctuates a lot. The problem is less damaging in cases where the majority of users synchronize transmissions with each other. However, the active hidden client still hinders performance by forcing heavy service fluctuations in the rest of the clients.




Figure 11. Hidden Node Problem Effect on Performance

The trade-off for the single-radio node is thus to maximize the throughput by disabling the mechanisms to handle the hidden node problem. This adversely affects the overall performance when the hidden node situation actually appears. While the choice to not enable RTS/CTS might prove better overall for the network, it is very detrimental to the nodes placed in the adverse situation.

4.6 VoIP Quality and Capacity

VoIP mean opinion score quality was tested with IxChariot [10] and consisted of client mixes with and without simultaneous background traffic. Traffic load was generated with a continuous 2Mbps UDP data stream on the downlink direction. The resulting

maximum number of calls using the G.711 codec with toll call quality was 6 to 7 for VoIP only scenarios, which is lower than the normal 802.11g capacity [5]. In scenarios with background traffic load, the maximum number of users was reduced to 6. Studies in [19] describe similar mesh results. Figure 12 and Figure 13 summarize the VoIP results.




Figure 12. VoIP Capacity without Simultaneous Data Traffic




Figure 13. VoIP Capacity with Simultaneous Data Traffic

5. CONCLUSIONS

Extensive measurements were carried out to evaluate and characterize the single-radio architecture, and some of its limitations. The test results prove that single-radio mesh architecture is not suitable for building large outdoor networks. There are multiple issues that can hinder performance and limit its usability. Despite of this, numerous networks have been deployed in different cities based on the single-radio architecture. Its performance is likely to be far from optimal as it is shown in measurements from live networks. We believe a dual- or multi-radio architecture should be used instead as an alternative technology.

The test bed results show that the capacity decrease rate is at its best 1/(number of hops), and probably lower in real life deployments. Also, the bandwidth share for clients is not distributed equally. When the clients are located at different nodes, the total available bandwidth is lower. In such cases, the clients located closer to the gateway are able to receive a higher bandwidth. The maximum number of users at different nodes appears to be three. With three simultaneous users, the overall performance decreases. Likewise, the furthestmost client is not able to receive any service most of the time. However, when the maximum bandwidth per user is limited to 1Mbps the problem seems to be partially solved, but at the cost of higher latency.

Finally, the hidden node problem is very damaging. Even with only one hidden client, the network performance is affected drastically, and can drop other users' throughput considerably. VoIP quality is also affected by elastic background data traffic. With simultaneous background data traffic, the maximum number of simultaneous calls with adequate voice quality is six.

6. ACKNOWLEDGMENTS

We would like to acknowledge the suggestions and contributions of Jan Berglund, Zheng Li, Antti Ylä-Jääski, Graciela Bueno and Harri Paloheimo.

7. REFERENCES

- [1] Alvarion, <http://www.alvarion.com>
- [2] A. Arjona, and S. Takala, "The Google Muni WiFi Network – Can it Compete with Cellular Voice?", AICT'07, Morne Mauritius, May 13-19, 2007.
- [3] BelAir Networks, <http://www.belairnetworks.com>
- [4] BelAir Networks, "Capacity of Wireless Mesh Networks – Understanding Single Radio, Dual Radio and Multi Radio Wireless Mesh Networks" BelAir Networks 2006.
- [5] L. Cai, X. Shen, J. Mark, L. Cai, and Y. Xiao, "Voice Capacity of WLAN with Unbalanced Traffic" QShine'05, Orlando USA, Aug. 22-24, 2005.
- [6] R. Cole, and J. Rosenbluth, "Voice over IP Performance Monitoring" ACM SIGCOMM'01.
- [7] Comptek Aqua, "Comparative Analysis of Different Media Access Methods in Metropolitan Area Networking", Available at http://aqua.comptek.ru/test/HiddenNode/hidden_node_en.html
- [8] Google, "Google WiFi Mountain View", Google 2007. Available at <https://wifi.google.com/support/>
- [9] ITU-T, "Rec. G.107 – A Computational Model for Use in Transmission Planning" ITU-T 2003.
- [10] Ixia, "IxChariot", Product Specification. Available at <http://www.ixiacom.com/products/display?skey=ixchariot>
- [11] B. Mandeville, C. Mandeville, and A. Bouard, "Wireless Mesh: Ready!" Light Reading June 20, 2006. Available at http://www.lightreading.com/document.asp?doc_id=96200
- [12] Meraki, <http://www.meraki.net>
- [13] Motorola Canopy, <http://motorola.canopy.com>
- [14] Y. Kim, J. Yu, S. Choi, and K. Jang, "A Novel Hidden Detection Mechanism in IEEE 802.11 WLAN", IEEE Comm. Letters vol. 10 issue 8, Aug. 2006, pp 608-610.
- [15] RoofNet, <http://pdos.csail.mit.edu/roofnet/doku.php>
- [16] SkyPilot, <http://www.skypilot.com>
- [17] Strix Systems, <http://www.strixsystems.com>
- [18] W. Terrill, "Wireless LAN Mesh: Wireless Wireless", Network and Telecom Strategies In-Depth Research Report, Burton Group June 02, 2005.
- [19] K. Ting, Y. Ko, and M. Sim, "Voice Performance Study on Single Radio Multihop IEEE 802.11b Systems with Chain Topology", In Proceedings of ICN'05.
- [20] F. A. Tobagi and L. Kleinrock, "Packet switching in radio channels: Part II - the hidden terminal problem in carrier sense multiple-access modes and the busy-tone solution," IEEE Trans. Commun., vol. COM-23, no. 12, pp. 1417-1433, 1975
- [21] Tropos Networks, <http://www.tropos.com>