Simulating the Performance of Building Area Networks as a Communication Bridge to Emergency Responders

Stephen Treado, David Holmberg, Steven Cook

National Institute of Standards and Technology

Gaithersburg, MD 20899

E-mail: stephen.treado@nist.gov

Abstract

This paper describes the development and use of OPNET models to simulate the use of building network infrastructures as a communication bridge between emergency responders inside buildings and emergency personnel outside of buildings. Emergency responders frequently encounter difficulties with radio communication when located within buildings because of the shielding effects of the structure. One possibility for overcoming this problem is to use the existing building network infrastructure, including wireless access points, radio transceiver gateways and the IT network, to provide a communication path from the portable radios inside the building to the public safety network outside the building. Methods for modeling such a system were developed and quality of service issues evaluated.

Introduction

Radio communications are an essential element of emergency responder activities. One of the problems with radio communications, especially those utilizing higher frequencies like police radios (800 MHz), is that signals do not propagate well within buildings, due to shielding by building structural and envelope components (the "no cell phone service" effect). This severely hampers the ability of emergency personnel to perform their duties, and threatens their personal safety, as well as the safety of the general public. In addition, it is envisioned that in the future, radios will function as communication links for voice, video and data, which places even more demands on their capability.

In order to help provide more robust radio communication within buildings, existing building network infrastructures could be used to provide a bridge between emergency responders located inside the building and their counterparts outside of the building. These existing building networks, including IT networks, fire alarm networks and phone networks, if suitably configured and provisioned, could be used to carry emergency responder communications in addition to their normal traffic [1]. Figure 1 illustrates the concept of using a building area network (BAN) as a communications bridge.

Security issues, such as authentication and authorization of users, are not specifically addressed in this analysis, but is the subject of related work and standards activities. One additional desirable characteristic of a radio communication path utilizing building networks would be an open, standard architecture, enabling interoperability among various jurisdictions and manufacturer's equipment.

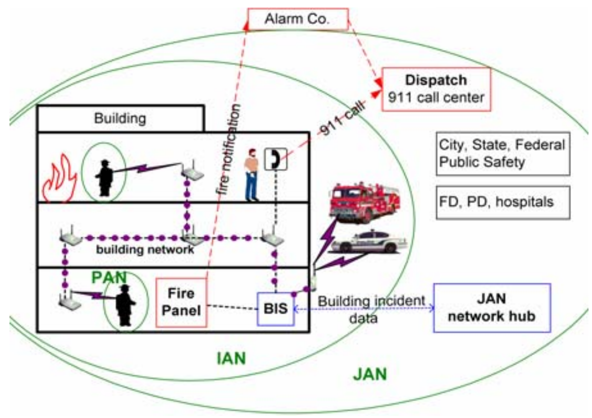

(JAN- jurisdictional area network, IAN- incident area network, BIS- building information server, PAN- personal area network, FD- fire department, PD- police department)

Figure 1. Building area network as a communication bridge

In order to help improve interoperability among public safety communications, communication protocol standards are being developed. The primary activity in this regard is Project 25 (P25) [2]. Radio equipment that is compatible with P25 standards will allow users from different agencies or areas to communicate directly with each other. This will allow agencies on the federal, state/provincial, or local level (or any other agency) to communicate more effectively with each other when required during emergencies, law enforcement actions, etc.

The main goal of this paper is to describe an approach for using OPNET to model communication with emergency responders using building network infrastructures. The communication flow is based on using Project 25 compliant digital radios, RF-IP gateways and building IP networks. While P25 radios have been standardized to some extent, and IP communication protocols are well established, the RF-IP gateways have not been standardized. Figure 2 illustrates the communication flow to be modeled.

Project 25 Overview

Project 25 is a set of standards produced through the joint efforts of the Association of Public Safety Communications Officials International (APCO), the National Association of State Telecommunications Directors (NASTD), selected federal agencies and the National Communications System (NCS), and standardized under the Telecommunications Industry Association (TIA). P25 is an open architecture, user driven suite of system standards that define digital radio communications system architectures capable of serving the needs of public safety and government organizations. The P25 suite of standards involves digital land mobile radio (LMR) services for local, state/provincial, and national (federal) public safety organizations and agencies. P25 open system standards define the interfaces, operation, and capabilities of any P25 compliant radio system [2].

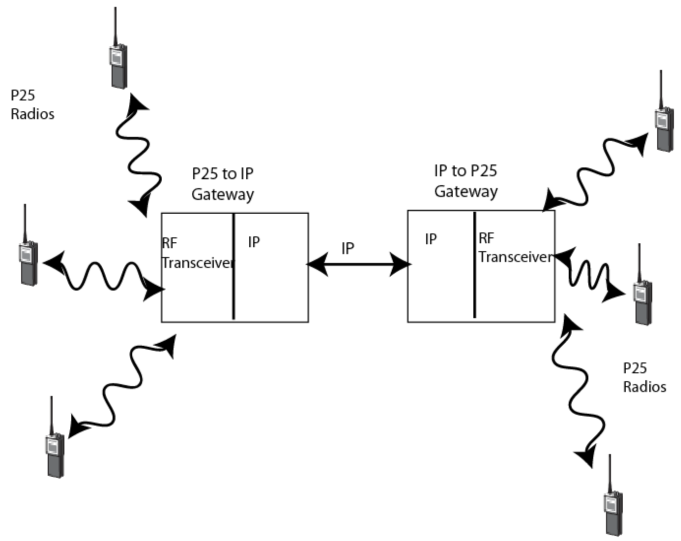

Figure 2. Radio to IP communication flow

Public safety radios are primarily used for voice communications, although data and video applications are also being considered, and in some cases, being utilized. There are several commercial products available that provide voice over IP fuctionality, albeit in a proprietary manner. Video and data transfer system are under development.

Modeling the Communications Flow

OPNET was used to develop a simulation model for the P25 radio to IP communication, including voice, video and data transfer applications. Voice and video are the most demanding applications, video because of the high bandwidth requirements, and voice because of the intolerance for data loss. Figure 3 shows an example of the P25 data structure for digital voice transmission, consisting of header information, including addressing and control information, followed by segments containing the encoded audio stream.

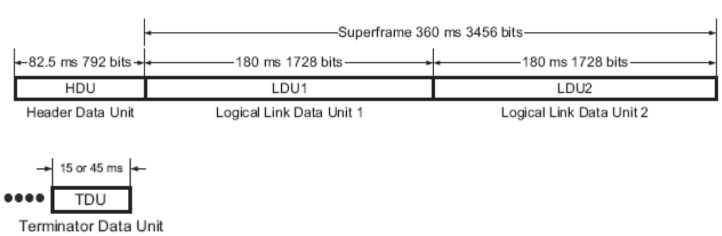

Figure 3. Data structure for P25 voice transmission

Similar data structures are specified for data transfer. The protocol for digital fixed station data transfer is based on Ethernet and IP standards, as shown in Figure 4.

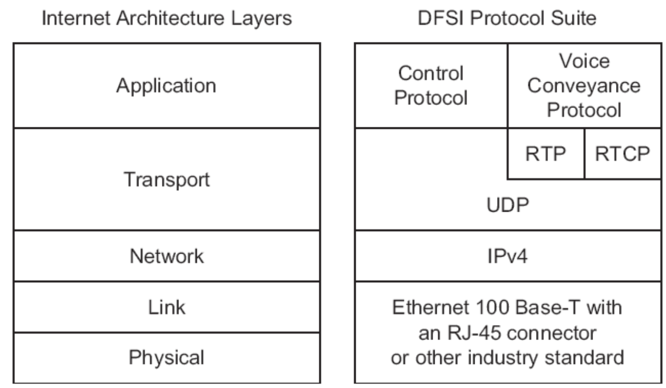

Figure 4. P25 Digital Fixed Station Interface protocol

Messages can be exchanged using either confirmed or unconfirmed operations. The respective data structures are shown in figure 5.

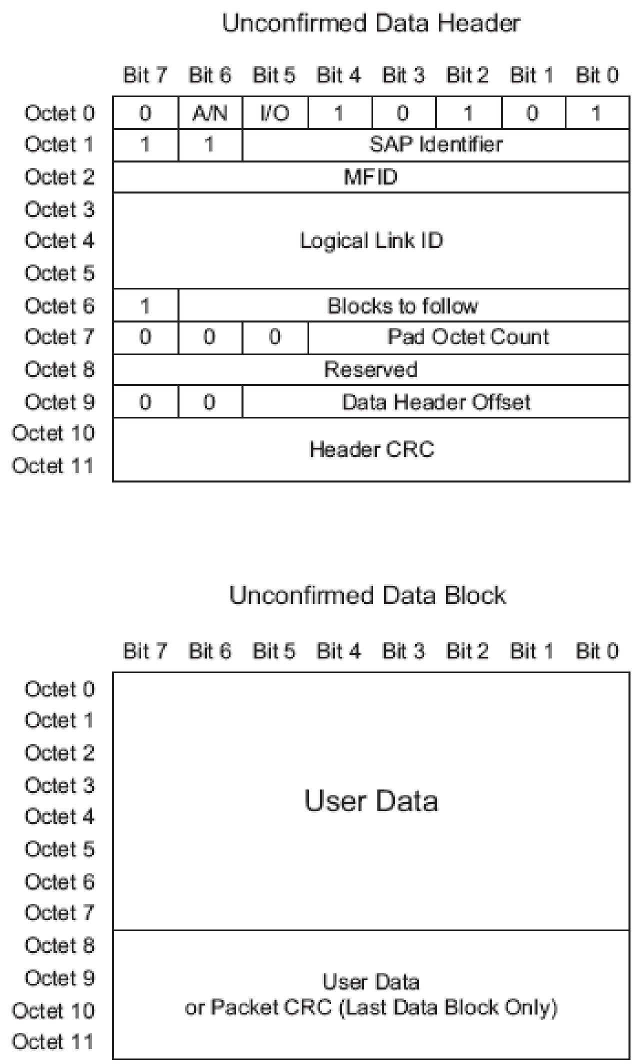

Figure 5. P25 Data block structures

P25 radios were modeled as wireless devices communicating with access points using wifi (802.11g). This approach is based on the assumption that the P25 data structures are similar to those used for wifi, from the point of view of general data handling and manipulation. While there are obvious differences in specific formats between the two protocols, the effect on quality of service and performance should be similar. Efforts are underway to develop specific OPNET models of P25 radios. P25 has not produced a standard for video communication, so a generic model was used as described below. Typical emergency response scenarios were developed, and modeled, also as described below.

The simulations consisted of first responders (FRs) that connect via 802.11g wifi to wireless access points (WAPs). These access points in turn connect through two building network switches to the Internet. Voice, video, and data traffic related to the accident and the first responders' health and status is then sent over this network into the internet, as per the diagram in Figure 6:

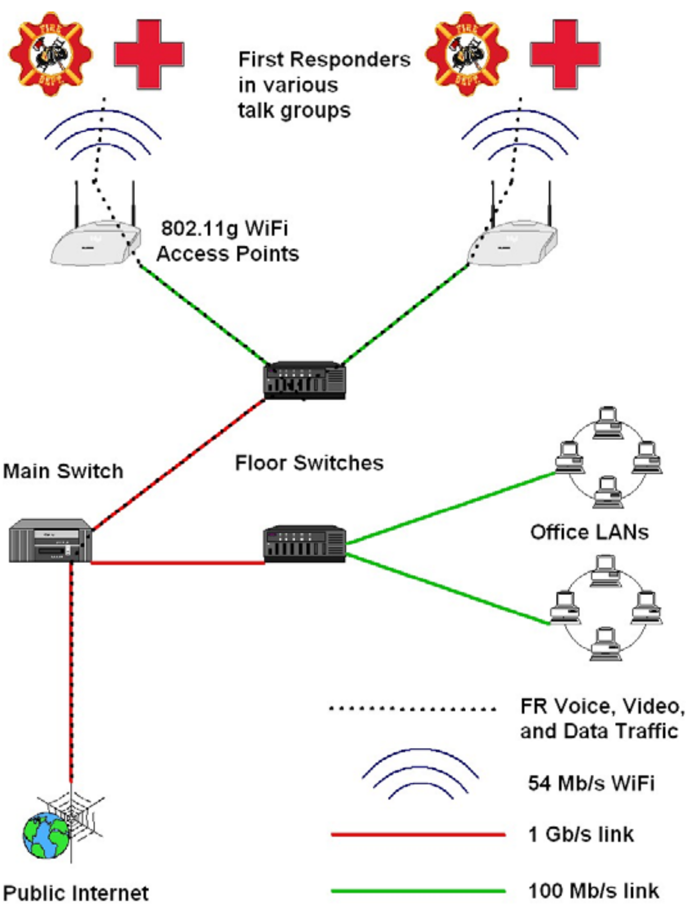

Figure 6. Simulation configuration for emergency response scenarios

Simulation Data Used

Speech traffic was defined based on research listening to actual voice data from dispatch channels. OPNET's default Voice application was used as the base, with the following changes made:

- incoming/outgoing silence: exponential (2 sec)
- incoming/outgoing voice: normal (5.8 sec, 23.5 sec)
- best effort

Video traffic was defined as a custom application that would initialize and then send a certain number of packets of a certain size every second. This simulates sending data at a near-constant defined speed. The information used was:

- MPEG-2 encoding: 1358 byte packet, UDP, 1.52 Mbps (140 packets/sec), best effort
- MPEG-4 encoding: 600 byte packet, UDP, 768 Kbps (160 packets/sec), best effort

Health / status traffic was likewise custom-made:

- 20 Kbps constant traffic (1@2500 byte packet / sec), UDP, best effort

Network:

- 1024000-bit buffers are used for both access points and first responders.

The first responders connect via an 802.11g, 54 Mbit/sec connection to the wireless access points, which are connected by 100 Mbps lines to a "small" switch, then by 1000 Mbps to a "large" switch, and finally to the public Internet.

Simulations

Deployment Scenario Simulation:

In this simulation, a total of 43 FR's are deployed in two phases in response to a scenario of an industrial explosion in the lowest level of a large laboratory building, with multiple fire and medic units arriving on the scene. The model's two phases involve four locations: Phase 1 is the original arrival team and Phase 2 introduces engine crews six and seven, along with the medics who set up the triage area on the first floor. The first responders deploy as follows:

Phase 1

Basement Area 1 – 12 FR's, video on
Floor Above, Rear – 12 FR's, video on, 4 rescue squad, video on
Lobby – 1 Battalion Chief, video off

Phase 2 (additions)

Basement Area 2 – 8 FR's, video on
Lobby – 6 Medics, video on

The basement area is divided into two parts; part one is where the original first responder team recovers the unconscious occupants; part two is where the team that arrives in Phase 2 sets up to fight the fire. All personnel transmit voice and biometric data; only those listed above additionally transmit video data.

Stress Test Simulation:

This simulation is designed to test the maximum load a single access point can handle before data loss and delay grow to unacceptable levels. Only one access point is left enabled for this simulation, which connects to the same building network as did the multiple access points in the deployment scenario. In this simulation, more and more first responders converge on a single access point; the effects of four, eight, twelve, and finally sixteen first responders broadcasting video, voice, and health status are measured.

Measurements Taken

As the goal of these simulations is to measure the effectiveness of the network under emergency first responder loads, the relevant metrics are traffic throughput, data loss, and network delay. The "global" wireless local area network (LAN) statistics in the OPNET model are the measurements used herein. Network delay across the wired portion of the network was found to be miniscule (on the order of microseconds during the highest load), and was thus ignored. On a lower-speed network with much more background traffic, this could grow to be of significant size, but that is beyond the scope of this research.

The simulations all start with an "off" period of nearly two minutes. During this time there is not yet any load on the network. The simulations then proceed to add network traffic as outlined in the two scenario descriptions given above. Simulations were conducted with both MPEG-2 and MPEG-4 video compression, with identical loading parameters.

Results

Deployment Scenario Simulation (MPEG-2) (Figure 7):

In this simulation, Phase 1 begins at about the 1.75 minute mark, and Phase 2 follows at the five minute mark. A throughput bottleneck occurs at about 24 Mbps (see inset box) during Phase 1, and little additional throughput is obtained during Phase2. Should this example prove accurate, a large response to a fire scenario would necessitate a more efficient video encoding mechanism:

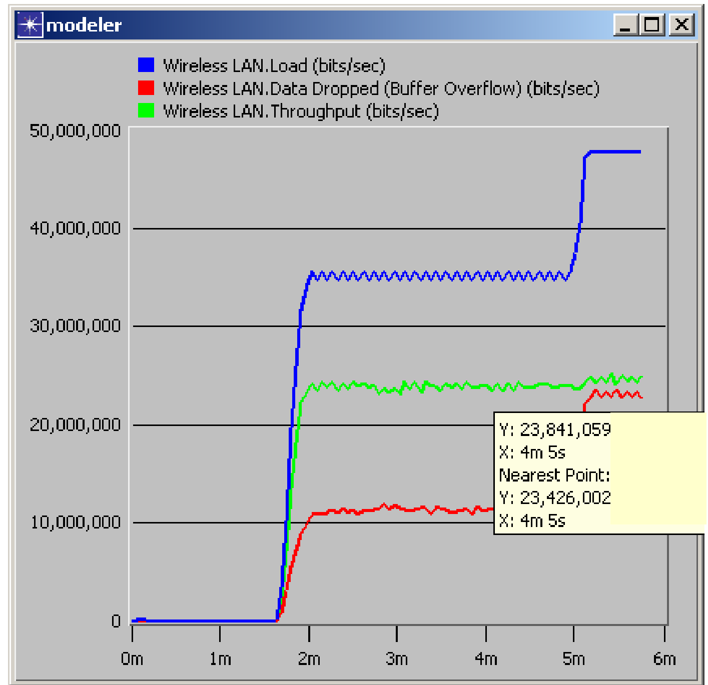

Figure 7: Deployment scenario using MPEG-2 video compression

Also of interest is the network delay resulting from this load. The overall wireless LAN delay (end-to-end delay for a single packet) was found to reach a steady state value of about 1.75 seconds, which is unacceptable for most video applications.

Deployment Scenario Simulation (MPEG-4) (Figure 8):

In this simulation, Phase 1 begins at about the 1.75 minute mark, and Phase 2 follows at the 2.5 minute mark. This test shows that great improvement in throughput is possible by using a more efficient codec (MPEG-4). A far larger percentage of the traffic reaches its final destination (approximately 50% for MPEG-2 compared to 75% for MPEG-4). Also, the jump to Phase 2 is less pronounced:

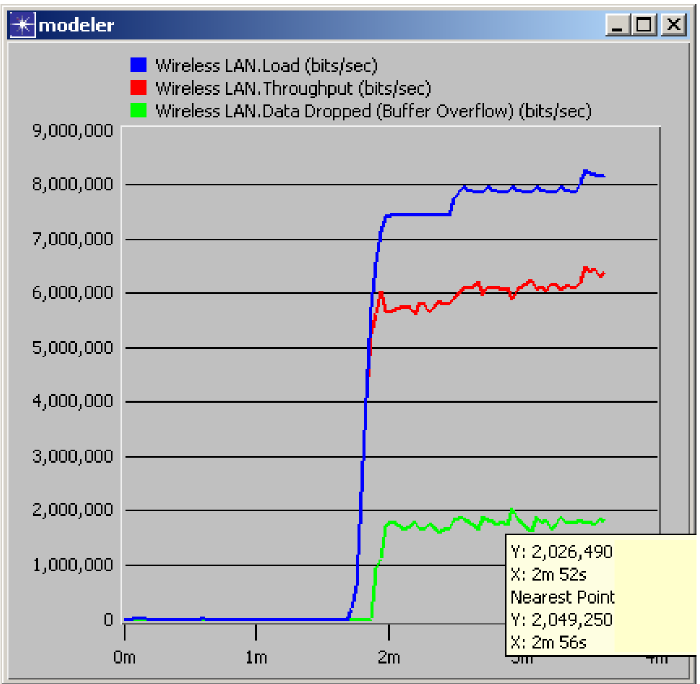

Figure 8: OPNET result for deployment scenario using MPEG-4 compression

The delay was also found to be much more manageable; it oscillates around 0.2 sec, which while not ideal, is much more reasonable for the types of applications that would need to be run on this system. Another interesting point to note is that the load is not as high as with MPEG-2, but data is still getting dropped. This appears to be due to buffer overflow conditions; despite the abnormally large size of the buffers used in this experiment, the data is still coming in faster than it can be processed.

Single Access Point Stress Test (MPEG-2) (Figure 9):

In this test, network traffic is generated in steps of 4, 8, 12 and 16 FR's, creating a stair step pattern in network load. The results of this stress test indicate that with the simulation data being used, a large fraction of data starts to be lost as soon as the first four responders begin transmitting. This loss is due to buffer overflow, a condition where packets cannot be sent as fast as they come in. The saturation point for throughput appears to be 12.5 Mbps, with a delay of at least two seconds at the lowest load levels. This delay constitutes an unacceptable level of performance for voice applications:

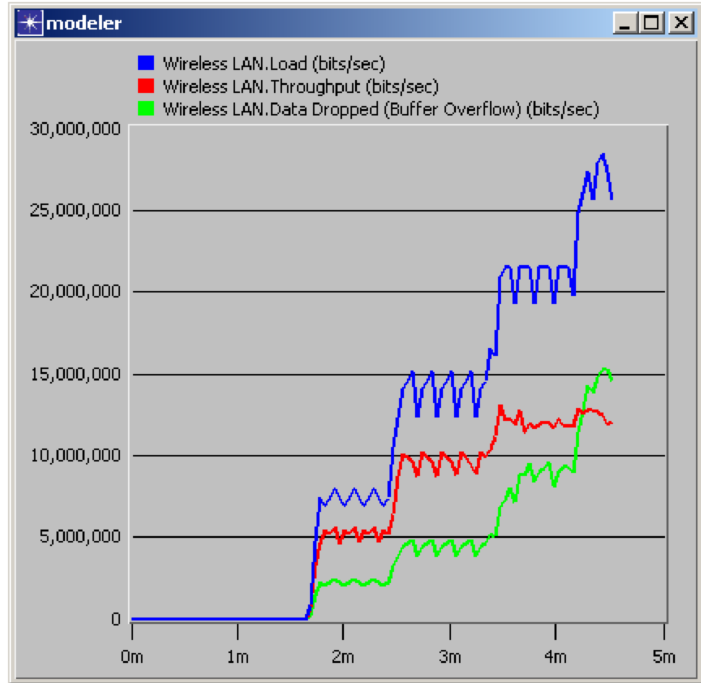

Figure 9: Single access point stress test with MPEG-2 encoding

The rate at which data throughput falls off as a function of load is displayed in Figure 10; the rate of successful transmission levels off somewhat, but gradually falls to about 45%.

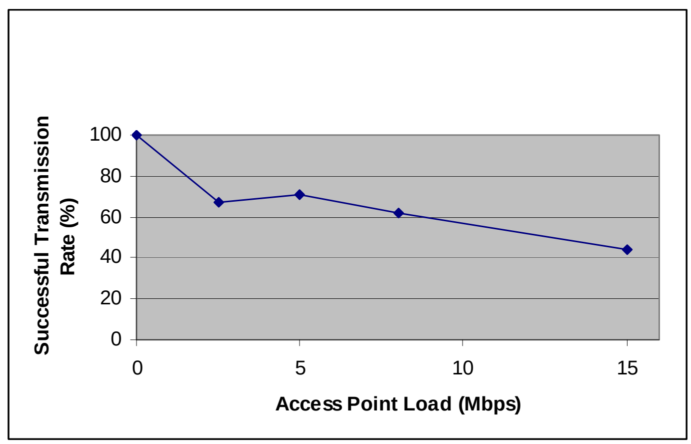

Figure 10 Transmission success rate versus network load for single access point stress test with MPEG-2 encoding

Single Access Point Stress Test (MPEG-4) (Figure 11):

In this test, network traffic is generated in steps of 4, 8, 12 and 16 FR's, creating a stair step pattern in network load.. The improvement demonstrated over the MPEG-2 standard is such that no noticeable data loss occurs until all sixteen first responders are broadcasting. This would set the maximum number of first responders who can run off of a single access point without loss between twelve and sixteen. A graph of transmission percentage vs. load was not included as there is no loss to speak of until the load reaches 4 Mbps:

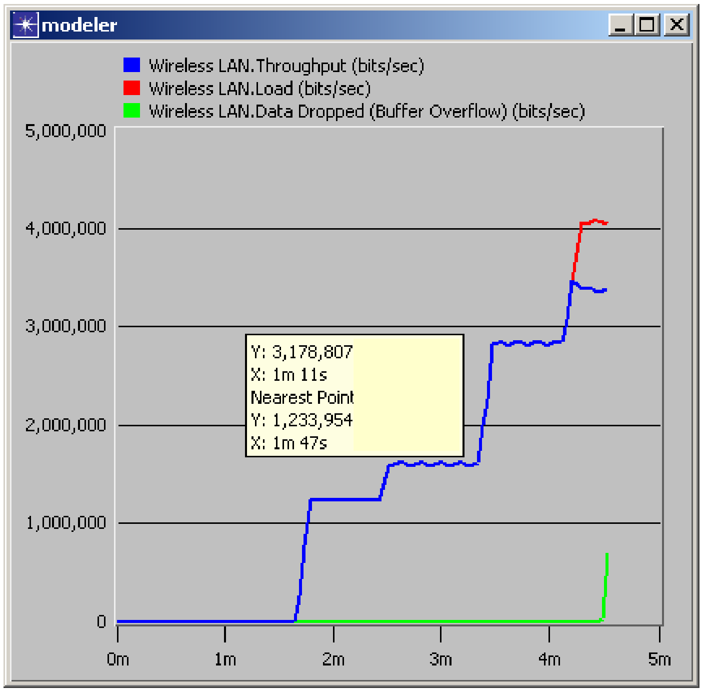

Figure 11: Single Access Point Stress Test with MPEG-4 encoding

The delay is also of particular interest here: instead of growing steadily, it stays put at near zero until the last level (16 FR's). At this point the delay rises dramatically to near 2.25 seconds; this would likely render voice communications useless.

Conclusion

The results of initial simulations demonstrate the feasibility of using building network infrastructures as a communication bridge to emergency responders located in buildings and otherwise shielded from radio communication. However, bandwidth limitations will need to be addressed to enable acceptable performance for large incidents or high amounts of video content.

References

[1] Holmberg, D.G., Davis, W.D., Treado, S. J., Reed, K. A., 2006, "Building Tactical Information System for Public Safety Officials," NISTIR 7314, January, 2006

[2] P25 Radio Systems- Training Guide, Daniels Electronics, Ltd., Jan. 2007