Network Measurement Methods

Homeland Security

DHS-TR-PSC-07-04

Department of Homeland Security

Public Safety Communications
Technical Report

This page intentionally left blank.

Homeland Security

Science and Technology

Command, Control and Interoperability: Communication, Interoperability and Compatibility

Defining the Problem

Emergency responders—police officers, fire personnel, emergency medical services-need to share vital voice and data information across disciplines and jurisdictions to successfully respond to day-to-day incidents and large-scale emergencies. Unfortunately, for decades, inadequate and unreliable communications have compromised their ability to perform mission-critical duties. Responders often have difficulty communicating when adjacent agencies are assigned to different radio bands, use incompatible proprietary systems and infrastructure, and lack adequate standard operating procedures and effective multi-jurisdictional, multi-disciplinary governance structures.

OIC Background

The Department of Homeland Security (DHS) established the Office for Interoperability and Compatibility (OIC) in 2004 to strengthen and integrate interoperability and compatibility efforts to improve local, tribal, state, and Federal emergency response and preparedness. Managed by the Science and Technology Directorate, and housed within the Communication, Interoperability and Compatibility thrust area, OIC helps coordinate interoperability efforts across DHS. OIC programs and initiatives address critical interoperability and compatibility issues. Priority areas include communications, equipment, and training.

OIC Programs

OIC programs, which are the majority of Communication, Interoperability and Compatibility programs, address both voice and data interoperability. OIC is creating the capacity for increased levels of interoperability by developing tools, best practices, technologies, and methodologies that emergency response agencies can immediately put into effect. OIC is also improving incident response and recovery by developing tools, technologies, and messaging standards that help emergency responders manage incidents and exchange information in real time.

Practitioner-Driven Approach

OIC is committed to working in partnership with local, tribal, state, and Federal officials to serve critical emergency response needs. OIC's programs are unique in that they advocate a "bottom-up" approach. OIC's practitioner-driven governance structure gains from the valuable input of the emergency response community and from local, tribal, state, and Federal policy makers and leaders.

Long-Term Goals

- Strengthen and integrate homeland security activities related to research and development, testing and evaluation, standards, technical assistance, training, and grant funding.
- Provide a single resource for information about and assistance with voice and data interoperability and compatibility issues.
- Reduce unnecessary duplication in emergency response programs and unneeded spending on interoperability issues.
- Identify and promote interoperability and compatibility best practices in the emergency response arena.

This page intentionally left blank.

Public Safety Communications Technical Report

Network Measurement Methods

DHS-TR-PSC-07-04
November 2007

Reported for: **The Office for Interoperability and Compatibility**
by NIST/OLES

Homeland Security

This page intentionally left blank.

Publication Notice

Disclaimer

The U.S. Department of Homeland Security's Science and Technology Directorate serves as the primary research and development arm of the Department, using our Nation's scientific and technological resources to provide local, state, and Federal officials with the technology and capabilities to protect the homeland. Managed by the Science and Technology Directorate, the Office for Interoperability and Compatibility (OIC) is assisting in the coordination of interoperability efforts across the Nation.

Certain commercial equipment, materials, and software are sometimes identified to specify technical aspects of the reported procedures and results. In no case does such identification imply recommendations or endorsement by the U.S. Government, its departments, or its agencies; nor does it imply that the equipment, materials, and software identified are the best available for this purpose.

Contact Information

Please send comments or questions to: S&T-C2I@dhs.gov

This page intentionally left blank.

Contents

This page intentionally left blank.

Abstract

This report presents the results obtained using the network model described in Section 5 of the Public Safety Statement of Requirements (PS SoR) Volume II [1], which describes example speech and video applications that could be carried over the network. Graphs in this report plot either of the following in a public safety communications network:

- Packet loss probability versus the number of public safety communications devices (PSCDs) sharing the network
- Expected delay versus the number of PSCDs sharing the access network

The results presented here are the basis for the network performance requirements given in Section 7 of the PS SoR Volume II [1].

Key words: application data rate, asymmetrical network path, channel data rate, network channel data rate, end-to-end network delay, network load, network speech application, network video application, network packet loss, symmetrical network path

1 Introduction

This report includes a series of graphs to illustrate network performance. Each graph assumes a particular value for the network channel data rate and, in the case of the video application, a particular source compression algorithm (ITU-T H.264 or MPEG-2). Loss curves include results for Slotted Aloha but not time division multiple access (TDMA). The TDMA worst-case loss probability is zero if the offered load is less than the channel data rate, and 1 if it is greater; the delay graphs show this cutoff point.

2 Network Paths

Figure 1 illustrates the following symmetrical and asymmetrical network path types in Table 1.

Table 1: Symmetrical and Asymmetrical Network Path Types

Path	Description
A Symmetrical	Public safety communications device (PSCD) via first responder's vehicle (FRV) to PSCD (PAN-IAN-PAN)[a] (involves one IAN; two wireless links)
B Symmetrical	PSCD via jurisdiction communication tower to PSCD (PAN-JAN-PAN) (involves one JAN; two wireless links)
C Symmetrical	PSCD via FRV to jurisdiction communication tower to FRV to PSCD (PAN-IAN-JAN-IAN-PAN) (involves two IANs and one JAN; four wireless links)
D Symmetrical	PSCD via FRV to jurisdiction communication tower to EAN to another jurisdiction communication tower to FRV to PSCD (PAN-IAN-JAN-EAN-JAN-IAN-PAN) (involves two IANs, two JANs, and one EAN; six wireless links)

Table 1: Symmetrical and Asymmetrical Network Path Types

Path	Description
E Symmetrical	PSCD to jurisdiction communication tower to EAN to another jurisdiction communication tower to PSCD (PAN-JAN-EAN-JAN-PAN) (involves two JANs and one EAN; four wireless links)
F Asymmetrical	PSCD via FRV to jurisdiction communication tower to PSCD (PAN-IAN-JAN-PAN) (involves one IAN and one JAN; three wireless links)
G Asymmetrical	PSCD via FRV to jurisdiction communication tower to EAN to another jurisdiction communication tower to PSCD (PAN-IAN-JAN-EAN-JAN-PAN) (involves one IAN, two JANs and one EAN; five wireless links)

a. Public safety network types: extended area network (EAN), incident area network (IAN), jurisdictional area network (JAN), personal area network (PAN).

Figure 1: Hierarchical Reference Paths Based on "Natural Network Hierarchy"

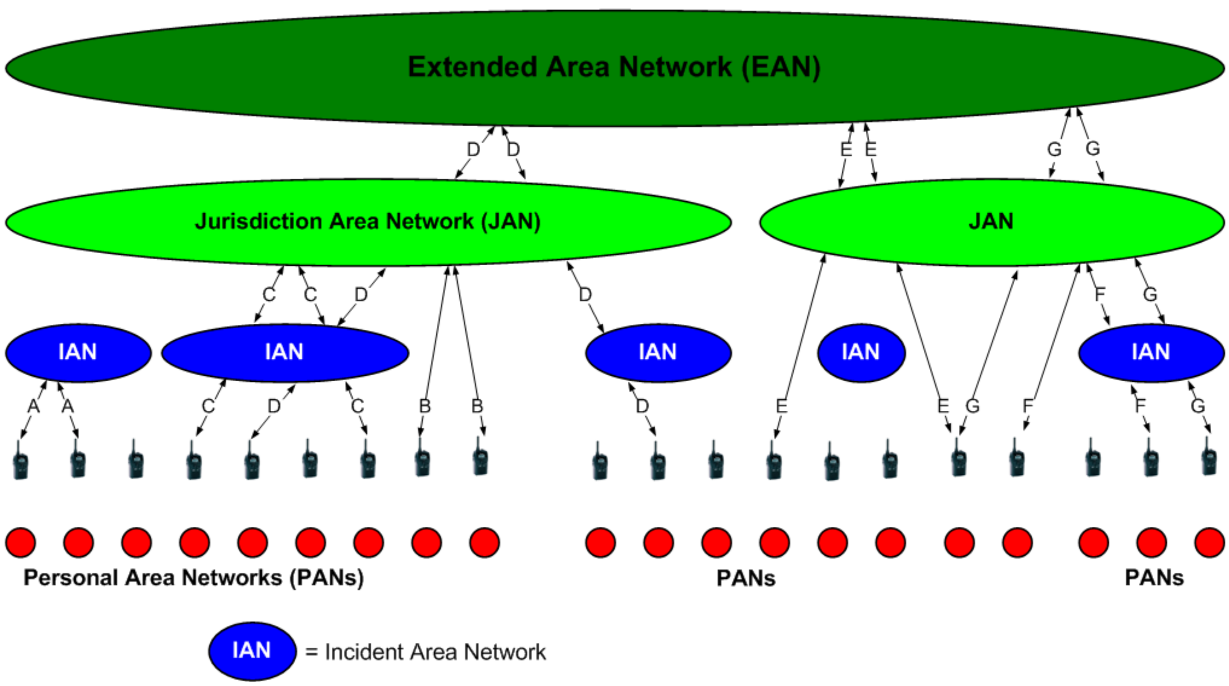

The graphs in sections 4 and 5 plot one of the following in the network:

- Packet loss probability versus the number of PSCDs sharing the network
- Expected delay versus the number of PSCDs sharing the access network

When a graph includes curves with multiple FRVs, such as for paths C (Figures 6 and 7), D (Figure 8 and 9), F (Figures 12 and 13), and G (Figures 14 and 15), the x-axis contains the number of PSCDs in the IAN per FRV. For example, a curve labeled FRV=20 at a point equal to 10 for the number of PSCDs in IAN means that there are 10 PSCDs for each of the 20 FRVs. Thus these 20 FRVs are generating to the JAN (jurisdictional communication tower) an aggregation of 200 PSCDs. For paths that include a JAN, each curve we plot is associated with a different value for the number of FRVs attempting to communicate on the JAN. A legend to the right of each graph identifies the parameters associated with each plot curve.

3 Modeling User Applications

From the view of the network, a user application is modeled as a traffic generator. In other words, at regular time intervals the user application produces a certain-sized packet (in bytes). Using these packet size and generation interval parameters, we can calculate an average application data rate, which gives us an estimate of the network traffic load, or the offered load. Packet size and generation interval parameters are sufficient if the user application is a constant bit rate service. Additional parameters may be necessary to model a user application other than a constant bit rate service.

The next two sections describe tactical speech and video as constant bit rate user applications.

4 Example Speech Application

ITU-T G.711 [2] defines a speech encoding scheme known as PCM (Pulse Code Modulation), which produces an 8-bit sample every 125 microseconds. This results in an application data rate of 64 kbps. For our speech application example, we assume the network is a packet network, and not a circuit-switched phone system. This means a number of samples must be grouped together to form an application packet. The packet is given to the network for delivery. The packet has Real-time Transport Protocol (RTP), User Datagram Protocol (UDP), and Internet Protocol version 6 (IPV6) headers applied. We assume that neither the size of the packet containing the original speech sample(s), nor the number of speech samples in a packet, changes over any link or section in the transmission path. That is, no fragmentation occurs along any link on the path. We use two speech sampling packet sizes: 80 samples per packet and 320 samples per packet.

This section graphs the performance of an example PCM speech application using G.711 encoding. We examine the application's performance as a function of the number of PSCDs and FRVs in the networks, as characterized by the loss and delay metrics described in Section 6 of the Public Safety Statement of Requirements (PS SoR) Volume II [1], for each of the path types listed in Section 1. This section also describes the effect of the channel data rate on the loss and delay performance.

4.1 Path Comparison

This section plots, for path types A through G, packet loss probability and average delay vs. the number of PSCDs, including the source PSCD, that share the source PSCD's area network. When PSCDs, FRVs, and a JAN are on a path, the graphs in Figure 2 through Figure 15 plot performance as a function of the number of FRVs competing for access to the JAN tower.

You can make several observations from the graphs in this section (Figure 2 through Figure 15). First, note that the loss curves in Figure 2, Figure 4, and Figure 10 are identical, indicating that the loss performance for Paths A, B, and E is the same. This follows from the fact that each of these paths features two 11 Mbps links that use either TDMA or Slotted Aloha. The difference between Paths A, B, and E arises from the delay performance, which is plotted in Figure 3, Figure 5, and Figure 11, respectively. The access delay to the three paths are identical, but the link and node delays become progressively larger from Path A with its single FRV to Path E, which incorporates two JAN towers and an EAN. You can expect the towers to be more distant from the PSCDs than an FRV would be, and thus the links to them would have greater propagation delays. The EAN adds processing delays at various routers and also features potentially long transit times between its internal nodes.

You can observe the same kind of relationship between Paths C and D, and between Paths F and G. By looking at Figure 6 and Figure 8 we see that the loss performance does not change if we add an EAN to Path C to get Path D. This is because the EAN uses dedicated resources, which eliminates packet loss due to contention. The addition of the EAN has a significant effect on the application's delay performance, however. Figure 7 and Figure 9 show an increase of about 50 ms in the end-to-end delay. Note that the number of PSCDs beyond which Pr{loss} = 1 and the end-to-end delay is infinite, and a plot of the links using TDMA (labeled as "Dedicated" in the legend to the right of the figure graphs) is the same in Figure 7 and Figure 9. This is because the IAN bandwidth is much less than that of the EAN and it is the limiting factor.

Figure 12 and Figure 14 illustrate the loss performance for the asymmetric paths F and G. Path G, which is essentially Path F plus an EAN, exhibits the same performance as Path F, for the same reason that Path D's loss performance is identical to that of Path C. As was the case with Paths C and D, the delay for Path G is about 50 ms greater than the delay for Path F, for all combinations of number of FRVs, PSCDs, and packet sizes.

Figure 2: Packet Loss for Speech on Path A at 11 Mbps

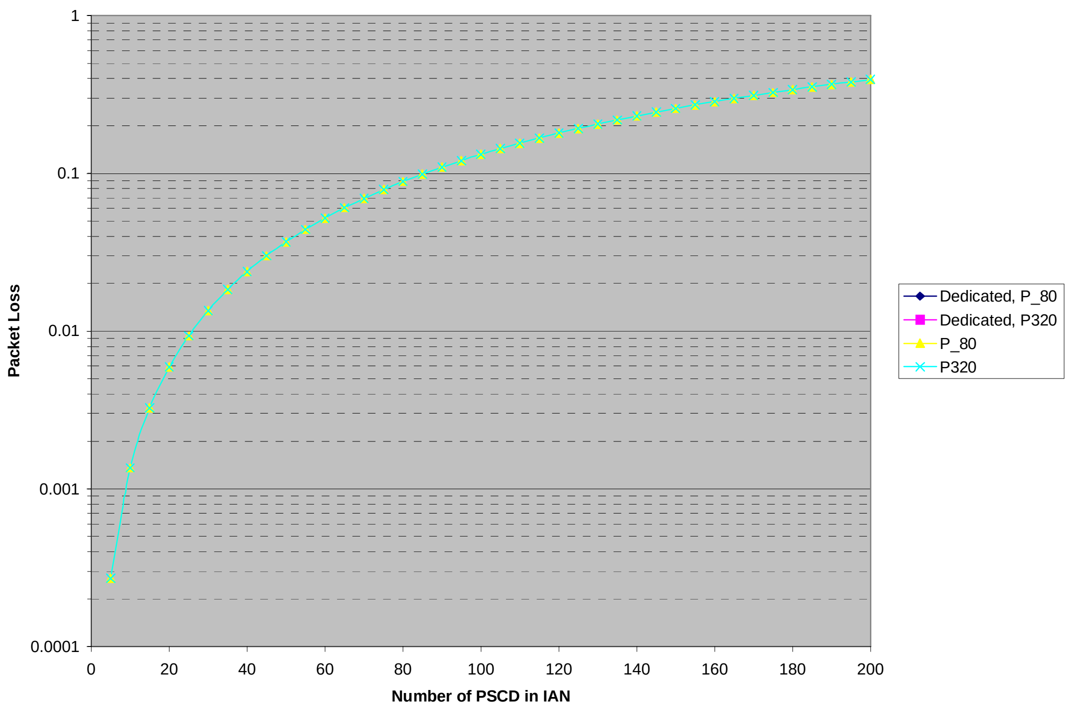

Figure 3: Delay for Speech on Path A at 11 Mbps

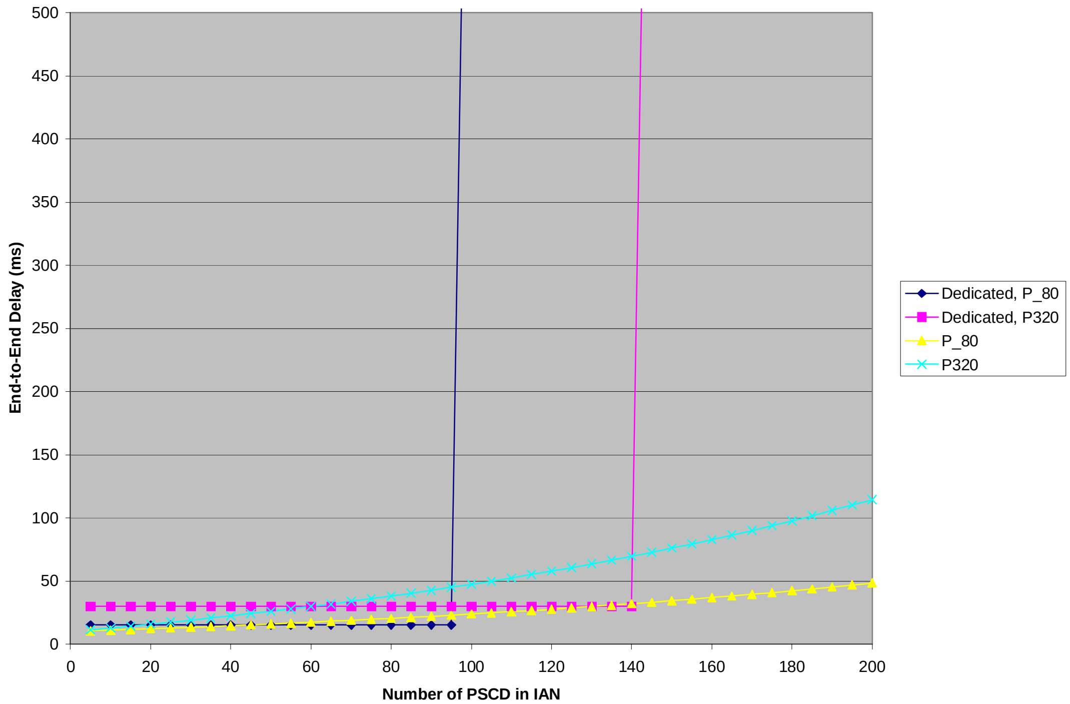

Figure 4: Packet Loss for Speech on Path B at 11 Mbps

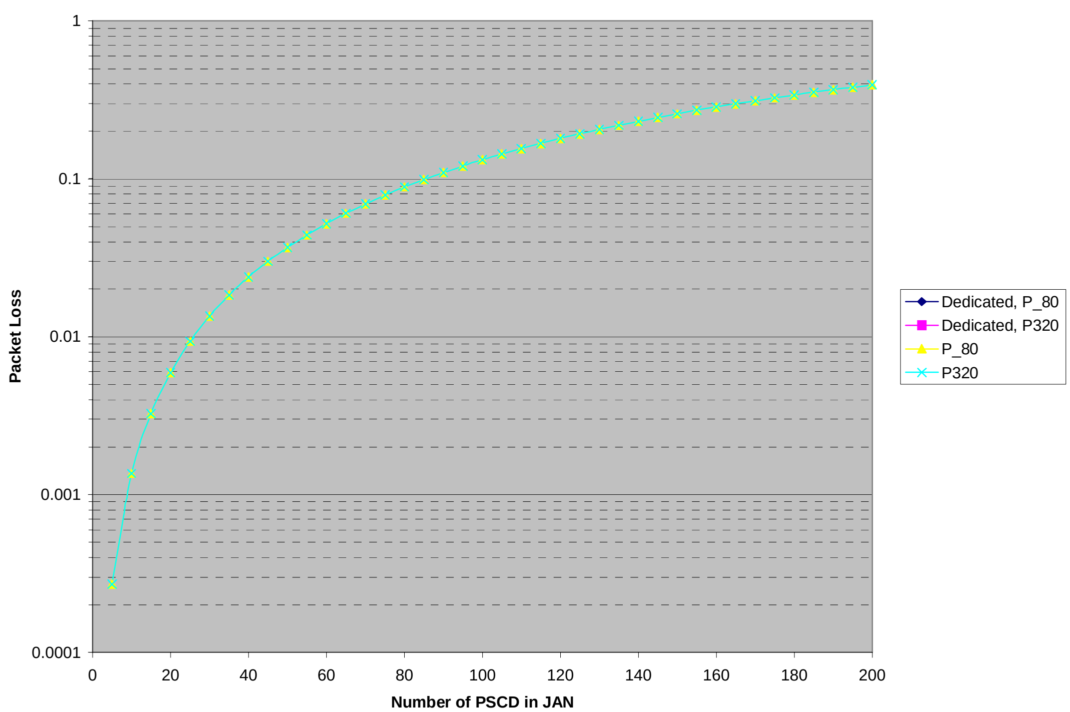

Figure 5: Delay for Speech on Path B at 11 Mbps

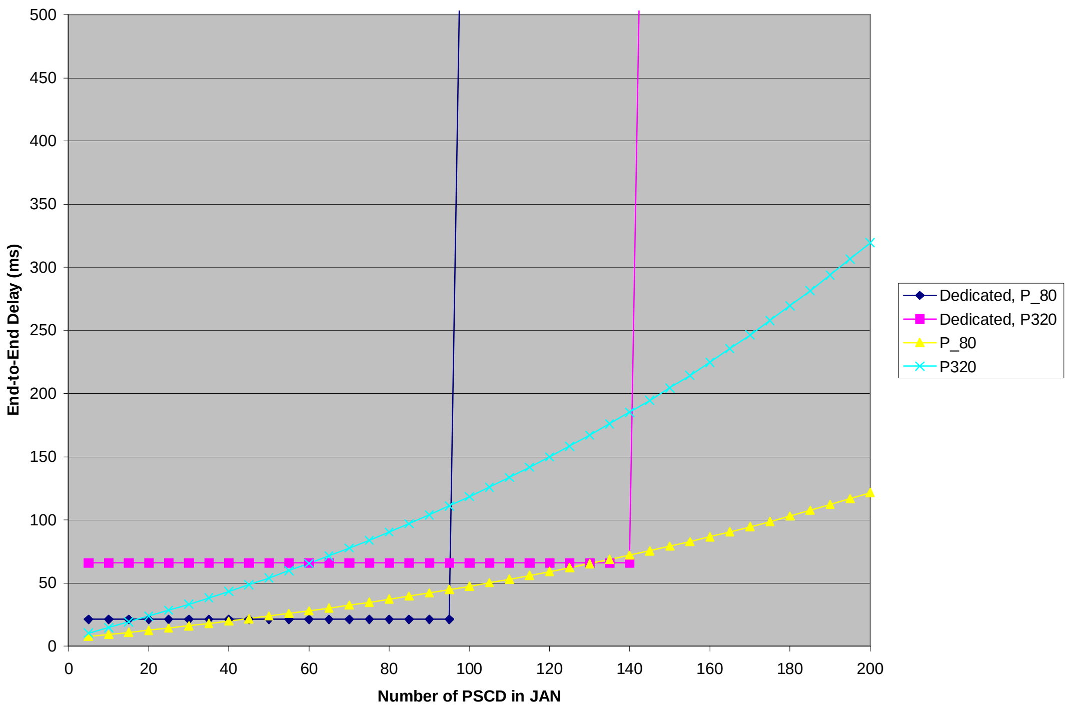

Figure 6: Packet Loss for Speech on Path C at 11 Mbps

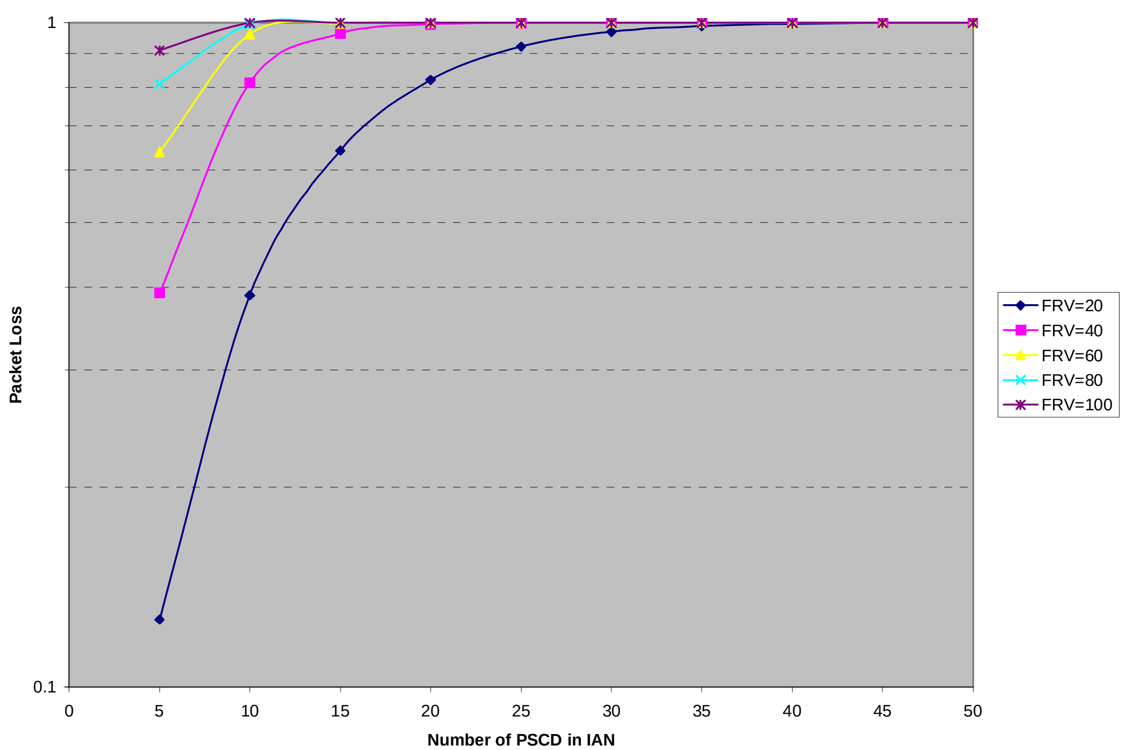

Figure 7: Delay for Speech on Path C at 11 Mbps

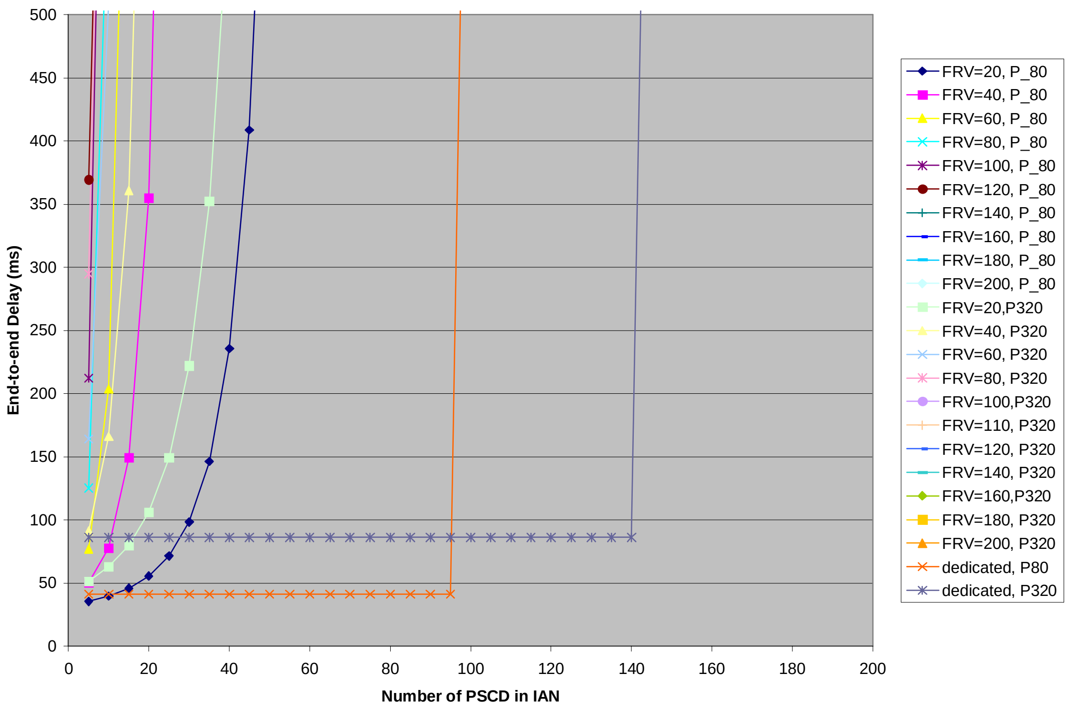

Figure 8: Packet Loss for Speech on Path D at 11 Mbps

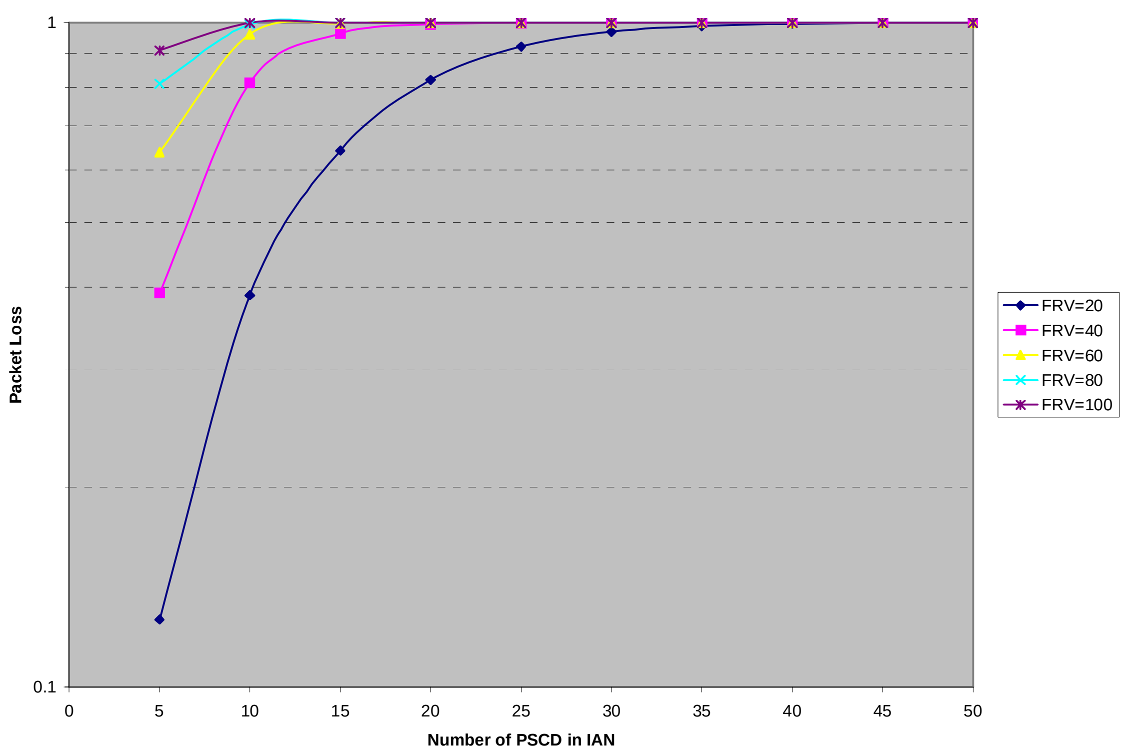

Figure 9: Delay for Speech on Path D at 11 Mbps

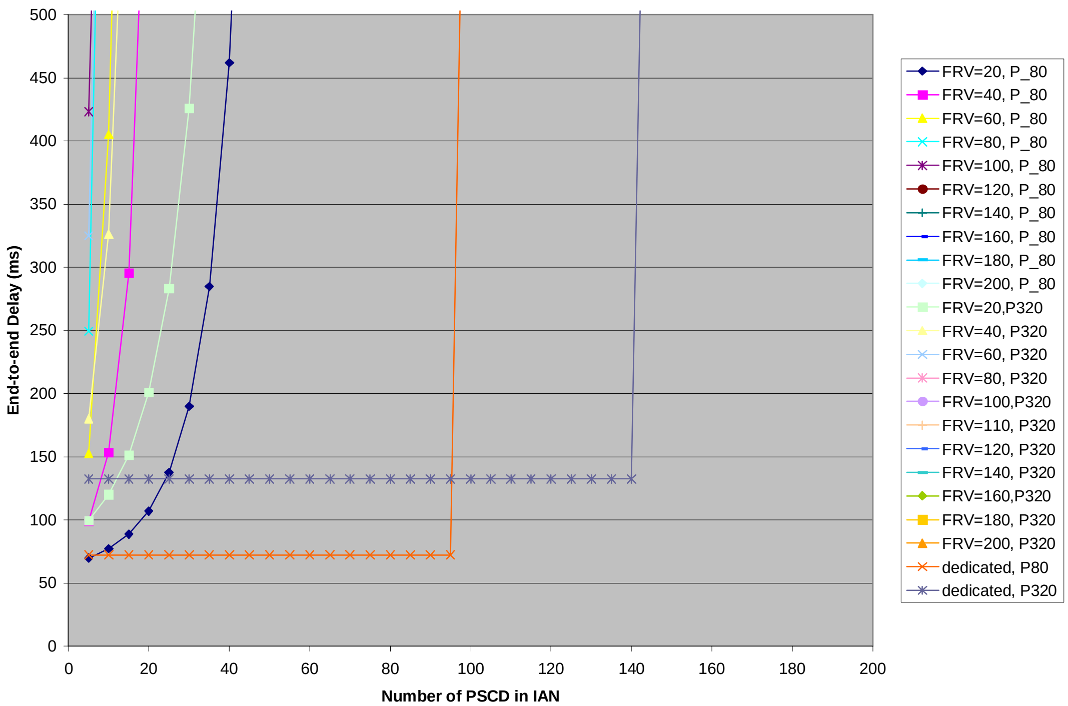

Figure 10: Packet Loss for Speech on Path E at 11 Mbps

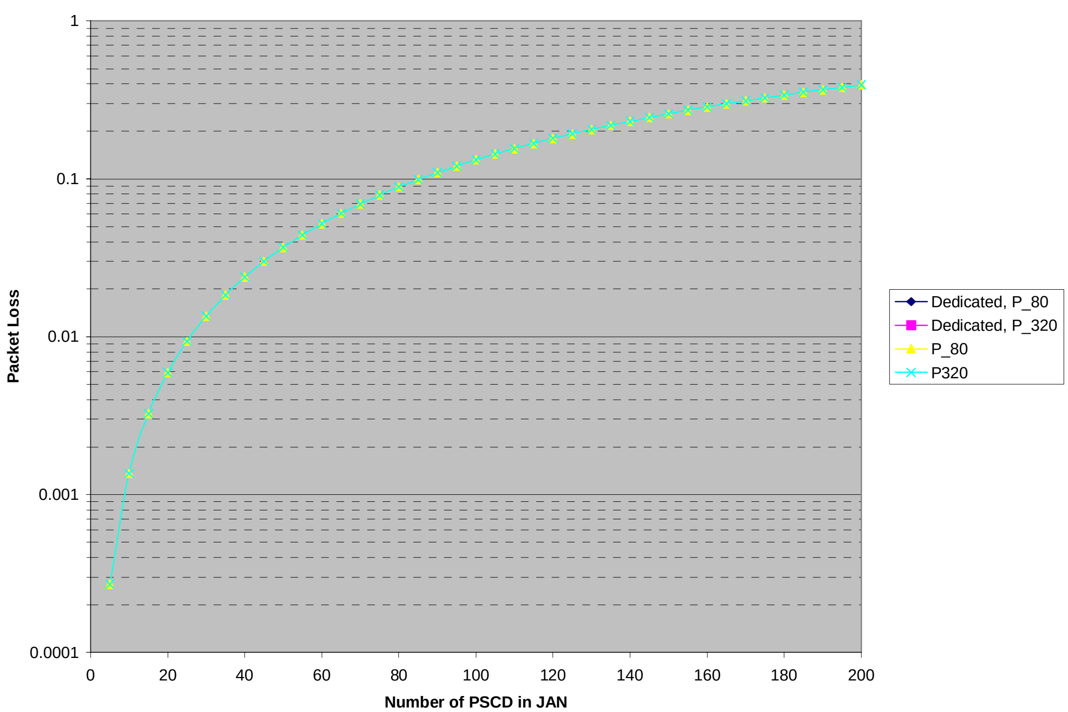

Figure 11: Delay for Speech on Path E at 11 Mbps

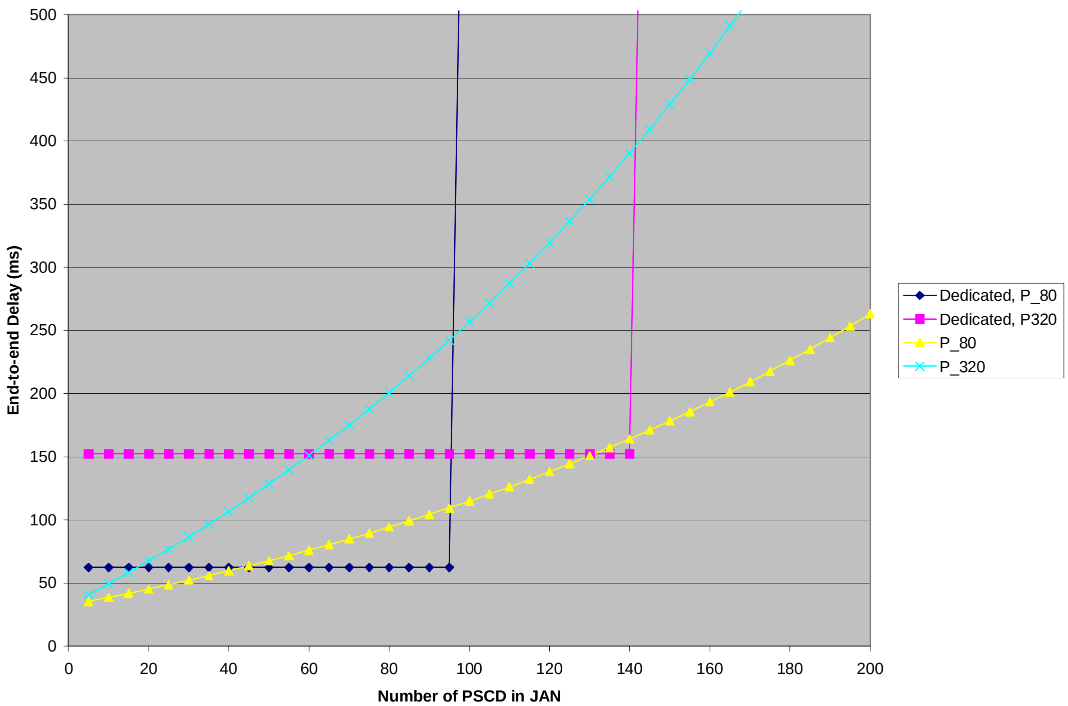

Figure 12: Packet Loss for Speech on Path F at 11 Mbps

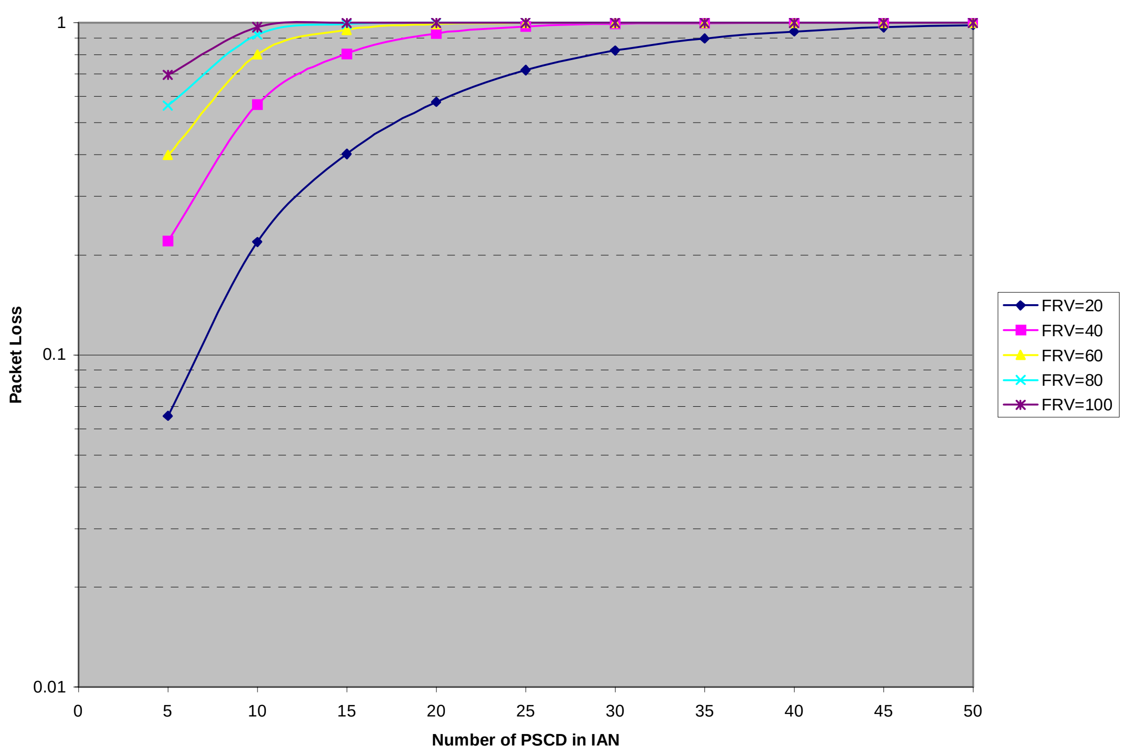

Figure 13: Delay for Speech on Path F at 11 Mbps

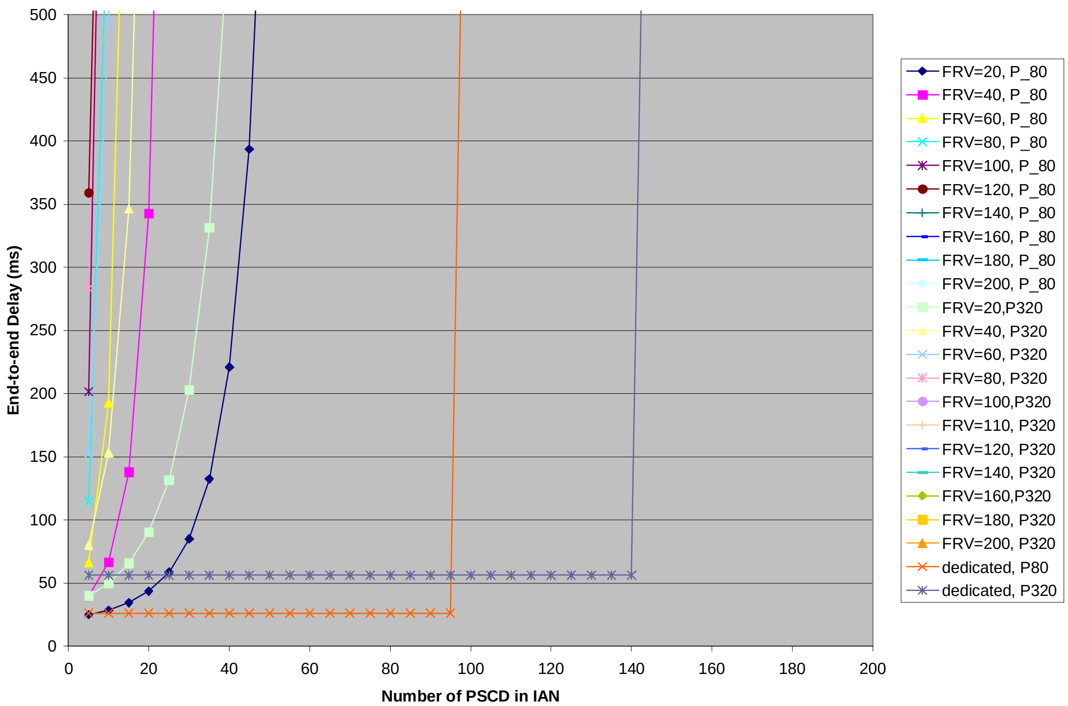

Figure 14: Packet Loss for Speech on Path G at 11 Mbps

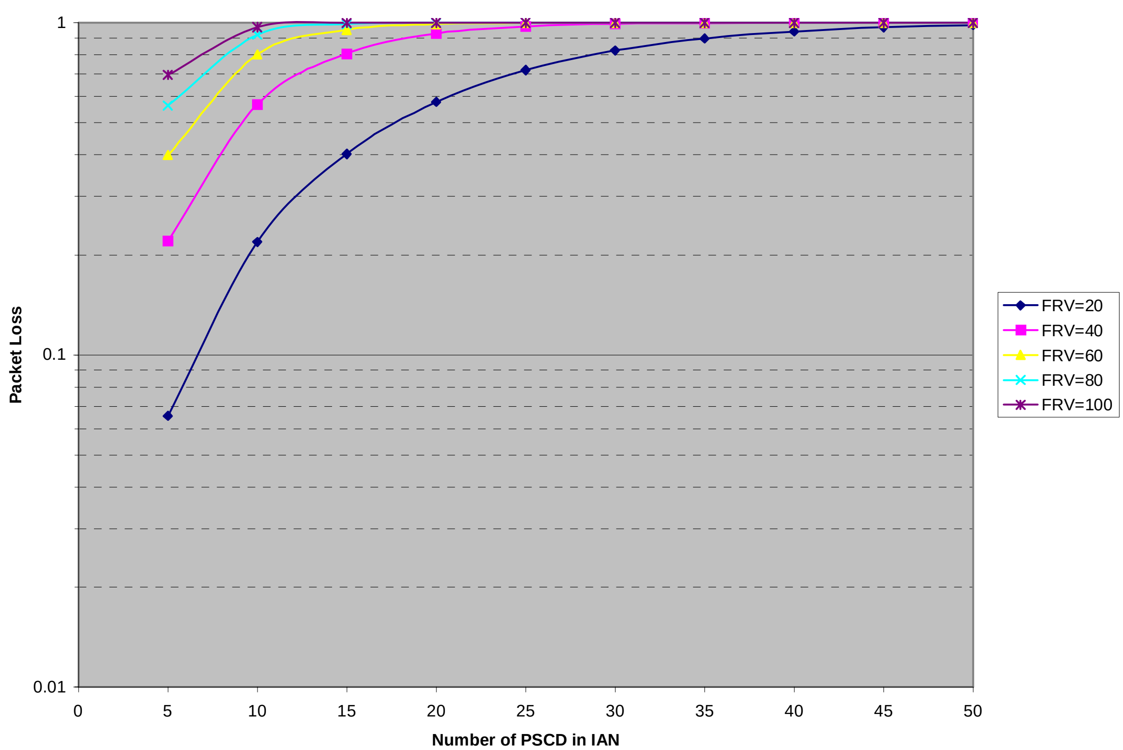

Figure 15: Delay for Speech on Path G at 11 Mbps

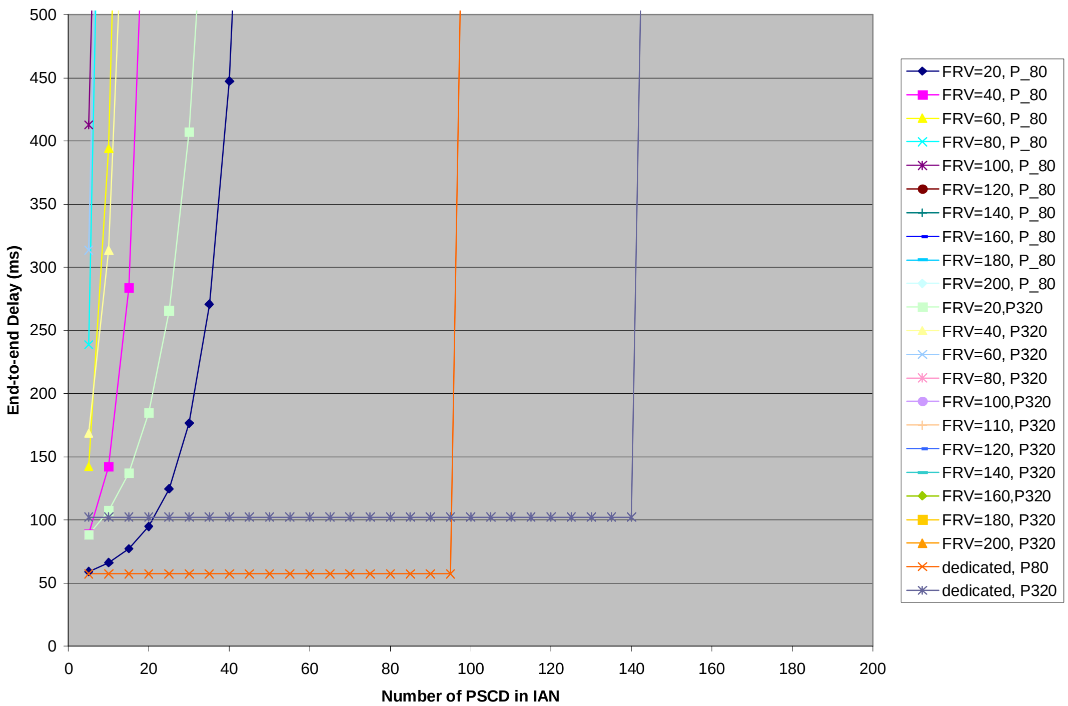

4.2 Effect of Channel Data Rate

This section describes the effect of increasing the channel data rate on the links traversed by the path. We choose 54 Mbps as an example data rate; it could potentially be used in the future network because it is part of the IEEE 802.11g high-speed wireless LAN standard. However, the network model does not depend on any particular value for the channel data rate.

Using Path C as a point of comparison, we increased the data rate on the IAN-JAN and JAN-IAN links only, while keeping the data rate on the PAN-IAN and IAN-PAN links at 11 Mbps. Similar effects occur with the other paths (A, B, D, E, F, G). Comparing Figure 6 and Figure 16 illustrates that where the wireless links use Slotted Aloha, the loss performance improves due to the additional available bandwidth. If the network is lightly loaded (20 FRVs and 5 PSCDs), a loss rate of less than 1 percent is possible.

Similarly, Figure 7 and Figure 17 show the reduction in the end-to-end delay resulting from the nearly five-fold increase in the available bandwidth. An example involves the curves associated with 20 FRVs and an 80-sample packet, which show that the maximum number of PSCDs the network can support while keeping the end-to-end delay less than 100 ms. The maximum rises from 30, when all links run at 11 Mbps, to nearly 160 when the data rate of the JAN is increased to 54 Mbps.

Figure 16: Packet Loss for Speech on Path C at 11 Mbps on IAN and 54 Mbps on JAN

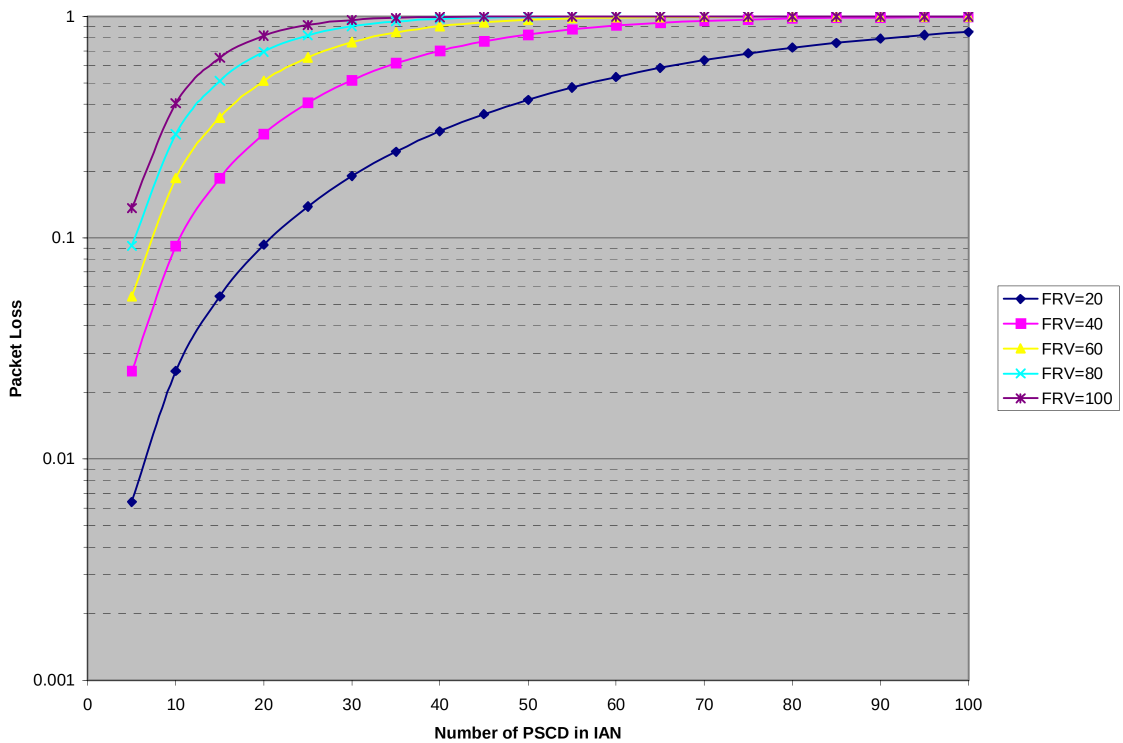

Figure 17: Delay for Speech on Path C at 11 Mbps on IAN and 54 Mbps on JAN

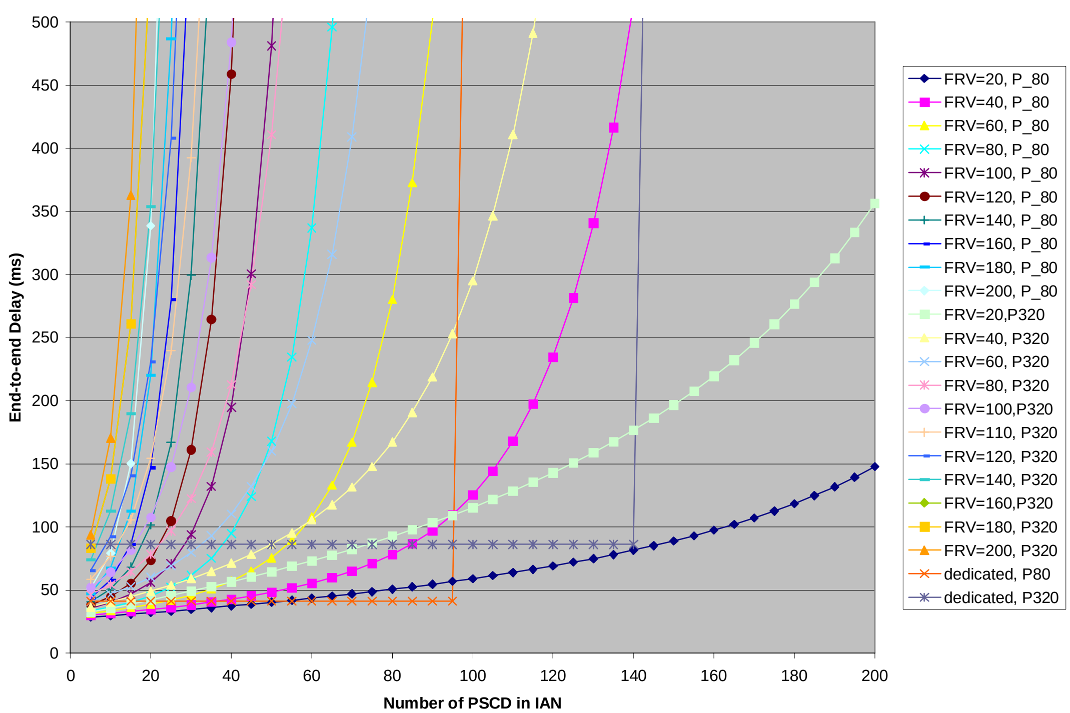

5 Example Video Application

The H.264 [3] and MPEG-2 [4] standards define different video encoding and transmission schemes. The output application data rate is modified so that it produces an average constant bit rate. For our example video application, we assume a 600-byte packet for H.264 video, and a 1358-byte packet for MPEG-2. The video application packet is encapsulated using the protocol stack of RTP (IETF RFC 3550) over UDP over IPv6. We also assume that the size of the packet containing the original video application packet does not change over any link or section in the transmission path. That is, no fragmentation takes place along any link on the path.

This section graphs the example video application to examine network behavior where the application has greater bandwidth requirements. Network performance depends on the type of video encoder and the channel data rate. As in Section 4.2, this section compares network performance over Path C.

5.1 Effect of Coding Scheme

Graphs in this section plot the loss or delay performance of H.264 and MPEG-2 video coding, where available bandwidth on every link of Path C is 11 Mbps. Figure 18 and Figure 19 show the packet loss and delay performance, respectively, of the H.264 encoder. When there are as few as 10 FRVs in the JAN, the loss performance is unacceptable in the EAN for any number of PSCDs larger than or equal to 5. This shows the unsuitability of a simple contention resolution protocol like Slotted Aloha when the channel utilization is high. Delay performance is similarly poor. Steep rises in all the delay curves indicate that acceptable performance is possible only with dedicated links or with small numbers of PSCDs and FRVs.

Figure 18: Packet Loss for H.264 Video on Path C at 11 Mbps

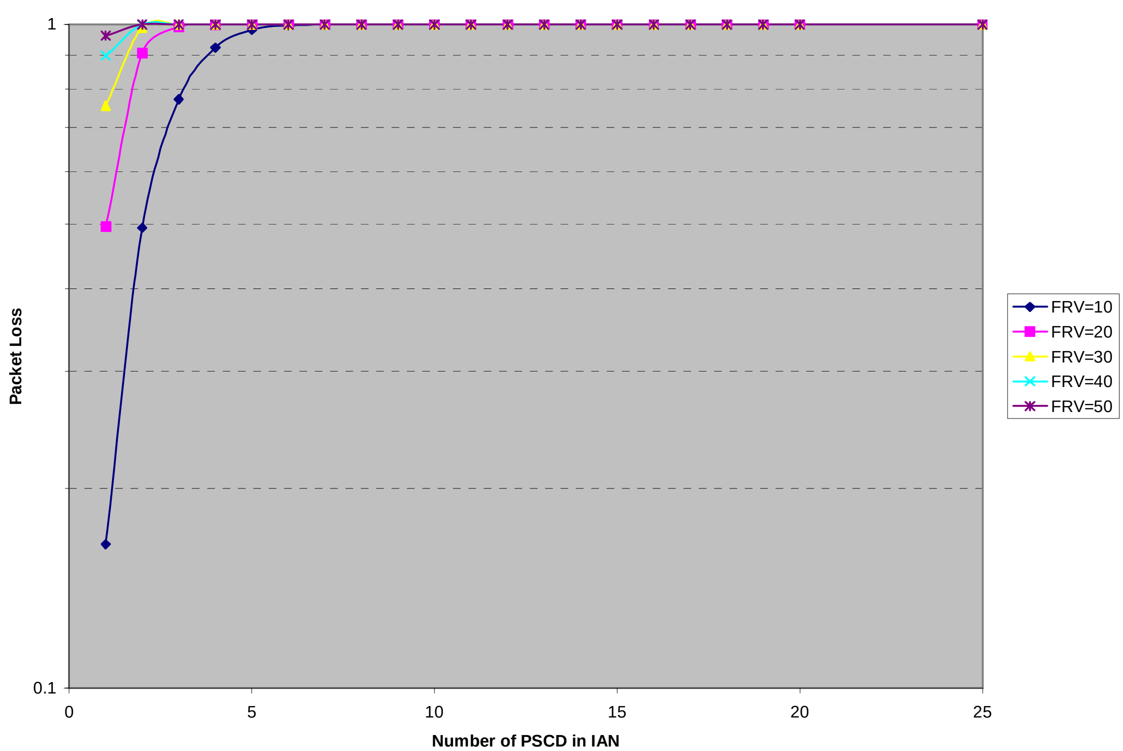

Figure 19: Delay for H.264 Video on Path C at 11 Mbps

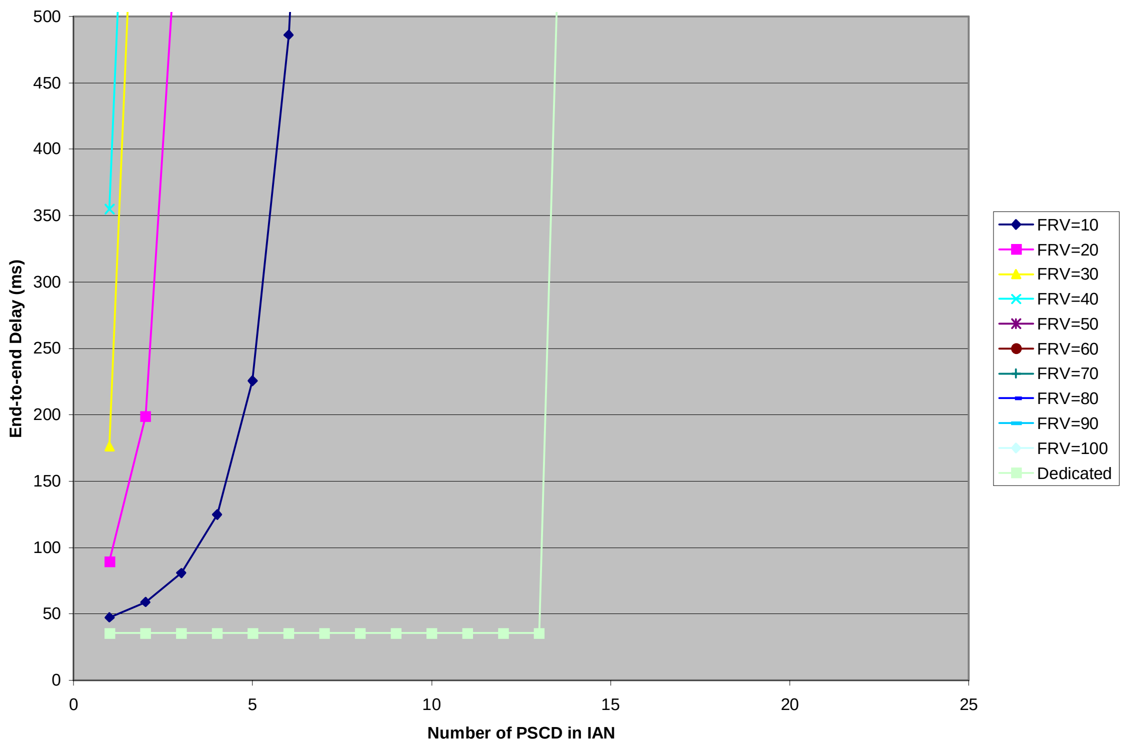

The MPEG-2 plots in Figure 20 (loss) and Figure 21 (delay) show that MPEG-2 loss and delay performance, respectively, are worse than the performance of H.264 video, because of the greater overhead associated with MPEG-2. For example, Figure 20 shows a packet loss probability of approximately 50 percent for the lowest load scenario we considered (5 PSCDs per IAN and 10 FRVs per JAN). That same scenario was the only one that yielded an average end-to-end delay that was less than 200 ms, as Figure 21 shows.

Figure 20: Packet Loss for MPEG-2 Video on Path C at 11 Mbps

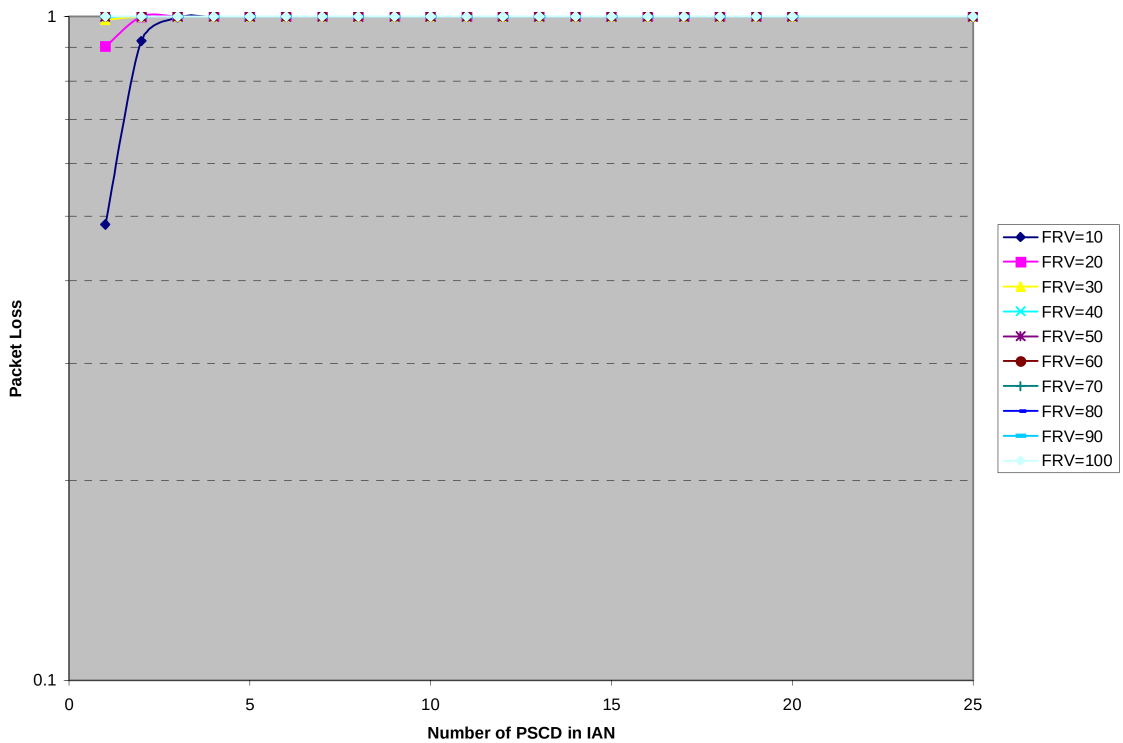

Figure 21: Delay for MPEG-2 Video on Path C at 11 Mbps

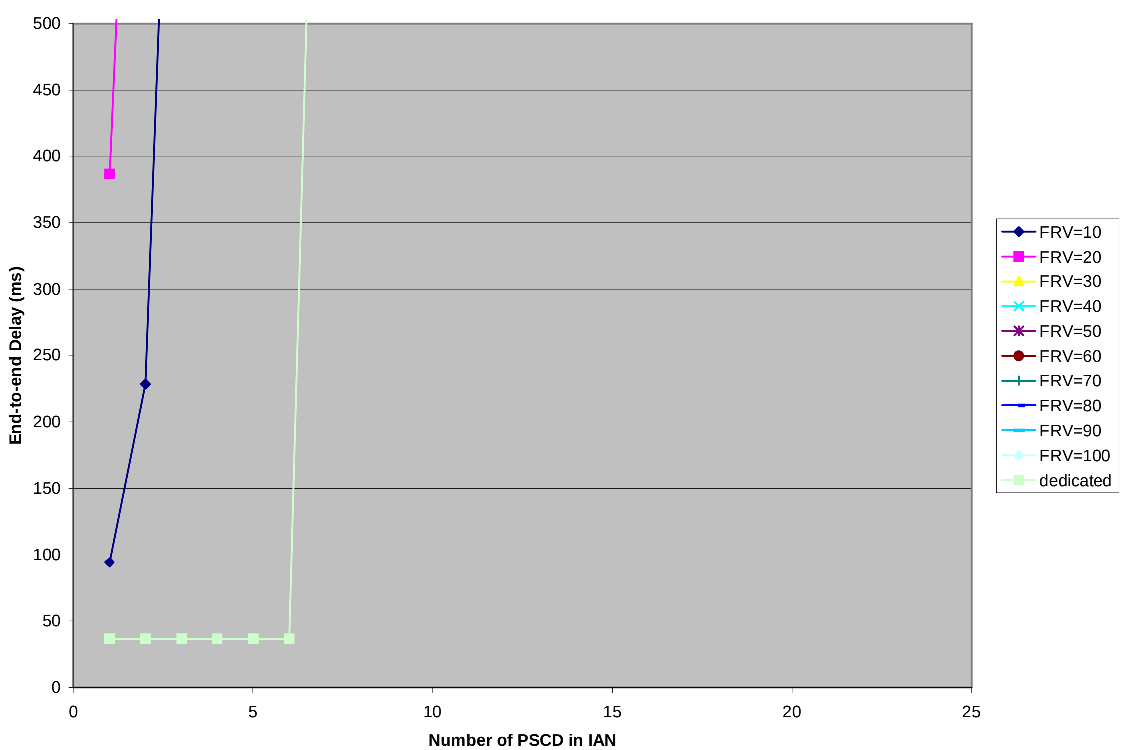

5.2 Effect of Channel Data Rate

Similar to Section 4.2, this section describes how an increase in the channel data rate can produce significant improvements in the application performance. The H.264 plots in Figure 22 (loss) and Figure 23 (delay), respectively, show the performance for H.264 video over Path C where the data rate on all links has been increased to 54 Mbps.

Although many of the usage scenarios in the graphs illustrate unacceptable performance, reasonable performance results are still obtainable when the network is lightly loaded. This means, (again, that the number of FRVs per IAN is 10, and the number of PSCDs per JAN is 5.) When the number of FRVs in the JAN is 10, it is possible to deploy nearly 20 PSCDs per IAN and still have a packet loss rate of less than 10 percent using Slotted Aloha. You can observe a similar degree of improvement in the delay performance, such that end-to-end delays of less than 100 ms are possible with 20 FRVs per JAN and 13 PSCDs per IAN.

Figure 22: Packet Loss for H.264 Video on Path C at 54 Mbps

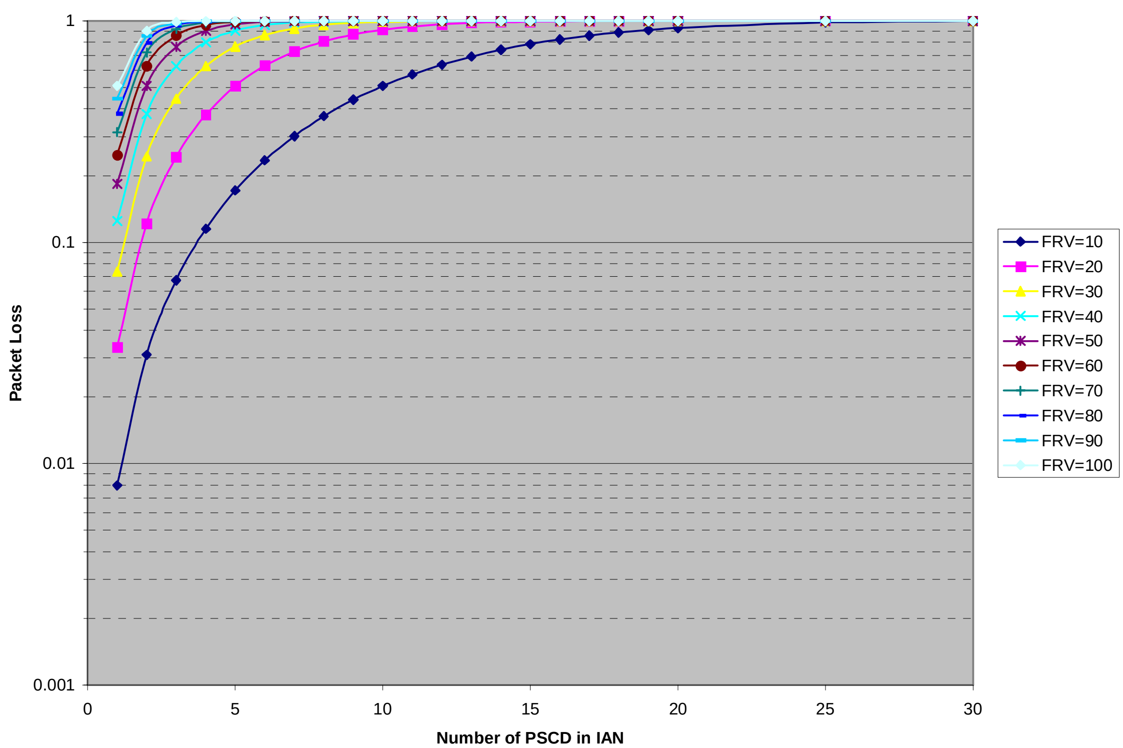

Figure 23: Delay for H.264 Video on Path C at 54 Mbps

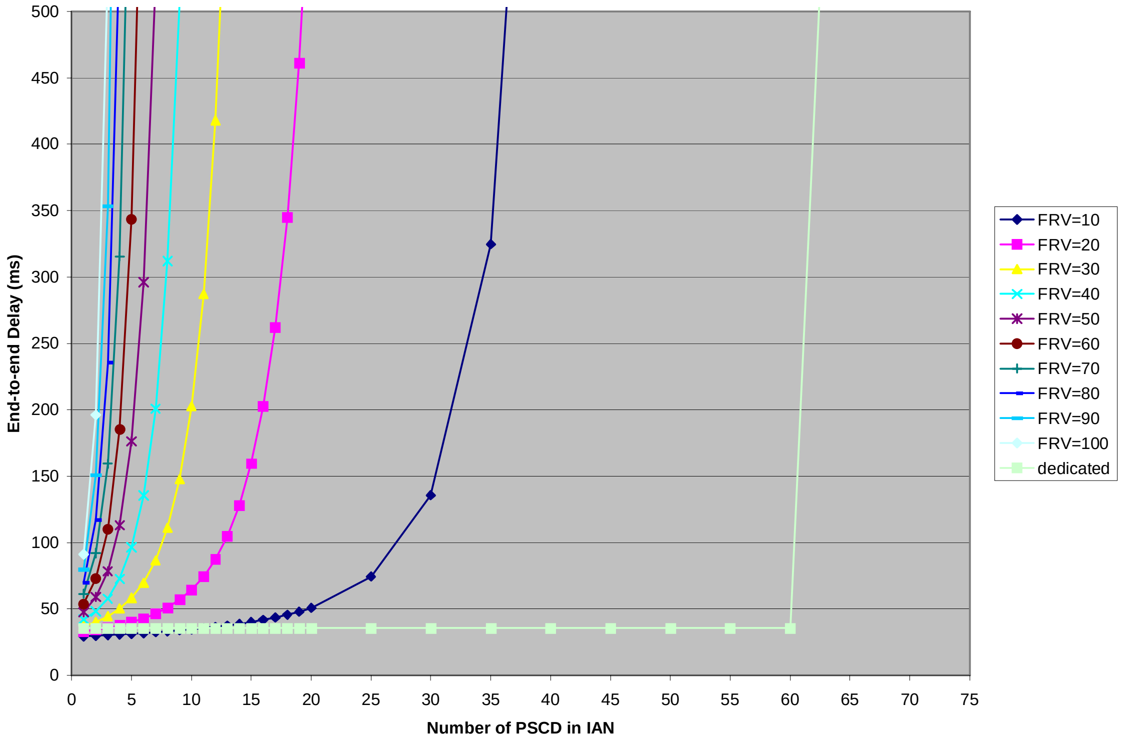

An increase in the channel data rate can produce improvements in MPEG-2 encoder application performance, although again the performance is not as good as what we can achieve with H.264. Figure 24 illustrates poor MPEG-2 encoder loss performance, but marginal performance is attainable using Slotted Aloha, although only when the network is very lightly loaded. Figure 25 shows the primary area of MPEG-2 improvement—delay. It is possible to achieve reasonable application performance using a MPEG-2 encoder even when the number of FRVs per JAN is on the order of 40, as long as the number of PSCDs per IAN remains at a reasonably small number.

Figure 24: Packet Loss for MPEG-2 Video on Path C at 54 Mbps

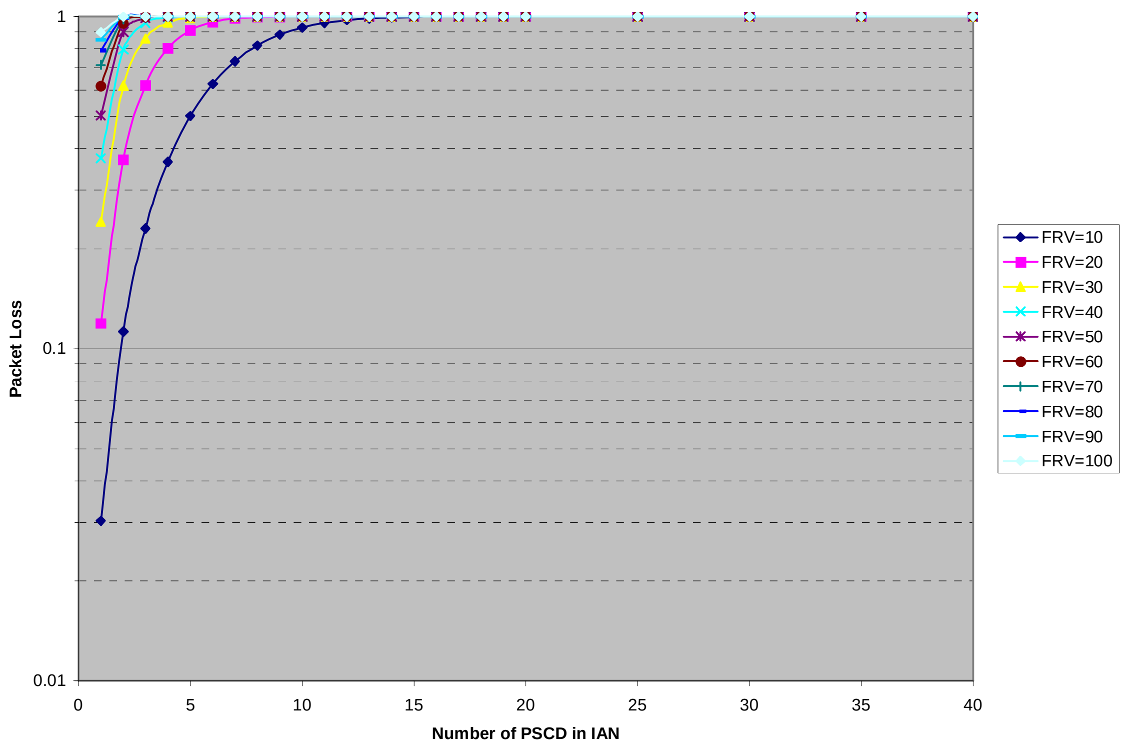

Figure 25: Delay for MPEG-2 Video on Path C at 54 Mbps

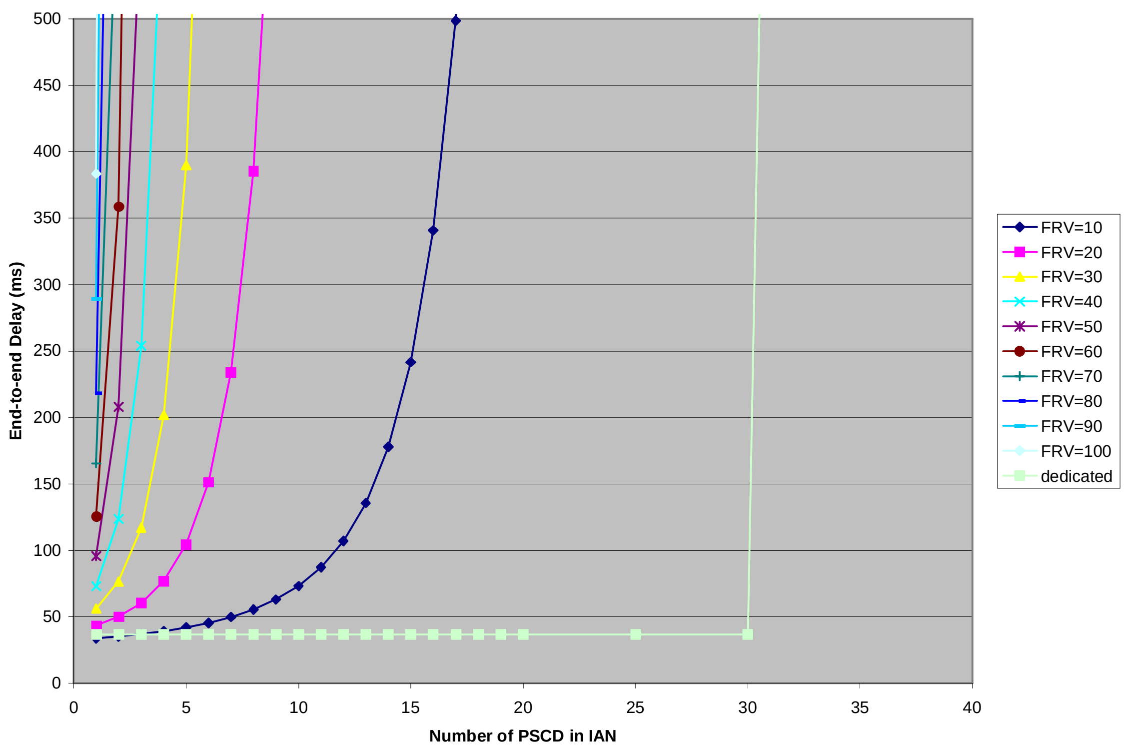

6 References

[1] *Public Safety Statement of Requirements for Communications & Interoperability*, Volume II: Quantitative, Version 1.1, November 2007. Available at: http://www.safecomprogram.gov/SAFECOM/.

[2] Recommendations of the International Telecommunication Union, Telecommunication Standardization Sector, ITU-T Recommendation G.711, 1998, "Pulse Code Modulation (PCM) of Voice Frequencies."

[3] Recommendations of the International Telecommunication Union, Radiocommunication Sector, ITU-T Recommendation H.264, 2005, "Advanced Video Coding for Generic Audiovisual Services."

[4] International Organization for Standardization, ISO/IEC 13818 (commonly known as MPEG-2), 2000, "Information Technology—Generic Coding of Moving Pictures and Associated Audio Information."

This page intentionally left blank.