Performance Measurements of Multimedia Flows over a WiFi-based Mesh Network

Roberto Riggio, Daniele Miorandi,
Giovanni Iacovoni and Imrich Chlamtac
CREATE-NET
Via dei Solteri 38, 38100, Trento, Italy
Email: name.surname@create-net.org

Nicola Scalabrino and Enrico Gregori
Italian National Research Council (CNR) - IIT
Via G. Moruzzi 1, 56124, Pisa, Italy
Email: name.surname@iit.cnr.it

***Abstract*— Wireless mesh networks are currently becoming one of the most promising approaches to provide ubiquitous broadband Internet access. In order to successfully make their way as access architecture for the next-generation Internet, mesh networks need to provide mechanisms able to efficiently support emerging broadband multimedia services. In this work, we report some performance measurements obtained on an experimental WiFi-based mesh testbed running at CREATE-NET premises. The tests aim at characterizing the suitability of current mesh networking solutions to support multimedia flows. The obtained performance is compared to those obtained by means of a conventional star-shaped topology based on the use of access points. The results show that mesh architectures are able to offer some advantages, in terms of fairness and lower packet loss rate, with respect to a standard access points based architecture.**

I. Introduction

As the trend toward broadband ubiquitous networking gains momentum, new networking paradigms are needed to fit the peculiarities of such novel scenarios. Wireless mesh networking has recently emerged as one of the most promising access architectural paradigms, being able to address a wide range of application scenarios, including home broadband Internet access, enterprise networking and metropolitan area networks [1], [2].

Wireless Mesh Networks (WMNs) rely on a multi-hop wireless backbone for delivering high-speed services to end-users without the need for deploying any fixed infrastructure. With respect to conventional star-shaped access network architectures, WMNs offer advantages in terms of enhanced robustness (in that no single points of failure are present and redundant links are encompassed) and flexibility (without the need for deploying cables, connectivity may be provided only where and when needed/economically attractive). With respect to conventional ad hoc networks [3], WMNs differ for (i) the *goal*, in that they are being intended as access architecture, not stand-alone systems (ii) the *heterogeneity of the devices*, in that there might be dedicated devices (with more powerful radio systems, multi-band capabilities etc.) acting as pure wireless routers. As an example, we may consider a wireless interconnection of hot spots, providing enhanced coverage without the need of having all of them wired to the Internet.

In order to be successful, WMNs must cope with current trends in services. It is indeed widely acknowledged that the next-generation Internet will be characterized by an extreme variety of multimedia broadband services. Without the ability to successfully support the peculiarities of these services, WMNs run the risk to remain a niche market. Unlike "pure data" applications like FTP or HTTP, next-generation services are characterized by requirements in terms of network support, i.e., bandwidth, latency, packet delay jitter etc. On the other hand, these constraints fit badly the decentralized architecture of WMNs, where smart solutions are needed to provide such performance guarantees. It is therefore a primary need to perform performance measurements on real-world testbeds, in order to characterize the ability of WMNs to support multimedia flows and gain insight into the critical points of such systems, therefore providing smart guidelines for the design of innovative solutions able to boost WMNs deployment.

In this work, we report some performance measurements obtained at our CREATE-NET testbed on a small-scale (7 nodes) IEEE 802.11-based WMN.

The tests aim at characterizing the ability of current WMN technology to support multimedia flows. The literature provides already some performance studies on WMN testbeds, from which our work differs in (i) the network architecture, in that we employ a single-tier architecture and (ii) the evaluation methodology, in that the performance of a mesh architecture is compared to that obtained by a standard star-shaped single-hop architecture.

Most works in the literature focus on outdoor metropolitan-scale deployments. For instance, [4] reports an analysis of the possible sources of packet loss in an outdoor WMN, assessing the effect of link distance and signal-to-noise ratio on the link quality statistics. Results show that a sharp dichotomy between working and not working link cannot be found, the majority of the links being characterized by an intermediate loss rate. In [5], the performance of an outdoor WMN is evaluated, discussing the effect of node density on network connectivity and throughput. Compared with a star-shaped network, the mesh architecture improve both the connectivity and the throughput. Results for an indoor environment are reported in [6], where the performance of multimedia flows over an IEEE 802.11-based two-tier WMN are given in terms of packet latency, loss rate, inter-flow fairness and jitter for different network configurations. Results show that the number of multimedia flows that can be supported by the network is constrained by the application packet rate, therefore, performance can be enhanced by aggregating multiple audio samples in a single packet. The impact of Request To Send/Clear To Send (RTS/CTS) is analyzed by comparing the number of video flows supported by the network with RTS/CTS and without. RTS/CTS turns out to limit the performance of the network in terms of number of concurrent video flows. An indoor scenario is considered in [7], where a routing metric exploiting multiple radio devices is shown to achieve higher throughput than other metrics (such as those based on the shortest path algorithm).

The remainder of this paper is organized as follows. In Sec. II we describe the mesh networking paradigm and report on the current state-of-the-art in WMN deployments. In Sec. III we describe the experimental settings and the traffic patterns used for the performance measurements. Section IV reports the outcomes of the measurements and discusses the ability of current WMNs to support multimedia flows. Section V concludes the paper pointing out directions for future work.

II. Wireless Mesh Networks

A Wireless Mesh Network consists of several nodes, interconnected via wireless links (possibly using multiple radio technologies/interfaces [7]) to the Internet through one or multiple gateway(s). Communications take place by means of multihopping, in that the nodes in the network cooperate to forward packets (by means of store-and-forward operations) to/from the Internet from/to the end node.

Nodes in a WMN can play two different logical roles, i.e., mesh clients and mesh routers [1]. Mesh clients can be the source/destination of connections, while mesh routers are in charge of forwarding packets to and from the Internet. A single node can play both roles at the same time, as in standard ad hoc networking paradigms [3]. Multi-tier architectures can be envisaged [2], with mesh routers providing multihop backhaul connectivity to the Internet, while the clients act just as sources/destinations of Internet connections. It is worth stressing that, from our standpoint, WMNs are to be thought as *access network architectures*, and not as stand-alone “ad hoc” systems. Nonetheless, they share many features with conventional ad hoc networking paradigms. In particular, self-organization is expected to play a key role in mesh networking due to both (i) technical reasons, since it allows the deployment of unplanned networks while keeping at the same time backward compatibility with existing WLAN installations and (ii) economical reasons, since it helps in lowering the entrance barrier to the ISP market, providing opportunities for SMEs to deploy backhaul networks in an incremental fashion.

Depending on the hierarchy introduced by the differentiation of nodes functionalities, WMN architectures can be classified according to the following taxonomy [1]:

- *Infrastructure/Backbone WMNs.* In infrastructure/backbone WMNs, as depicted in Fig. 1, wireless routers realize a self-configuring and self-healing mesh backbone, providing the clients with the opportunity to connect to a remote Internet gateway. Typical applica-

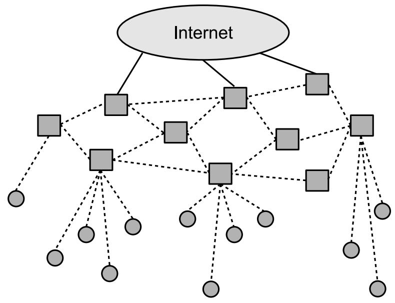

Fig. 1. Infrastructure/Backbone WMNs.

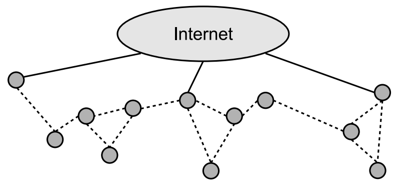

Fig. 2. Client WMNs.

tions of this architecture are in community/neighborhood networking and in wireless mesh ISPs, where mesh routers are placed on the roof and a local in-home distribution service (either wired or wireless) is added to provide end-user connectivity. Examples of such architecture include the MIT's Roofnet [8] and the (commercial) LocustWorld [9] deployments.
- *Client WMNs.* In client WMNs, sketched in Fig. 2, client nodes organize themselves into a flat architecture for providing Internet access by means of store-and-forward operations. This solution adapts well to extensions of indoor WLANs. On the other hand, it is not suitable for metropolitan-level networks due to the obvious scalability problems. The Microsoft's Mesh Connectivity Layer [10] fell into this category.
- *Hybrid WMNs.* Hybrid WMNs represent the combination of the two aforementioned solutions, as depicted in Fig. 3.

Besides, in terms of routing protocols, the most successful approach has been to re-use existing standards for ad hoc networks and adapt them to the peculiarities of the mesh environments [11]. The performance obtained by such systems are clearly far from optimal, and a lot of efforts are needed to enhance and optimize such solutions. In terms of protocol architectures, two solutions can be envisaged to forward and route packets on the mesh. In the first, the routing protocol is implemented directly at level three of the ISO/OSI stack, therefore (partially) modifying standard IP operations. In the second case, a 2.5-level routing protocol is provided, so that, to higher layer, the WMN appears like a LAN. The protocol stack of the two possible solutions is sketched in Fig. 4(a) and Fig. 4(b), respectively. The first choice provides more space for optimization and performance enhancements, but its implementation may not be trivial and may result platform-dependent. The second approach has the advantage of being transparent to standard networking stacks, so that it can be readily implemented over (virtually) any platform. On the other hand, it adds some overhead, thus lowering the network performance.

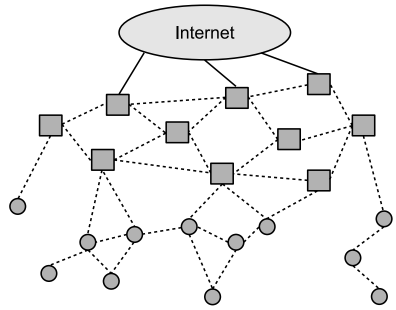

Fig. 3. Hybrid WMNs.

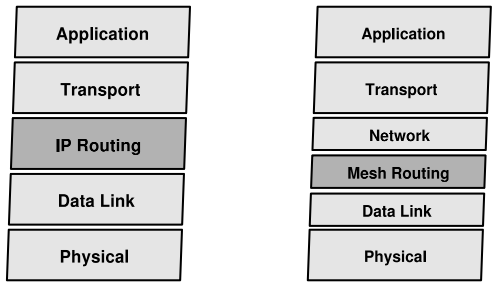

(a) L3 Routing (b) L2.5 Routing

Fig. 4. WMN Protocol Architectures.

III. System Configuration

A. Network Configuration

The experimental data has been collected exploiting a 7-nodes wireless testbed deployed in a typical office environment implementing a single-tier structure, as sketched in Fig. 5. Testbed's nodes are all Dell notebook model D600/D610/D810 equipped with a 1.86 GHz Intel Pentium M processor with 512 MB of memory. All nodes run Microsoft Windows XP Professional. Each node has a single Intel 2915ABG or a Dell 1470 Wireless adapter with RTC/CTS disabled. For the infrastructured test we used a Cisco Aironet 1200 Access Point (AP) [12] that supports both 802.11a and 802.11b/g operation mode. The AP is equipped with 2 omni-directional antennas with a gain of 2.14 dB. The default maximum output power of the access point is 50 mW. However, we decided to reduce this value to 20 mW (which is the maximum output power of our wireless adapters) in order to have the same operating conditions for both the infrastructured and the mesh modes.

During our measurements, functionalities provided by node number one are twofold. In the mesh scenario, it acts as gateway to the Internet, with the routing protocol running on it. In the infrastructured scenario, it is attached through an Ethernet connection to the AP. All measurement are run using IPv4 with statically assigned addresses and IEEE 802.11 operating in “g” mode. In order to increase the reliability of our results, we have exploited the AP’s site survey tool in order to detect the presence of interference caused by other 802.11 devices. The operating channel for both the AP and mesh scenarios has been chosen according to this analysis. Mesh connectivity is realized using the Microsoft Mesh Connectivity Layer [10].

The Mesh Connectivity Layer (MCL) is a loadable Microsoft Windows driver. It implements an interposition layer between layer 2 (the link layer) and layer 3 (the network layer) of the standard ISO/OSI model. It is sometimes referred to as layer 2.5. To the higher layers, MCL appears to be just another Ethernet link, albeit a virtual one. To the lower layers, MCL appears to be just another protocol running over the physical link. MCL routes using a modified version of DSR [13] called Link Quality Source Routing (LQSR) [7]. LQSR assigns a weight to each link. This weight is the expected amount of time it would take to successfully transmit a packet of some fixed size on that link. In addition, the channel, the bandwidth, and the loss rate are determined for every possible link. This information is sent to all the nodes. Based on this information, LQSR uses a routing metric called Weighted Cumulative Expected Transmission Time (WCETT) to define the best path for the transmission of data from a given source to a given destination.

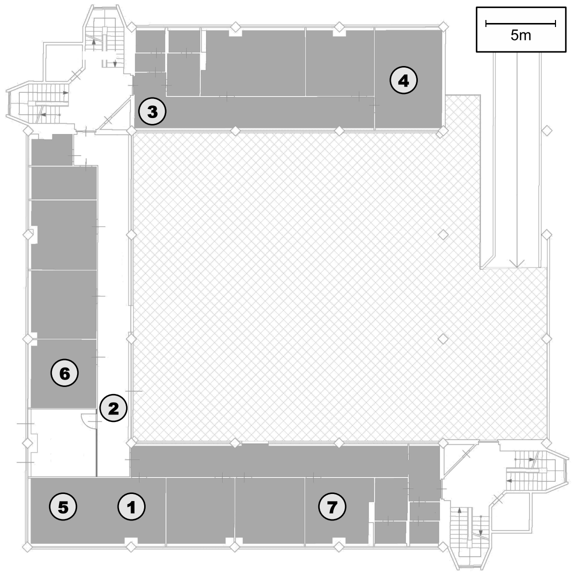

Fig. 5. Testbed planimetry. In the mesh scenario node number one acts as gateway, while in the infrastructured scenario it is attached through an Ethernet connection to an AP.

B. *Multimedia Traffic Patterns*

The experimentation had been performed by using synthetic traffic generated by means of the Distributed Internet Traffic Generator (D-ITG), a freely available software tool [14]. D-ITG can generate and inject different traffic patterns over TCP and/or UDP sockets. The traffic is then collected at the receiver side where suitable tools can provide a great variety of statistical analysis. By means of D-ITG it is possible to simulate many traffic scenarios originated by a large number of users and network devices, whereas other traffic generators have limited capabilities in terms of performance and range of source models.

Looking at multimedia communication, we focused on a video conference application due to: (i) its widespread use (e.g. Skype 2.0) and (ii) its strong requirements in terms of Quality-of-Service. Actually, we chose such a real-time service since it is one of the most demanding in terms of loss and delay constraints. Therefore, it is particularly suited to stress the network, especially when dealing with mesh structures, where multihop communication could introduce unacceptable delays.

We have emulated each video conference service by continuously transmitting two UDP packet flows

	Video (H.264)	**Audio (G.729.3)**
Rate (Packets/sec)	10	33
Payload length (Bytes)	800	42

TABLE I

FLOWS CHARACTERIZING THE VIDEO CONFERENCE TRAFFIC PATTERN (UDP)

	Best Effort (FTP)
Rate (Packets/sec)	2000
Payload length (Bytes)	1240

TABLE II

FLOW CHARACTERIZING THE BEST-EFFORT TRAFFIC PATTERN (TCP)

at the same time: a voice stream and a video stream. For the former one, we have considered the G.729.3 codec [15], a worldwide used speech codec for VoIP applications, with each packet containing three voice samples and without Voice Activity Detection. The video stream has been generated according to the recently approved H.264 standard [16], well-known for its compression performance. We assumed that a good video quality can be attained by coding the video using 10 frames/sec [17], in such a way that one frame can be carried in one packet. The properties of both flows are summarized in Table I. As it can be red from Table I, each video conference application requires 75 kbit/sec including RTP headers. On the other hand, best effort traffic (in our case persistent TCP connections) is modeled considering a TCP socket working in saturation regime, according to the parameters reported in Table II. In order to collect reliable measure of delays, before each experiment we synchronized each node with a common reference using NTP [18].

IV. PERFORMANCE MEASUREMENTS

In this section we report the outcomes of some experimental tests run with the equipment and settings described in Sec. III-A. As said before we compare the results obtained exploiting our mesh architecture, with the ones achieved using the infrastractured scenario. Due to the preliminary nature of this work, the tests reported refer to downlink traffic only. The nodes are activated according to the numbering in Fig. 5, so that when N flows are active hosts $2, 3, \dots, N+1$ are downloading from host 1.

As outlined in Sec. III-B two traffic patterns are considered. First, we will focus on data traffic only, where persistent TCP connections are emulated. In this case, we will consider the average throughput experienced by each node. We will consider as performance metrics both the mean aggregated throughput (which, roughly speaking, shows the ability of the system to efficiently use the available bandwidth) and the fairness, defined according to the classical Jain's index [19]:

$$f = \frac{(\sum_{i=1}^{N} x_i)^2}{N \sum_{i=1}^{N} x_i^2}, \qquad (1)$$

where x_i denotes the average throughput experienced by node $(i+1)$. The fairness index f is an indicator of how fairly the overall bandwidth is shared among competing connections. In the infrastructured mode, this depends mainly on the different channel conditions encountered on the links, exacerbated by the dynamics of TCP's congestion control mechanism, which has the overall effect of penalizing the hosts far away from the AP. On one hand we can expect the mesh architecture to provide a higher level of fairness, in that hosts far away from the AP could exploit relays to enhance their throughput. On the other hand, in the mesh case, links may be shared by multiple connections, giving rise to problems of buffer overflows with possibly negative effects on the overall performance.

The results for the aggregated throughput and the fairness index are plotted in Fig. 6 and Fig. 7, respectively. The infrastructured mode provides better performance in terms of aggregated throughput. However, the higher bandwidth utilization is achieved at the expenses of nodes with poor channel conditions. This is shown in Fig. 7, where mesh architecture performs slightly better than infrastructure mode in terms of fairness. In Table III, we reported the average throughput experienced by each node for the case of six best-effort flows. There is a higher variance for the AP with respect to the mesh architecture, thus confirming the results in terms of fairness.

The other tests refer to video conference applications, modeled according to the parameters detailed in Sec. III-B. In this case we look at packet delays and losses as the two main QoS metrics. We expect the mean packet delay to be higher in the case of mesh architecture, due to the processing and

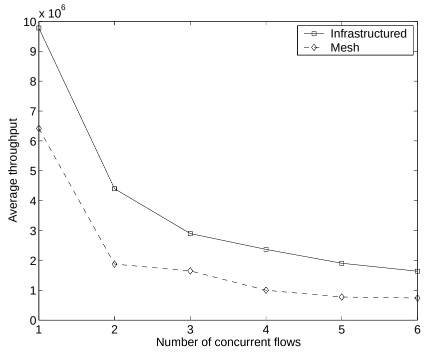

Fig. 6. TCP average throughput versus number of concurrent flows using the two different network architectures.

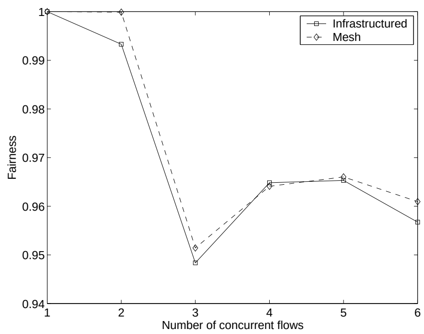

Fig. 7. TCP fairness versus number of concurrent flows using the two different network architectures.

Node Id	Infrastructured Average bitrate (kbit/s)	Mesh Average bitrate (kbit/s)
2	1967	783
3	1657	680
4	1062	1065
5	2018	605
6	1355	640
7	1966	773

TABLE III

AVERAGE THROUGHPUT FOR SIX BEST-EFFORT FLOWS USING THE TWO DIFFERENT NETWORK ARCHITECTURES

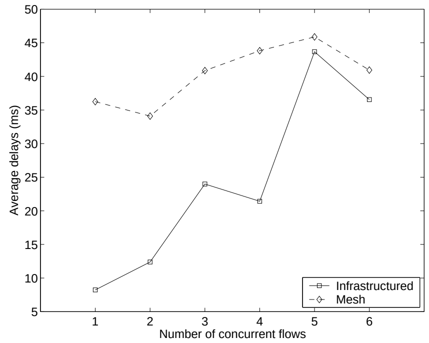

Fig. 8. Average delays versus number of concurrent flows using the two different network architectures.

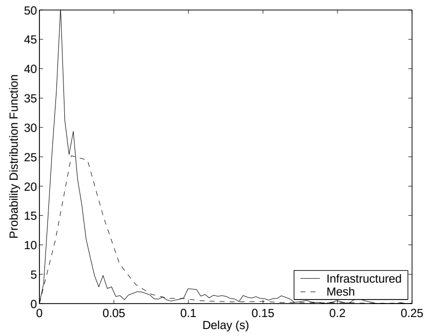

Fig. 9. Packet Delays's Probability for six concurrent Multimedia Flows using the two different network architectures.

buffering at each node necessary to perform store-and-forward operations. This is confirmed from the results plotted in Fig. 8, which reports the average delay vs. number of concurrent multimedia flows. On the other hand, the effect of such operations on the detailed statistics, i.e., Probability Distribution Function (PDF), is hardly predictable. Indeed, on one hand we expect the buffers at intermediate nodes to act as "integrators", smoothing the delay PDF. On the other hand, the buffering could introduce unpredictable delays, worsening the overall performance. The results are reported, in terms of delay PDF in Fig. 9. As it may be seen, the mesh architecture presents a smoother delay PDF. This is generally acknowledged to have a beneficial effect on multimedia flows, in that it facilitates the design and dimensioning of playout buffers. In Table IV we reported the sample mean and sample standard deviation for the packet delay, for both infrastructured and mesh mode, in the case of six concurrent multimedia flows. Finally, Fig. 10 reports the mean packet loss rate for both considered architecture. It can be seen that the mesh architecture presents a lower packet loss rate than the infrastructured architecture when the number of concurrent multimedia flow is rather large. This suggests that the mesh architecture presents better scalability properties than conventional access-points based WLANs.

V. CONCLUSIONS

In this paper, we have reported some preliminary results on the performance of multimedia and data flows in wireless mesh networks. Such results have

	Sample Average Delay (ms)		Sample Delay Standard Deviation (ms)	
Node	**Infrastrucured**	**Mesh**	**Infrastrucured**	**Mesh**
2	20	55	14	207
3	26	38	7	30
4	35	27	39	37
5	17	64	21	149
6	107	32	80	55
7	9	24	6	20

TABLE IV

AVERAGE DELAYS FOR SIX MULTIMEDIA FLOWS USING THE TWO DIFFERENT NETWORK ARCHITECTURES

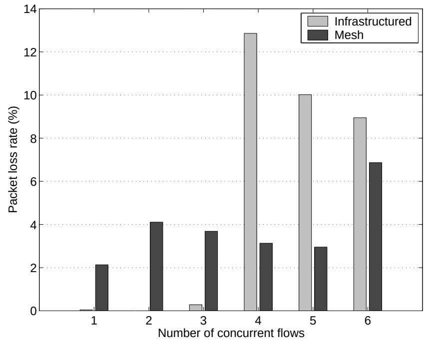

Fig. 10. Packet loss versus number of concurrent flows using the two different network architectures.

been obtained by emulating data and multimedia flows over a WMN testbed. The results, compared with the conventional infrastructured architecture, based on the use of access points, have shown that mesh architectures may represent a feasible solution for providing indoor broadband support to multimedia flows. In particular, mesh architectures have shown to be able to attain a better fairness in bandwidth sharing with respect to conventional star-like topologies. Further, the effect of buffering at intermediate nodes is to smooth the packet delay PDF, with a beneficial impact on the performance of jitter-sensitive multimedia flows.

The results presented in this paper, while promising, represent just a first step toward the understanding of the real capabilities of mesh-based architectures. In particular, it is worth recalling that our measurements were based on standard off-the-shelves devices and a freely available software package for mesh networking. We believe that the performance measures we obtained could be highly improved by optimizing the protocols for the application scenarios we considered. Our future plans include the extension of the testbed size (in terms of number of connected devices and network coverage) to study the scalability of mesh architectures and the optimization of the protocol stack for enhancing system performance.

REFERENCES

[1] I. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: a survey," *Elsevier Computer Networks*, vol. 47, no. 4, pp. 445 – 487, Mar. 2005.

[2] R. Bruno, M. Conti, and E. Gregori, "Mesh Networks: Commodity Multihop Ad Hoc Networks," *IEEE Communications Magazine*, vol. 43, no. 3, pp. 123 – 131, Mar. 2005.

[3] I. Chlamtac, M. Conti, and J. Liu, "Mobile ad hoc networking: imperatives and challenges," *Elsevier Ad Hoc Networks*, vol. 1, no. 1, pp. 13 – 64, Jan. 2003.

[4] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, pp. 121 – 132, Oct. 2004.

[5] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," in *Proc. of ACM MOBICOM*, Cologne, Germany, 2005.

[6] Y. Sun, I. Sheriff, E. M. Belding-Royer, and K. C. Almeroth, "An experimental study of multimedia traffic performance in mesh networks," in *Proc. of WitMeMo*, Seattle, Washington, 2005.

[7] R. Draves, J. Padhye, and B. Zill, "Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks," in *Proc. of ACM MOBICOM*, Philadelphia, Pennsylvania, USA, 2004.

[8] MIT, "Roofnet." [Online]. Available: http://pdos.csail.mit.edu/roofnet/

[9] "LocustWorld." [Online]. Available: http://www.locustworld.com/

[10] Microsoft, "MCL." [Online]. Available: http://research.microsoft.com/mesh/

[11] R. Draves, J. Padhye, and B. Zill, "Comparison of Routing Metrics for Static Multi-Hop Wireless Networks," in *Proc. of ACM SIGCOMM*, Portland, Oregon, USA, 2004.

[12] "Cisco Aironet 1200 Series Access Point." [Online]. Available: http://www.cisco.com/en/US/products/hw/wireless/ps430/index.html

[13] IETF, "The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR)." [Online]. Available: http://www.ietf.org/internet-drafts/draft-ietf-manet-dsr-10.txt

[14] "D-ITG, Distributed Internet Traffic Generator." [Online]. Available: http://www.grid.unina.it/software/ITG/

[15] "ITU-T Recommendation G.729 Annex B, A silence compression scheme for G.729 optimized for terminals conforming to Recommendation V.70," Nov. 1996.

[16] Joint Video Team (JVT) of ISO/IEC MPEG& ITU-T VCEG, "Joint Final Committee Draft (JFCD) of Joint Video Specification (ITU-T Rec. H.264 ISO/IEC 14496-10 AVC)," Jul. 2002.

[17] P. Brun, G. Hauske, and T. Stockhammer, "Subjective Assessment of H.264/AVC Video for Low-Bitrate Multimedia Messaging Services," in *Proc. of IEEE Image Processing*, Singapore, 2004.

[18] "Simple Network Time Protocol (SNTP) Version 4." [Online]. Available: http://www.apps.ietf.org/rfc/rfc2030.html

[19] R. Jain, D. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," DEC Research Report, Tech. Rep., 1984.