

MULTIMEDIA SUPPORT IN IEEE 802.11 BASED WIRELESS MESH NETWORKS

Karina Gomez, Tinku Rasheed, and Roberto Riggio

April 23, 2009

Create-Net Technical Report N. 200800024

CREATE-NET
Via Solteri 38
38100 – Trento (Italy)
www.create-net.it

Abstract

The extended radio coverage provided by wireless mesh networking (WMN) technologies represents a key advantage compared to the last mile solutions based on standard IEEE 802.11 hot spots. Nevertheless, modern requirements to wireless connectivity include mandatory QoS guarantees for a wide set of real-time applications. In this paper, we first propose a methodology for evaluating multimedia applications over real world WMN deployments and then we report the results of an extensive measurement campaign performed exploiting IEEE 802.11-based WMN testbed. Later, we evaluate the QoS performance of the reference multimedia applications exploiting three different link scheduling disciplines. Results show that, using opportunistic scheduling techniques capable of providing performance isolation among competing flows can significantly improve system capacity.

KEYWORDS: wireless networks, IEEE 802.11, mesh architecture, multimedia applications, testbed, roofnet

1 Introduction

In the last years, considerable efforts from both academic and industrial worlds have been devoted to the wireless mesh networking paradigm [1]. An important role in this context is played by multimedia traffic, which also represents one of the most resource consuming transmissions. Companies and providers from all business sectors are discovering the marketing power of video streaming communication that reaches thousands of users at their home or work. However, supporting multimedia applications such as VoIP, conference-call, video streaming etc. over an IEEE 802.11-based Wireless Mesh Networks (WMNs) is a challenging task. To the best of our knowledge, even though there has been a plethora of efforts to measure the performance of WMN in real settings, there are very few works that evaluate and characterize the performance of real-time traffic in WMNs.

Ultimately, the economic convenience of a WMN is measured by the number of customers that the provider can sustain for a given network deployment. One major challenge for the research in this field is indeed to increase the scalability of WMNs. The two aspects are correlated since, the larger the area, the larger the number of users that can access the network. But, little efforts are known to be dedicated to investigate efficient techniques providing and assessing QoS in WMNs. Vendors selling wireless mesh solutions do implement some form of QoS policies, but they are obviously reluctant to release those information. Hence, research in WMNs lack from a comprehensive QoS perspective.

The contribution of this work is twofold. First we propose a methodology for evaluating multimedia applications over real world WMN deployments. Second, based on the defined methodology, we report the results of an extensive measurement campaign performed exploiting an IEEE 802.11-based WMN testbed deployed in a typical office environment. We focus our research on three mainstream multimedia applications: VoIP, Video Conference, and Video Streaming. Two single-hop star-shaped network topologies (with symmetric and asymmetric links) and a multi-hop string topology have been exploited in order to provide a comprehensive evaluation of the testbed's performances.

Finally, we evaluate the QoS performance of the multimedia applications over the WMN by performing measurements with three different link scheduling disci-

plines: First-Come First-Serve (FCFS), Deficit Round Robin (DRR), and Airtime Deficit Round Robin (ADRR). The latter discipline, proposed in [2] exploits measurable routing metrics typically available in WMNs in order to compute the optimal schedule list and can be readily implemented using off-the-shelf components. The outcome of the measurements campaign show that the ADRR scheduler is capable of providing performance isolation in IEEE 802.11-based WMNs especially in the case of UDP traffic. As expected, the performance figures for the reference multimedia applications degrades when either the number of hops or the number of concurrent flows increases. However, this performance degradation is shown to be smoothed by the ADRR scheduling discipline.

The rest of the paper is structured as follows. In Sec. 2 we discuss some related works. Section 3 sketches the testbed's architecture. The evaluation methodology including traffic patterns, network topologies and performance metrics are discussed in Sec. 4. Section 5 describes outcomes of our measurements campaign and Sec. 6 concludes the paper.

2 Related Work

Albeit WMNs are a relatively new research area, there is considerable prior work devoted to characterize their performances. Several recent studies have evaluated the performance of WMN in real environments with dedicated testbeds [3, 4, 5, 6]. The work in [4] studies the performance of a single radio IEEE 802.11-based WMN testbed in both outdoor and laboratory environments. The authors characterize the performance of the network using different applications (VoIP and elastic traffic) and discuss key challenges such as bandwidth allocation, fairness and the hidden node phenomena.

The voice capacity of IEEE 802.11 clusters connected to an IEEE 802.16 backhaul is studied in [5]. Results show that the performance bottleneck lies in the WLAN segment. The authors propose a multiplex-multicast scheme capable of doubling the voice capacity of the system. The multiplexer combines multiple downstream VoIP packets into a single multiplexed packet and then multicasts the multiplexed packet to the wireless stations. In [7], experimental investigation to improve quality of VoIP in WMNs is presented. A MAC-waiting based distributed packet aggregation

strategy is implemented and evaluated using an experimental testbed.

Multiple aspects of an outdoor WMN deployment are analyzed in [8]. In particular the paper investigates the impact of node density on connectivity and throughput and the characteristics of the links selected by the routing protocols. Similarly, experimental characterization of mesh networks in a home environment is reported in [9]. The authors observe significant benefits of mesh network topologies within a home, reporting performance improvements when used alone or in combination with other wireless configuration parameters.

The causes of packet loss in a large outdoor WMN is analyzed in [6]. The authors conclude that the loss rate distribution is substantially uniform across the whole range of loss rates and that a large number of links are characterized by intermediate loss rates. Such links can greatly affect the performances of the whole network with special regard to the nodes experiencing good channel conditions.

We differ from the aforementioned works in that we aim at assessing the capability of current WMNs to support broadband multimedia services like VoIP, Video Conference, and Video Streaming. As the demand for rich-media content continues to increase, such applications are certainly important to be analyzed in WMN environments. Our measurements campaign has been carried out exploiting an IEEE 802.11-based WMN deployed in a typical floor office environment.

3 Testbed characteristics

3.1 Network architecture

The WING testbed¹ is an experimental IEEE 802.11 wireless mesh network. The network consist of 10 nodes deployed in a typical office environment and implementing a flat network architecture. Testbed's nodes are based on the PCEngines ALIX processor board [10]. Each node is equipped with a 500MHz CPU, 256MB of RAM, and two IEEE 802.11a/b/g wireless interfaces with RTC/CTS disabled. The first interface builds and maintains the wireless backhaul, while the second interface is configured in *Master mode* providing a standard IEEE 802.11a Access Point.

¹On line resources including traffic trace files and processing scripts available at <http://www.wing-project.org/>.

Our WMN is based on the Roofnet platform developed by the MIT and deployed in Cambridge, MA, USA [8]. Roofnet routes packets using a DSR-like routing protocol called *SrcRR* exploiting the Estimated Transmission Time (ETT) as routing metric [11] and optimized for network scalability and throughput rather than for supporting mobility. The ETT metric aims at estimating the amount of time required to transmit an unicast packet over a wireless link (including re-transmission). The ETT metric is computed as follows:

$$M_{\text{ETT}} = \frac{1}{P_{\text{ACK}}R}$$

Where R is an estimate of the highest effective throughput achievable in the forward direction, and P_{ACK} is the delivery probability of the ACK signal in the reverse direction (d_{rev}). Being d_{fwd} the link delivery probability in the forward direction, and x the transmission rate in Mb/s, we can write:

$$R = \max(r_1, r_2, r_{5.5}, r_{11}); \quad r_x = d_{\text{fwd}}x$$

In order to compute the forward (d_{fwd}) and reverse (d_{rev}) link delivery ratios each node periodically broadcast a sequence of five probes: one short probe aimed at modeling the ACK transmission and one long probe for each available transmission rate (1, 2, 5.5, 11 Mb/s)². Each node keeps track of the number of probes received during an observation window W . At any time, d_{rev} is then given by:

$$d_{\text{rev}}(t) = \frac{\text{count}(t - W, t)}{w/\tau}$$

Note that $\text{count}(t - W, t)$ is the number of probes received during the observation window W and w/τ is the number of probes that should have been received. Finally each probe sent by a node contains the number of probe packets received by the same node from all its neighbors during the last observation window. Such a design choice allows the receiver to compute the forward delivery ratio d_{fwd} toward the node from which the probe was originated.

Routing is implemented using the Click router [12]. A Click router is built by assembling several packet processing modules, called elements, forming a directed graph. Each element is in charge of a specific function such as packet classification, queuing, and interfacing with networking devices. We extended the default Roofnet configuration by implementing additional elements responsible for packet scheduling.

²Broadcast frame are not acknowledged nor retransmitted by IEEE 802.11 devices.

3.2 Opportunistic scheduling

WMNs are known to be particularly susceptible to the “*IEEE 802.11 performance anomaly*” [13], refer to the sudden performance drop that occurs when nodes transmitting at low bit-rates due to poor channel conditions capture the wireless medium for longer periods of time at the expense of the nodes transmitting at higher bit-rates.

In such a context, *airtime* fairness can be provided by the Airtime Deficit Round Robin (ADRR) link scheduling discipline. ADRR enhances the Deficit Round Robin (DRR) scheduling discipline by taking into account the channel quality which in time prevents a node affected by high packet loss from monopolizing the wireless channels thus lowering the performance of the whole system. The ADRR link scheduling discipline, described in detail in [2], does not requires modification to the standard IEEE 802.11 MAC and can be readily implemented using off-the-shelf components.

Figure 1 depicts the main building block implementing the ADRR scheduling policy. The scheduler maintains a list of currently backlogged queues. Incoming data frames are first classified according to their next hop and then fed to the corresponding queue. If such a queue does not yet exist, it is created dynamically by the scheduler. Each queue is associated with a counter, called *Deficit Counter*, that indicates the amount of resources the link can use in a round. At each round, the deficit counter of the currently visited queue is increased by a fixed quantity called *Quantum*. The ADRR scheduler only serves packets whose expected transmission time is smaller than the deficit counter. Let L_{Probe} be the size of the probe used to compute d_{fwd} , the expected transmission airtime $TX_{AIRTIME}$ for a packet S bytes long is then given by:

$$TX_{AIRTIME} = M_{ETT} \frac{S}{L_{Probe}}$$

After a packet is sent the deficit counter is decreased by $TX_{AIRTIME}$. A frame whose transmission time exceed the deficit counter is held back until the next round. Empty queues are removed from the list and their deficit counter is set to zero.




Figure 1: Airtime DRR Scheduler's block diagram

4 Multimedia applications in WMNs

4.1 Traffic patterns

Looking at multimedia communications, we focused our attention on three reference applications: VoIP, Video Conference, and Video Streaming. The reasons for such a choice lie in the widespread use of these applications (e.g. Skype, YouTube) and in their strict requirements in terms of QoS. FTP-like traffic has been used to assess the capability of the network to support best effort traffic.

The reference multimedia services have been emulated using JTG [14] a freely available software tool capable of injecting different traffic patterns over UDP and TCP sockets. JTG can read the information about packet transmission intervals and packet sizes from files, allowing us to create an exact duplicate of a trace starting from a pre-recorder stream. Both uplink and downlink traffic has been considered in the VoIP and in the Video Conference scenarios. In order to collect reliable measures of delays, before each experiment we synchronized each node with a common reference using NTP [15]. The following traffic patterns have been used during our measurements campaign:

- *VoIP*. We have emulated each VoIP call as two UDP streams (one bi-directional voice channel) modeled according to the parameters of the G.729.3 codec [16]. The G.729.3 VoIP codec generates 33 pkts/s, each packet contains 3 voice samples (10 bytes each) producing a final bit-rate of 8 kbits/s.
- *Video Conference*. The video conference has been emulated as four UDP streams (one bi-directional voice channel plus one bi-directional video channel).

The voice channel has been modeled according to the parameters of the G.729.3 codec while the video channel exploited a trace file generated from a 15 minutes long RealVideo encoded video sequence. The video stream average and peak bitrates are respectively 153 kbits/s and 158 kbits/s.

- *Video Streaming.* Video streaming has been emulated as a single UDP stream exploiting a trace file generated from a 10 minutes long MPEG-4 encoded video sequence. The video stream average and peak bitrates are respectively 572 kbits/s 771 kbits/s.
- *FTP.* Best effort traffic (in our case persistent TCP connections) is modeled as TCP socket working in saturation regime.

4.2 Network topologies

In order to provide a comprehensive evaluation of the testbed performance, three different network topologies have been exploited during our measurements campaign:

- *Star-shaped (symmetric links).* In this scenario a star-shaped topology (see Fig. 2a) has been exploited in order to model the gateways' operating conditions.
- *Star-shaped (asymmetric links).* Measurements have been carried out exploiting the previously described star-shaped topology. However, albeit always in radio range, the mesh nodes have been deployed in three different configurations. More precisely, while node number 2 and 3 are kept close to the gateway (node number 1), while node number 4 is positioned in such a way to experience channel condition raging from *Good* to *Poor* with an intermediate *Medium* quality.
- *String.* Especially when used as last mile solution, WMNs are characterized by a tree topology where different flows interact only with the aggregation points. In such a scenario several string topologies (see Fig. 2b) are exploited in order to model the performance of a single branch neglecting the interference of the remaining tree.




Figure 2: Network topologies exploited during the measurement campaign.

For all measurements involving the single-hop star-shaped network topology flows are activated according to the node numbering reported in Fig. 2b; so when N flows are active nodes 2, 3, $N + 1$ are communicating with node 1. Moreover, unless otherwise specified, the performance figures discussed in this paper always refer to the ongoing communication between node 1 and node 3. For the DRR scheduling discipline the *quanta* has been set to 1500 bytes which is the MTU (Maximum Transmission Unit) supported by an Ethernet LAN.

4.3 Performance metrics

4.3.1 VoIP

We resort to the E-Model [17] as an objective method to evaluate VoIP quality. The outcome of an E-Model evaluation is the R-factor (R). The R-factor is a numerical measure of voice quality, ranging from 0 to 100, with 70 as lower bound for a VoIP call of acceptable quality.

4.3.2 Video Conference & Video Streaming

Video Conference applications imply full-duplex services with the added requirement that the audio and video must be synchronized within certain limits to provide *lip-synch*. In order to simplify the analysis methodology we decided to ignore synchronization issues and to evaluate the performances of the audio and video flows separately. Video streaming applications are expected to deliver better quality than both VoIP and Video Conference applications leading to tighter requirements in terms of packet loss. Performance targets for Video Conferencing applications have been suggested by ITU-T in [18] where the value of 150 ms for end-to-end delay and 1% of packet loss are considered the threshold for an acceptable user experience.




Figure 3: Performance figures for the VoIP scenario exploiting different scheduling disciplines. This graph reports the results for the star topology with up to 9 concurrent VoIP sessions.

5 Discussion

5.1 VoIP

5.1.1 Star-shaped topology (symmetric links)

In this scenario we evaluated the number of concurrent VoIP flows that can be sustained, i.e. the voice capacity of the system. Results, reported in Fig. 3, show that the ADRR scheduling discipline is capable of providing a significant performance boost with respect to both the FCFS and the DRR disciplines even when all node are experiencing good channel conditions. We postulate that the ADRR scheduler is capable of exploiting local channel fluctuations by opportunistically allocating more airtime when and where the channel is strong. Such considerations are supported by the theoretical findings in [19].

5.1.2 String topology

The results of this measurements set are reported in Fig. 4 as empirical cumulative distribution function of the packet's delays for an increasing number of hops. As expected, as the number of hops increases, the end-to-end delays performance progressively degrades. Moreover, as it can be seen from the Fig. 5, the performances of the system for a three hops topology drop suddenly when the number of concurrent calls is higher than 3.




Figure 4: Empirical CDF for the VoIP scenario with up to 3 concurrent VoIP sessions (string topology).




Figure 5: Performance figures for the VoIP scenario with an increasing number of calls. This graph reports the results for the three-hops string topology.

5.2 Video Conference

5.2.1 Star-shaped topology (symmetric links)

Results, reported in Fig. 6, show no significant performance improvement when either the DRR or the ADRR scheduler are used. As can be seen from the figure, the system is capable of supporting up to 4 concurrent conference calls before both the delay and packet loss figures drop to unacceptable levels.

5.2.2 String topology

The results of this measurements set are reported in Fig. 7 as empirical cumulative distribution function (CDF) of the packet's delays for an increasing number of hops. As expected, as the number of hops increases, the end-to-end delays performance for both the video and the audio flows progressively degrades. The system capacity for the two and the three hops string topologies is respectively 3 and 2 concurrent video conference sessions.

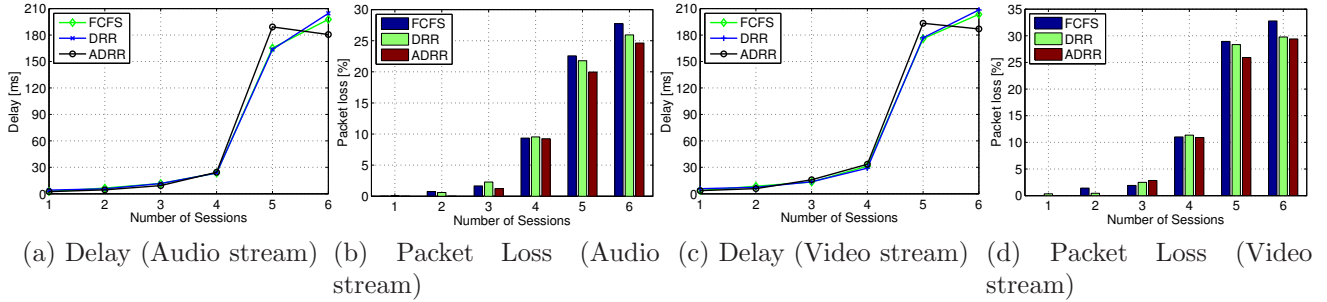


Figure 6: Performance figures for the Conference Call scenario exploiting different scheduling disciplines. This graph reports the results for the star topology with up to 6 concurrent Conference Calls.




Figure 7: Empirical CDF for the Video Conference scenario with up to 9 concurrent sessions (string topology).

5.3 Video Streaming

5.3.1 Star-shaped topology

Results show no significant performance improvement using the ADRR scheduler in both symmetric and asymmetric links scenarios. However, as it can be seen from the Fig. 8a, the performances of the system drop suddenly when the number of concurrent video streams is higher than 7. As previously pointed out, due to the non-conversational nature of the service, one-way delay requirements are not very stringent, hence we ignored them.

5.3.2 String topology

The results of this measurements set are reported in Fig. 8b and 8c. As expected the packet-loss increases with the number of hops. The system's capacity bounds lie at 5 and 4 concurrent streams in the case of respectively a two-hops and a three-hops

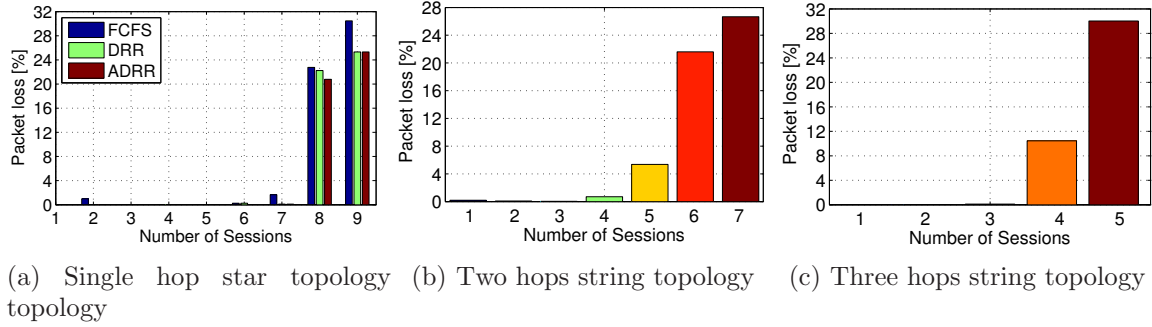


Figure 8: Performance figures for the Video Streaming scenario with an increasing number of concurrent sessions.




Figure 9: Average throughput for the FTP scenario exploiting different scheduling disciplines. This graph reports the results for the star topology with 3 concurrent flows. Node number 4 is positioned in such a way to experience variable channel condition.

string topology.

5.4 FTP

5.4.1 Star-shaped topology

Results, reported in Fig. 9a, 9b and 9c, show that the ADRR scheduler does not increase significantly the aggregated throughput of the system. We postulate that this behavior is due to a mutual interference between the ADRR adaptive scheduling policy and the TCP's congestion control feedback loop.




Figure 10: Average FTP throughput exploiting different scheduling disciplines for an increasing number of hops.

5.4.2 String topology

The results of this measurements set are reported in Fig. 10. As expected the end-to-end system throughput decreases with the number of hops. Although at any time no more than one outgoing link exists for each node, we decided to repeat this set of measurements exploiting different schedulers in order to prove that the ADRR computational overhead does not impact the performance of the system proving its lightweight design.

6 Conclusions

In this paper we proposed an evaluation methodology for assessing multimedia applications performance in IEEE 802.11-based WMNs. The results attainable using the proposed methodology can be exploited by network designers to devise innovative resource management techniques and accurate design guidelines (e.g. users to gateways ratio, number of radio interfaces, etc) for WMNs. The results of such an extensive measurement campaign allowed us to devise an opportunistic scheduling algorithm capable of improving the capacity of the system by providing performance isolation among competing flows.

Among the various possible research directions to be pursued to extend the current work, the most promising is to better evaluate the performance limits of our WMN testbed allowing us to fully characterize its scalability using real-world applications. Moreover, we plan to extend our evaluation methodology including additional reference applications such as HTTP and online gaming. Lastly, we are currently investigating the interference between the TCP congestion control

mechanism and the ADRR link scheduling discipline.

References

- [1] “Muniwireless.” [Online]. Available: <http://www.muniwireless.com/>
- [2] R. Riggio, D. Miorandi, and I. Chlamtac, “Airtime deficit round robin (ADRR) packet scheduling algorithm,” in *Proc. of IEEE MeshTech*, Atlanta, GA, USA, Oct. 2008.
- [3] V. Chavoutier, D. Maniezzo, C. Palazzi, and M. Gerla, “Multimedia over wireless mesh networks: Results from a real testbed evaluation,” in *Proc. of MED-HOCNET*, Corfu, Greece, Jun. 2007.
- [4] A. Arjona, C. Westphal, J. Manner, A. Ylä-Jääski, and S. Takala, “Can the current generation of wireless mesh networks compete with cellular voice?” *Elsevier Computer Communications*, vol. 31, no. 8, pp. 1564 – 1578, May 2008.
- [5] P. C. Ng, S. C. Liew, and C. Lin, “Voice over wireless LAN via IEEE 802.16 wireless MAN and IEEE 802.11 wireless distribution system,” in *Proc. of IEEE WIRELESSCOM*, Hawaii, USA, 2005.
- [6] 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.
- [7] D. Nicolescu, S. Ganguly, K. Kim, and R. Izmailov, “Performance of VoIP in a 802.11 Wireless Mesh Network,” in *Proc. of IEEE INFCOM*, Barcelona, Spain, Apr. 2006.
- [8] 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.
- [9] K. Papagiannaki, M. Yarvis, and S. Conner, “Experimental characterization of home wireless networks and design implications,” in *Proc. of IEEE INFCOM*, Barcelona, Spain, Apr. 2006.

-
- [10] “PC Engines.” [Online]. Available: <http://www.pcengines.ch/>
- [11] 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.
- [12] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, “The Click modular router,” *ACM Transaction on Computer System*, vol. 18, no. 3, pp. 263 – 297, Aug. 2000.
- [13] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, “Performance anomaly of 802.11b,” in *Proc. of IEEE INFOCOM*, San Francisco, California, USA, 2003.
- [14] “JTG.” [Online]. Available: <https://hoslab.cs.helsinki.fi/savane/projects/jtg/>
- [15] D. Mills, “Simple Network Time Protocol (SNTP) Version 4,” IETF RFC 4330, Jan. 2006, <http://www.ietf.org/rfc/rfc4330.txt>.
- [16] “Coding of speech at 8 kbit/s using conjugate-structure algebraic-code-excited linear prediction (CS-ACELP),” ITU-T Recommendation G.729, Jun. 2007.
- [17] “The E-model, a computational model for use in transmission planning,” ITU-T Recommendation G.107, Nov. 2005.
- [18] “End-user multimedia qos categories,” ITU-T Recommendation G.1010, Nov. 2001.
- [19] D. Tse and S. Hanly, “Multi-access fading channels: Part i: Polymatroid structure, optimal resource allocation and throughput capacities,” *IEEE Transactions on Information Theory*, vol. 44, no. 7, pp. 2796–2815, Nov. 1998.