Large Scale MANET Emulations using U.S. Army Waveforms with Application: VoIP

Brian J. Henz
U.S. Army Research Laboratory
APG, MD 21005

Travis Parker
ICF JASI
APG, MD 21005

David Richie
Browndeer Technology
Forest Hill, MD 21050

Lisa Marvel
U.S. Army Research Laboratory
APG, MD 21005

Abstract—**Large scale experimentation and analysis of mobile ad-hoc networks (MANETs) is an expensive and time consuming task. Even with the best planning, the environment at the time of the experiment is unpredictable, making large scale controlled experiments difficult to impossible to perform. At the U.S. Army Research Laboratory (ARL) we are using the EMANE (Extensible Mobile Ad-hoc Network Emulation) environment to supplement and extend live experiments in a controlled environment. Our emulation environment includes locally developed tools for real-time path loss calculations using the Longley-Rice irregular terrain model implemented on GPGPUs (General Purpose Graphics Processing Units). This real-time path loss calculation capability allows for the integration of virtual devices with live experiments and forces modeling and simulations. Live experiment integration is identified as one method to increase the perceived scale of the MANET to live participants and increase the repeatability of the experiment. This emulation system is currently being scaled to 1000s of CPU cores and emulated radios through the use of high performance computing assets and approaches. We will present this emulation platform in the context of the performance analysis of a MANET application, namely VoIP (Voice over IP).**

Keywords: Mobile Ad-Hoc Network, Emulation, GPGPU, RF Propagation

I. Introduction

In mobile networks, access points can move and coverage may vary widely in a region. At times, the access points will cluster together, leaving parts of the map with sparse coverage and parts with compromised service due to competition for available bandwidth as channel subscriptions become saturated. On the battlefield, relocating receivers to areas with better coverage may not be an option and calls have been made for tools to optimize and plan prior to mission as much as possible. Consider the problem of combat in an urban area. Narrow streets and buildings with metal roofs and reinforced concrete walls may interfere with radio reception and access to surveillance data, yet that surveillance information may be key to locating enemy combatants before they inflict casualties on friendly forces.

The Mobile Network Modeling Institute (MNMI) was established in FY2007 to exploit High Performance Computing (HPC) resources through the development of computational software. Thus enabling the Department of Defense (DoD) to design, test, and optimize networks at sufficient levels of fidelity and with sufficient speed to understand the behaviors of network centric warfare technologies in the full range of conditions in which they will be deployed. Operational goals include the development of scalable computational modeling tools for simulations and emulations, the ability to understand a priori the performance of proposed radio waveforms in the field, and to optimize the network for U.S. Army warfighters. The focus here will be on the HPC developments for MANET emulation including the use of general purpose graphics processing units (GPGPU) for real-time RF path loss predictions.

The scale and complexity of mobile ad hoc networks to the DoD is unique, and to the Army in particular as a mobile fighting force. The military is rapidly becoming a network-centric force, with substantial access to sensor-derived surveillance information as well as increasingly complicated application layer running over many different devices. This introduces significant advantages to the warfighter, but also brings in new dependencies and new risks from the rapid change in configurations of the MANETs that provide network access across the battlefield.

Mobile networks must be understood at a series of levels, from radio to packet network to communication infrastructure and its resulting impact on the warfighter. To make this difficult problem tractable, the MNMI has developed a four-prong approach. The first is Mobile ad hoc Network (MANET) simulation, where large-scale HPC assets can be used to test and optimize large radio deployments. Second is MANET emulation, where researchers can investigate performance of proposed radios high in the network stack (application layer) all the way to the lower physical layers. Third is MANET experimentation, where live and constructive exercises can capture real radio performance and test interoperability with real and virtual assets. This also results in data sets that can be later mined to fill data gaps and verify models. Fourth, a system that ties together all of these aspects and brings in support for visualization and data analysis. The MNMI is addressing all of these topics, and a DHPI (Dedicated HPC Project Investment) for the MNMI will greatly enhance all four of these endeavors with a special ability to greatly augment efforts in MANET emulation. The DHPI will be discussed in a following section. Throughout this MANET emulation effort the EMANE (Extendable Mobile Ad-hoc Network Emulator) from DRS (formerly Cengen Labs) is used. This software is open source and sponsored by NRL (U.S. Naval Research Laboratory) and ARL (U.S. Army Research Laboratory). EMANE can be obtained from DRS. [1]

U.S. Government work not protected by U.S. copyright

II. Large Scale MANET Emulation

One aspect of the MNMI is focused on developing a framework for large scale MANET emulations, up to 5000 emulated devices. This will allow for the research, development, and evaluation of network algorithms, applications and devices in a controlled environment. In addition to the use of unmodified software applications, with this environment it will be possible to integrate virtual devices into live experiments in order to augment the testing parameters and the perceived traffic by physical network devices in the field. This augmentation of live experiments enhances the experience of testers and increases the degrees of freedom that can be evaluated in an experiment. The achievement of these goals for large scale MANET emulations and live experiment integration requires the bridging of a number of technical gaps. One of the bridging technologies presented here is the development of a real-time RF propagation computation and the hi-fidelity software emulation of U.S. Army waveforms. Currently the MANET emulation framework used here includes radio models for 802.11 a/b/g, the soldier radio waveform (SRW) developed by Cengen, and a configurable model called rf pipe. On top of all of these we currently use the OLSRd ad-hoc wireless mesh routing daemon. [2]

The emulation of a large number of MANET devices has required the integration of a number of computational and cluster management technologies. The first technology used is virtualization in order to split a physical hardware node into multiple nodes, each supporting the emulation of a single network device. With virtualization it is possible to host a heterogeneous collection of guest operating systems on a single host including multiple flavors of Linux and Microsoft Windows. Unmodified applications are then able to run on these guest virtual machines (VM) with minimal overhead. Commodity servers from commercial vendors can currently have up to 64 physical CPU cores, each hosting a VM and sharing a high speed network interface. We have investigated multiple virtualization technologies including full virtualization with KVM, XEN, and VMWare and containers with OpenVZ. Each has their own advantages and disadvantages. For instance overhead and load balancing is better with containers such as OpenVZ because each execution thread is sceduled by the host scheduler. A drawback of OpenVZ is that it does not allow for the hosting of guest operating systems that are different from that the host. Full virtualization, on the other hand, is much more flexible in the operating systems supported through various host/guest combinations. This flexibility comes at the cost of some additional overhead and possibly less load balancing. The management of this large number of VMs, up to 5000, requires HPC techniques such as the netbooting of VM hosts and guests with common images and cluster management with Perceus. [3] Another technology used is the extensible discrete-event mobile networking data model and format (NetDMF) [4]. This data model and framework provides for the storage of MANET data from simulation, emulation and experimentation (SEE cycle). Having a common data format has resulted in the development of common tools for visualization, data mining, and analysis of modeling results and the sharing of scenarios across the SEE cycle.

III. Real-Time RF Path Loss Computations

RF wave propagation models play an essential role in the planning, analysis and optimization of radio networks [5]. For instance, coverage and interference estimates of network configurations are based on field strength predictions. Approaches for field strength prediction can be divided into (semi-) empirical and ray-optical models. For example, the semi-empirical COST-Walfisch-Ikegami model [6] estimates the received power predominantly on the basis of frequency and distance to the transmitter. Ray-optical [7] approaches identify ray paths through the scene, based on wave guiding effects like reflection and diffraction. Semi-empirical algorithms usually offer fast computation times but suffer from inherent low prediction quality. Ray-optical algorithms feature a higher prediction quality at the cost of higher computation times. For MANET emulation integration these algorithms must be computed in real-time for each of the propagation paths possible, with an $O(n^2)$ complexity for a flat network hierachy.

GPUs have been identified as a solution to provide the raw floating-point performance required to compute the RF propagation path loss in real-time. The GPU architecture is also ideally suited for accelerating ray-tracing algorithms that are found in the ray-optical approaches for RF wave propagation modeling. In addition, these RF propagation computations must be tightly coupled with the MANET emulation environment in order to provide real-time, less than 0.5 second, response to computation requests. In addition to potentially real time propagation modeling, a GPU-accelerated algorithm is the basis for adding higher fidelity modeling capabilities such as foliage [8] effects.

A. Real-Time Propagation Path Loss Estimation Using GPGPU Co-Processors

In order to meet the real-time requirements for RF propagation path loss estimation suitable for network emulation, an implementation of the Longley-Rice [9] Irregular Terrain Model (ITM) was developed for GPU co-processors. The implementation is based on the open-source C implementation available from the U.S. Department of Commerce [10]. The code development required significant re-factoring to employ algorithms suitable for the fine-grained parallelism of modern GPU architectures. A digital terrain extraction algorithm was also ported to a GPU to allow the entire propagation path loss calculation to be performed on a GPU with the GPS coordinates of each radio as an input, and the path loss matrix as the output.

Several challenges were encountered in the development of the GPU implementation. First, the C code upon which the implementation is based closely reflected the original FORTRAN implementation and contained many basic constructs ill-suited to modern processors, requiring substantial

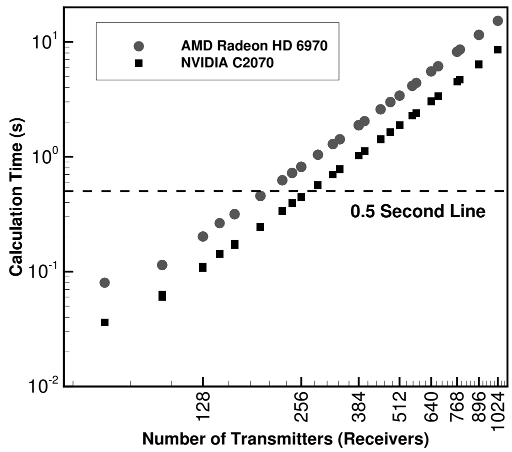

Fig. 1. Plot of total ITM (Longley-Rice) calculation time versus number of transmitters/receivers. The 0.5 second line represents the maximum time allowed for real-time computations.

reformulation. One of the challenges encountered in obtaining an efficient GPU implementation was the replacement of nested conditional control flow with predicated execution inside inner loops. An issue encountered during development of the GPU implementation was the use of single-precision that required the reformulation of several algorithms to address the numerical stability of transcendental functions at small angles. The port to GPU co-processors was enabled by the use of OpenCL™, which is an industry standard programming API for parallel programming of heterogeneous computing platforms. [11] The use of OpenCL ensures portability across modern multi-core and many-core processors and specifically supports the use of GPUs from AMD and NVIDIA. The ITM implementation here utilized the STDCL [12] interface, which greatly simplifies the use of OpenCL for complex HPC applications. The digital terrain extraction and ITM calculation were implemented using 10 OpenCL kernels executed in succession using a digital terrain map pre-loaded on the GPU co-processors. The computation is performed for each array of radio positions and results in the return of a path loss matrix for all transmitter/receiver pairs.

Through multiple interations the current ITM implementation supports multiple GPUs, has been tested using AMD and NVIDIA GPUs and is capable of providing path loss estimation results for 256 radios within the 1 second limit. The 1 second limit represents 0.5 seconds for GPGPU on device computations as defined by the real-time requirement for network emulation. For 256 radios the calculation can be performed in 0.44 sec using an NVIDIA C2070 [13] as compared with 0.82 sec using an AMD Radeon HD 6970. [14] Complete performance results are plotted in Fig. 1.

IV. Application: VoIP

We wish to investigate the performance of various VoIP configurations within the emulated MANET environment. This will allow us to explore the feasibility of VoIP within MANET as well as consider the impact of network configuration, routing as well as physical layer radio models and node movement.

A. VoIP Basics

VoIP enables participants to send and receive telephone calls using a computer network employing the Internet Protocol (IP) instead of traditional telephone lines and systems. VoIP is one of the fastest growing services on the Internet. The prevalence of broadband networks today as well as advances in hardware and typical computational power has enabled many to forsake their plain old telephone system (POTS) in favor of internet telephony.

A VoIP system utilizes several enabling technologies and standards. To make VoIP calls an network connection is necessary along with a computer running VoIP software, or a specialized phone or adapter. VoIP calls can also be made using cellular telephone thanks to an increase in bandwidth and network performance of 3G and 4G cellular wireless standards as well as increase in computational power of handheld devices.

Several types of signaling can be used to initiate and terminate VoIP calls including Signal Initiation Protocol (SIP), Extensible Messaging and Presence Protocol (XMPP), and H.323. There are several popular propriety VoIP systems such as Skype and the Google Voice web-based application but there are a multitude of open source VoIP softphones as well. One can think of signaling protocols as an analog to the traditional POTS switching or switchboard operators of the past that link calling parties prior to actual voice transmission.

Once the call has been established via the signaling protocol the audio information is transmitted. This may be performed by the same protocol or a separate transmission protocol. For instance XMPP and H.323 do both call signaling and control and media (audio) transport. However many systems, open and proprietary, employ the Real-time Transport Protocol (RTP) for the transmission of the actual voice data.

For our purposes we are using the SIP signaling protocol for call initiation and RTP as a means to transmit the actual audio packets. SIP is an application layer protocol described in RFC 3261 [15]. TCP or UDP can be used as the transport protocol. UDP provides low overhead and has no guarantee delivery of packets while TCP is used for reliable delivery with higher overhead. RTP is also an application layer protocol that uses the UDP transport protocol to provide packet delivery. RTP is defined in RFC 3550 [16] and is used for time sensitive streaming media where a best effort is made to transmit the packet. It is designed for content that is delay sensitive, meaning that if the arrival time of a data packet exceeds some threshold it is no longer usable to the recipient and will be discarded.

To transmit an audio signal representing voice over the Internet it must be converted into a digital signal which compressed and segmented into the RTP packets for transport. When these packets are received at the destination, the digital signal is decompressed and the digital representation is converted to

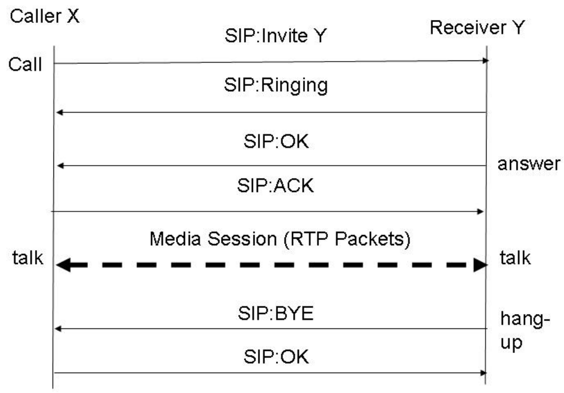

Fig. 2. VoIP Transaction

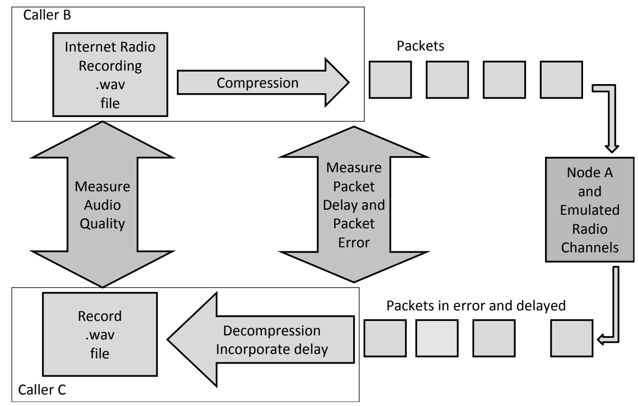

Fig. 3. Experiment End node Calculations

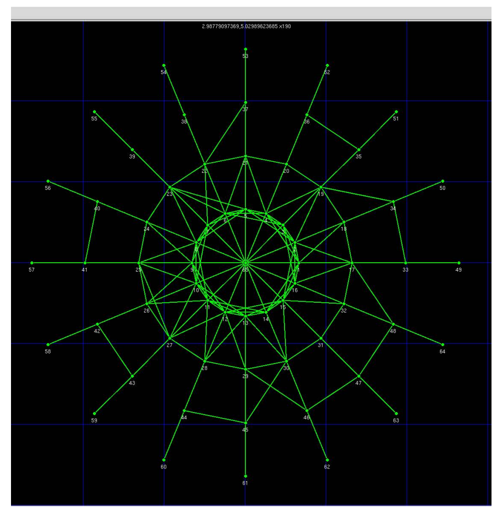

Fig. 4. 16 Outer Nodes Communicate through a Central Node

analog and played over a handset or computer speakers. There exists a plethora of voice compression standards and coders for audio that can be used in VoIP applications such as G.711 [17], G.722, G.723, G.728, and G.729 as well as iLBC (Internet Low Bitrate Vocoder) as defined in RFC 3951 [18].

B. VoIP on MANET

Our MANET emulation employs the Optimized Link State Routing Protocol (OLSR) as specified in RFC 3626 [19]. This is a proactive routing protocol that is optimized by Multi-Point Relay (MPR) flooding and messaging. It generates a constant overhead of control traffic and incurs no route lookup delay.

In our initial experiments, we assume that calls are made to IP addresses and that the IP address that is advertised in the OLSR tables will provide access to a person using that radio and any identity information can be resolve on each and every node independently using the IP/hostname affiliation. This will eliminate the need for the central server and expedite call initiation process while still keeping within the SIP/RTP standards.

In Fig. 2, the transaction of SIP and RTP packets between a caller during a successful VoIP call is illustrated. A basic process of the call is as follows: The caller places a call to the receiver by sending a SIP invitation packet. The receiver responds with a ringing packet and then sends and OK once the call has been accepted (answered); this is acknowledged by the caller. At this time the audio transmission begins and the RTP is used to send audio packets from caller to receiver and receiver to caller. The coordinated selection of the audio compressor/vocoder and the setting of the sender and receiver RTP communication ports occur during the caller SIP:Invite and the recipient SIP:OK(answer). To terminate a call one party, the receiver in this case, sends a SIP “BYE” message this is followed by a SIP “OK” from the other party.

C. Experimental Setup

We generate VoIP calls by playing recorded samples saved in the .wav format. This allows us keep the voice data constant while varying other parameters to determine the system’s performance. VoIP performance is measured by packet loss, end to end packet delay, and packet delay variation (PDV), also known as jitter, and computed as described in RFC 3550 [16]. In addition, we measure the perceptual call quality using Perceptual evaluation of speech quality (PESQ) Algorithm [20]. The experimental set up and performance measures are shown in Fig. 3.

The software implementation of SIP and RTP on each node is performed by PJSIP and PJMEDIA, respectively [21]. PJSIP.org website provides the Open Source, comprehensive, high performance, small footprint multimedia communication libraries written in C language for building embedded/non-embedded VoIP applications. These features are desirable for the MANET emulation.

D. Experimental Results

We are using this emulation setup to determine experimentally the Maximum Number of Voice Calls (MNVC) [22] supported with various radio models and complex MANET configurations. We design our initial configuration so that all calls will be routed through a common central node along with unique intermediate nodes as depicted in the screenshot from

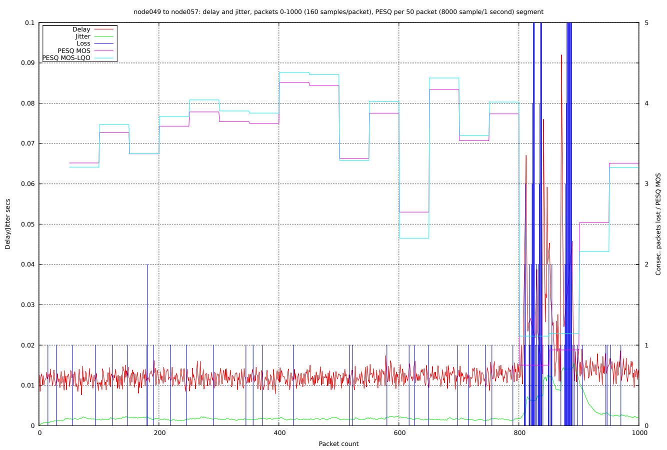

Fig. 5. VoIP Performance, single conversation

the network visualizer Fig. 4. The physical layer model in this case is the RF path loss model described in III. Green lines indicate connectivity between nodes during this snapshot in time.

Preliminary results of VoIP performance are shown in Fig. 5. Two-way VoIP conversations occurred simultaneously between the outer nodes on one side of the central node to the other side. Again all traffic traversed the central node. Although other connection paths existed, they were not the shortest path and therefore were not taken. (Traceroute was used to confirm this behavior.)

These results show packet loss as indicated as (blue crosses), end to end packet delay, and packet delay variation (PDV) (jitter) for a caller in the network. Note that jitter buffering was not utilized in this experiment. In addition, the PESQ metric was calculated to assess the audio quality experience by the caller.

In future experiments, we will vary the network size, physical layer radio model, vocoders, topology, node movement and other VoIP parameters to determine the performance limits of VoIP on MANET.

V. Dedicated HPC Project Investment

Through the HPCMP Office a dedicated HPC project investment was awarded to the the MNMI in order to reach the goals set forward by the MNMI. One of these goals is to emulate 5000 radios in real time. This emulation goal requires not only the computing power to host the virtual machines but the capability to compute the RF path loss in real time so that these emulated radios can be interfaced with live experiments or force modeling simulations. The DHPI system is therefore configured with over 5000 physical CPU cores and approximately 500 GPUs. This configuration gives the DHPI system a peak theoretical throughput approaching 1 PFLOPs. This would place it near the top 10 super computers in the world based on the November 2010 rankings. [23]

VI. Conclusions and Future Work

Through this effort we have developed a framework that leverages many technologies allowing for the emulation of large MANETs. The use of a common data file format such as NetDMF allows for the mining and visualization of data in a consistant manner with MANET experiments and simulations. At run-time, the use of HPC tools such as perceus and slurm for cluster management and job control, respectively, provides efficient use of our hardware resources and scalability to 1000's of virtual machines and hosts. The use of GPGPUs provides the computational horse power required to compute the RF propagation path loss ITM algorithm in real-time. With the installation of the DHPI hardware and this framework we will be able to emulate 5000 MANET nodes in real-time and provide critical data for application analysis and scenario development.

This framework allows the exploration of VoIP over MANET performance by varying numerous design parameters including physical layer radio models, network configuration (number of nodes and placement), topology, node movement, vocoder selection, multicast (conference calls), unicast (point to point calls), etc. We anticipate that this exploration will help lead to the design of an additional voice communication capability for the soldier on the battlefield.

VII. Acknowledgements

The authors would like to acknowledge the support received from the High Performance Computing Modernization Program Office (HPCMPO) under the Mobile Network Modeling Institute (MNMI).

References

[1] DRS-Cengen. EMANE Website. http://labs.cengen.com/emane. [Online; accessed 28-April-2011].
[2] A. Tønnesen. Impementing and extending the Optimized Link State Routing Protocol. Master's thesis, University of Oslo, 2004.
[3] Infiniscale. *Perceus 1.6 User Guide*, June 2010.
[4] Jerry Clarke and Ken Renard. Netdmf: A scalable active data hub for mobile network modeling. *HPCMP Users Group Conference*, 0:321–325, 2009.
[5] T. Rick and R. Mathar. Fast edge-diffraction-based radio wave propagation model for graphics hardware. In *Antennas, 2007. INICA '07. 2nd International ITG Conference on*, pages 15–19, 2007.
[6] E. Damasso, editor. *Digital mobile radio towards future generation systems*. Office for Official Publications of the European Communities, Luxembourg, 1999.
[7] Henry L. Bertoni. *Radio Propagation for Modern Wireless Systems*. Prentice Hall Professional Technical Reference, 1999.
[8] Feinian Wang and K. Sarabandi. A physics-based statistical model for wave propagation through foliage. *Antennas and Propagation, IEEE Transactions on*, 55(3):958 –968, March 2007.
[9] A. G. Longley and P. L. Rice. Prediction of Tropospheric Radio Transmission Loss over Irregular Terrain, A Computer Method-1968. Technical report, Environmental Sciences Services Administration, Boulder, CO, 1998.
[10] U.S. Department of Commerce. Irregular Terrain Model (ITM) (Longley-Rice). http://flattop.its.bldrdoc.gov/itm.html. [Online; accessed 26-April-2011].
[11] U.S. Department of Commerce. OpenCL Overview. http://www.khronos.org/opencl. [Online; accessed 28-April-2011].
[12] Browndeer Technology. COPRTH. http://www.browndeertechnology.com/coprthr_stdcl.htm. [Online; accessed 28-April-2011].

[13] Benchmarks used the OpenCL implementation provided by the NVIDIA CUDA Toolkit v3.2.

[14] Benchmarks used the OpenCL implementation provided by the AMD ATI Stream SDK v2.3.

[15] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler. SIP: Session Initiation Protocol. RFC 3261 (Proposed Standard), June 2002. Updated by RFCs 3265, 3853, 4320, 4916, 5393, 5621, 5626, 5630, 5922, 5954, 6026, 6141.

[16] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A Transport Protocol for Real-Time Applications. RFC 3550 (Standard), July 2003. Updated by RFCs 5506, 5761, 6051, 6222.

[17] ITU-T Recommendation G.711, Pulse code modulation (PCM) of voice frequencies. Recommendation G.711 (11/88), 1988.

[18] S. Andersen, A. Duric, H. Astrom, R. Hagen, W. Kleijn, and J. Linden. Internet Low Bit Rate Codec (iLBC). RFC 3951 (Experimental), December 2004.

[19] T. Clausen and P. Jacquet. Optimized Link State Routing Protocol (OLSR). RFC 3626 (Experimental), October 2003.

[20] Perceptual Evaluation of Speech Quality (PESQ): An Objective Method for End-To-End Speech Quality Assessment of Narrow-Band Telephone Networks and Speech Codecs. Recommendation P.862, 2001.

[21] PJSIP. Open source SIP stack and media stack for presence, im/instant messaging, and multimedia communication, December 2010.

[22] J. Cheah, R. North, W. Wang, and R. Perkins. Thoughts on adjunct cell phone support within jtrs network. In *MILITARY COMMUNICATIONS CONFERENCE, 2010 - MILCOM 2010*, pages 1025 –1031, 31 2010-nov. 3 2010.

[23] TOP500.Org. Top 500 Supercomputer Sites. http://top500.org/. [Online; accessed 26-April-2011].