



# The Impact of Packet Loss on an IPTV Network

by

Michael G. Luby  
&  
Jerry W. Miller

January 2007  
(revision 1)

## Introduction

With the introduction of digital video technology, telcos are presented with the opportunity to provide full-screen video services to consumer televisions using IPTV. However with new technology come new technical challenges. One of the most challenging for deploying quality video and audio services is data loss, due to the readily-noticeable artifacts created (see Figure 1).



Figure 1. Example of the effect of minor packet loss on a video image.

Viewers have had decades of broadcast television and years of digital transport to define their expectations for video image quality, audio quality, and presentation delays, which includes channel change times (zapping times) and startup times. The introduction of HDTV further increases the quality expectations from viewers. Any data loss can have dire consequences in trying to meet these expectations, and viewers have shown a low tolerance for poor quality services. A study was conducted at the NTT East Regional IP network in Japan over their FTTH and ADSL networks taken Dec 2003 to Jan 2004. Packet loss received was not simulated, although packet restoration at the STB was randomly assigned per user. As the packet loss rate increased, viewers showed little interest in continuing to watch as seen in Figure 2 below.




Figure 2. Telco study showing viewer response when subjected to packet loss.

Data loss, in the form of packet loss, is an inherent property of all IP networks, due to the fundamental design of IP as an economical “best-effort” packet delivery platform independent of application. When such a network cannot properly transmit packets of data, it will simply “throw out” some packets. It is then the task of the application to properly deal with the effects of lost packets to avoid causing a whole host of reception problems for the intended user service.

In IPTV roll-outs where DSL provides the last mile connectivity, customer studies have shown impulse noise that overwhelms physical layer correction resulting in packet loss. The most common reception problems are degradation to the video and audio quality, but packet loss can also lengthen channel change times and in extreme cases, it can even cause the set-top box to reboot. Depending on the level of packet loss, these reception problems can be, at minimum, an unwanted nuisance, but can often lead to an unwatchable viewing experience and highly dissatisfied customers. Studies have shown that even a small amount of uncorrected packet loss can lead to significantly less overall viewing by customers.

This paper intends to highlight the sources of such loss, some of the issues they create, and the solutions available.

## **IPTV Challenges**

Telephone companies (telcos) deploy video in a characteristic environment that can present some challenges.

- *Closed networks*—Telcos transmit video (and voice and data) services over a proprietary, landline network that they maintain, thus allowing them complete control (in most cases) over the equipment used, including equipment selection, configuration, administration, and maintenance.
- *High-quality service*—Quality benchmarks for acceptable video and audio quality have already been set by satellite and cable service providers. Telcos must, at a minimum, meet these quality expectations in order to be competitive.
- *Standards-based system*—Compression standards are primarily confined to the MPEG variants (MPEG-2 and H.264/MPEG-4 part 10), both as a result of the content providers and the equipment available (receivers, encoders, STBs, etc.).
- *Large numbers of programs*—Typical telco offerings run on the order of 100-200 video programs with 30-50 audio programs.
- *Multiple service types*—Telcos provide voice, data, and video/audio services to customers over the landline network. Especially in triple-play offerings, these services must be run in parallel. Such services have different priorities and requirements, e.g., voice services require low data loss and low latency; data services allow higher loss and longer delays due to the use of TCP; and video/audio requires extremely low data loss and very high data rates, but can tolerate greater latency than voice.
- *Multicast protocols*—Not all network equipment and customer premises equipment (CPE) can handle all of the content at the same time, so bandwidth must be constrained. Multicast via UDP provides a simple means to accomplish this.
- *Last-mile bandwidth constraints*—xDSL deployments provide specific bandwidth limitations for a given loop length from the DSLAM. This is not currently an issue with FTTH (fiber-to-the-home) deployments.

Such a system has inherent risks, such as the following. This paper will delve into how some of these risks can be mitigated or even resolved by addressing the issue of packet loss in the network.

- Self-maintained equipment places the onus for proper configuration and maintenance on the operator. Any faults or issues will affect live customers, so any protection against faults can minimize or prevent customer complaints while the fault is corrected. Additionally, any faults outside the system that impact the overall user experience must still be addressed, such as impulse noise affecting DSL lines.
- High-quality service means, among other things, rapid channel change times and consistent video and audio quality. Digital video providers continue to struggle with how to provide high-quality video and acceptable channel change times simultaneously. HDTV makes this even more difficult.
- The current state of compression equipment effectively defines the bandwidth requirements for various programs. Based on the current state of technology, the MPEG codecs require from 3-6Mbps for MPEG-2 SDTV, 14-20Mbps for MPEG-2 HDTV, 1-4Mbps for H.264/MPEG-4 SDTV, and 6-14Mbps for H.264/MPEG-4 HDTV. While some improvements will be possible as next-generation encoders are made available, current bandwidth requirements are relatively fixed for a given encoder, input type, content, and quality target.
- Large numbers of programs place large bandwidth requirements on the network in and near the headend. Typically bandwidth is constrained by the system design, so any reduction in bandwidth required near the headend provides for the addition of revenue-generating services, such as video on-demand (VoD) or additional broadcast programs.
- Services (and programs) compete for available bandwidth if not prioritized or separated properly, and this allows the possibility of undesirable packet discards due to buffer overflows. “Well-designed” networks prevent this as much as possible, but a truly “error-proof” design introduces inefficiencies, wasted bandwidth, incurs significant maintenance costs, and may require additional expensive router technologies as it contradicts the best-effort transport principle of IP.
- IP multicast, while providing low overhead and scalability, does not allow use of retransmission protocols for reliable and timely packet delivery. Reliable and timely delivery is critical to provide video and audio services. Since all of a content’s packets are not guaranteed to be delivered to requesting CPE, reliability must be built into the system at the application layer.
- Last-mile constraints ultimately confine the number, type, and quality of services available to the customer base. It is common for telcos using xDSL to specify a target download bandwidth based on a number of given services and then determine the eligible customer base. However possible, increasing the available bandwidth increases the number of available services and/or the eligible customer base for those services. It is not generally acceptable to reduce quality to fit bandwidth; rather DSL reach or the number of simultaneous programs is reduced to fit existing bandwidth capabilities.

## **Sources of Packet Loss**

Packet loss sources include the following:

- Analog and electro-magnetic interference, such as impulse noise
- Short-term transient changes in bandwidth, causing packet buffer overflow.
- Equipment issues, failures and incompatibilities

*Analog and EM interference* are typically caused by external factors, including weather and nearby electric equipment, such as radios, televisions, and kitchen blenders. Such external factors can also involve the telco’s own equipment. Due to the intermittent nature of these causes, they can exceed the physical layer correction models resulting in packet loss. DSL configuration, including noise margin and impulse noise protection, can be engineered on a per line basis to mitigate such noise interference, but this results in reductions in available bandwidth and per-line engineering involves significant operational cost.

*Transient changes in bandwidth* arise from both configuration selection as well as equipment tolerances. QoS constraints necessarily involve selecting bandwidth limits that can produce packet loss if not managed properly. Excessively jittery or bursty traffic can exceed input buffers or packet processing capabilities, resulting in packet loss, or late loss as the packet is received too late to be of use. Crucially, even if events of this nature are rare they cause significant issues for video, whereas they might remain unnoticed for other services.

*Equipment issues* include such items as bad fiber connections, spurious messaging, and improper packet handling, e.g., inserting interpacket jitter. Incompatibilities include using hardware that is standards-based yet not fully interoperable with other models or vendor equipment such that dropped packets occur due to timing issues, for example. This could occur with new technologies when field-testing is limited.

These packet loss sources produce issues for the telco as follows:

- Reduced video/audio quality
- Overengineered networks/excessive overheads
- Reduced xDSL loop lengths
- Longer channel change times
- New equipment purchases

These issues will be explored further in the next section.

## **Issues Due to Data Loss**

Ultimately, uncorrected data loss leads to packet loss at the decoder. Once packet loss reaches the decoder, it is uncorrectable, producing video and audio issues, which are detailed below. Since such issues are undesirable, various mechanisms and designs are put into place to reduce or eliminate them. However, these designs produce undesirable side-effects, also as listed below.

### Video and Audio Quality

Video and audio quality are directly affected by packet loss, since the overwhelming telco install base uses multicast over UDP (vs. RTP), which provides no protection from packet loss (unlike the TCP/IP model used, for example, for web browsing). Packet loss of audio can be exhibited as dropouts, squeaking, chirping, or skipping. With video the characteristic mild result is pixelization or blocking, with stuttering, freeze frame, and STB lockup or rebooting as major examples. For video the degree of impact is also dependent upon the frame of video affected. Since I-frames (IDR frames in MPEG-4) serve as the reference for all frames in a group of pictures (GOP), loss of part or all of an I-frame propagates and can persist for the entire GOP (typically 0.5-1 seconds). Similarly, P- and B-frames can be referenced by other frames, such that issues with these being corrupted can also persist but usually to a lesser extent and less long (potentially up to 1 second). The more flexible inter-picture prediction of MPEG-4 can worsen this effect. Packet loss of as little as  $1 \times 10^{-4}$  (or one lost packet per minute on a MPEG-4 SD program) is generally considered unviewable and one lost packet per hour (or  $2 \times 10^{-6}$ ) is considered unacceptable per the DVB standard.

### Overengineered Networks

Prevention of data loss involves designing the network to minimize packet loss and bandwidth conflicts. A “well-designed” network will allocate sufficient bandwidth and priority for the various services and programs, ensuring there is generally adequate available bandwidth, but even well-designed networks can

suffer constructive interference when multiple variable bit rate services peak at the same time. The issues in and near the headend differ from those at the last-mile, particularly for xDSL deployments, but they are primarily due to Ethernet jitter and packet-handling.

Different equipment perform time-averaging differently, such that a “constant bit rate” (CBR) stream from one device may not be sufficiently CBR for another device, resulting in the possibility of discarded packets due to exceeding packet buffers over very short times (see Figure 3). This can occur during mode conversion from Ethernet to ATM, typical in ADSL environments in and/or near the headend, and is prevented by providing excess overhead for the conversion, in the realm of 20% overhead over IP vs. the theoretical 15-17% required for AAL-5 over OC-3c [1]. This is only a problem with xDSL environments using ATM/SONET/SSH and/or QoS (quality of service), as Ethernet QoS solutions become more common. Additionally, hardware limitations such as packet-processing and packet-handling can produce discarded packets if performance requirements are exceeded, for example, passing too many packets through a particular blade or set of ports, even if only for small periods of time.




Figure 3. Example CBR stream exceeding QoS bandwidth constraint.

In the last-mile of the transport network, bandwidth is tightly constrained. Operators with QoS available will set a maximum throughput rate for each video/audio program to ensure last-mile bandwidth

thresholds are not exceeded. Unfortunately video encoders, Ethernet network gear, and most packet-handling hardware can add jitter (burstiness), which could cause individual video/audio programs to exceed the allocated packet buffers. The net result is that available bandwidth may be exceeded and packets will be discarded.

### Reduced xDSL Loop Lengths

With xDSL deployments, operators select maximum reach based on a target bandwidth, e.g., 10Mbps, which is usually selected based on a combination of intended services (voice, data at some guaranteed rate, 2 MPEG-4 SD channels, 1 MPEG-4 HD channel, etc.). Customers outside this reach are not offered the services or are offered a reduced subset, e.g., data and voice only. xDSL equipment provides for physical layer forward error correction (FEC) to correct analog and EM noise by throttling back on the allowed bandwidth. So as the amount of noise increases, the xDSL equipment will dynamically reduce the errors by increasing physical layer FEC, which has the side effect of reducing bandwidth. Less aggressive profiles (due to this extra noise) reduce overall reach and limit the customer base. Figure 4 demonstrates different loop lengths for an ADSL2plus access network using different physical/link layer protection. The increase in overall reach area is equal to the square of the increase in loop length.

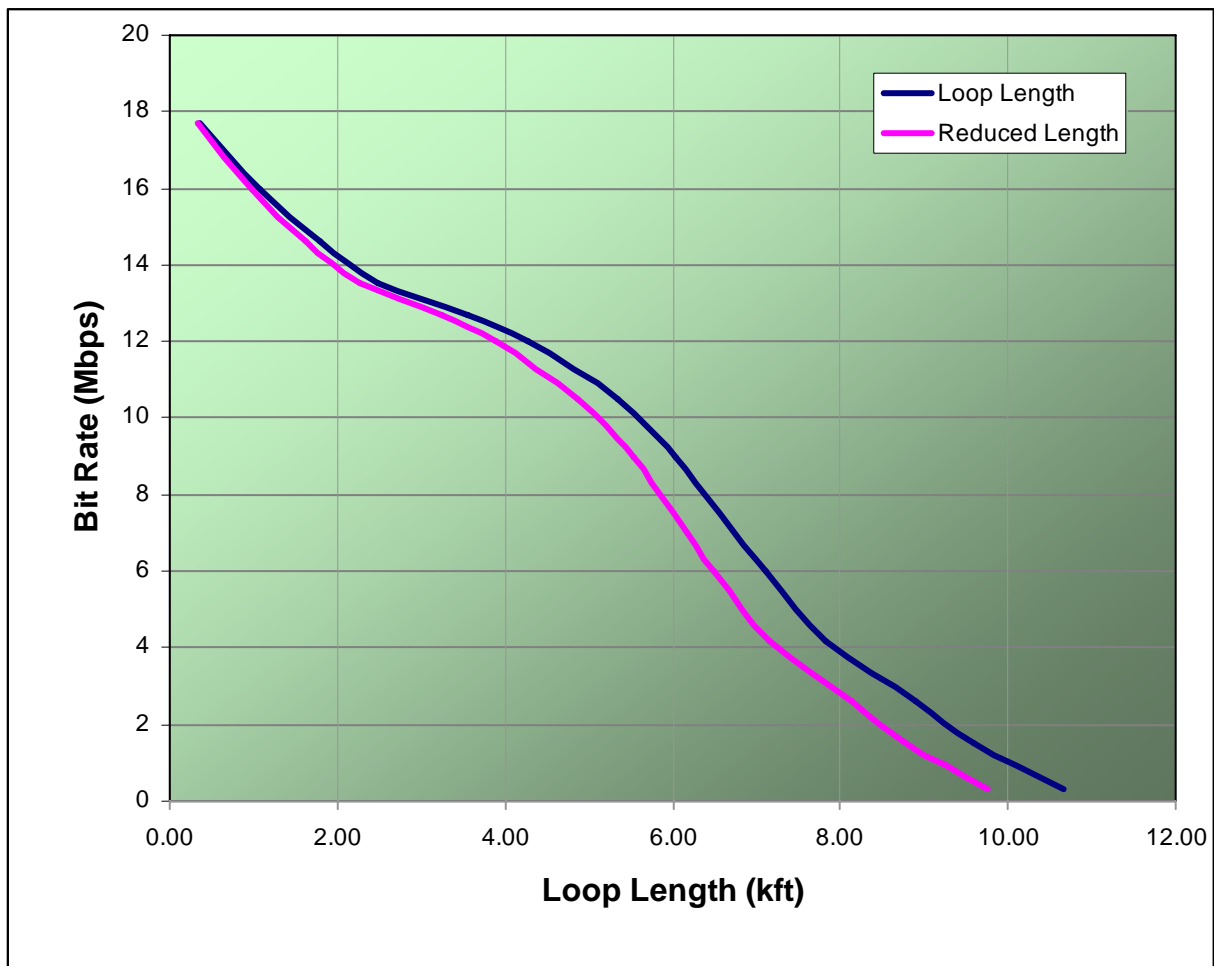


Figure 4. Reduced ADSL2plus loop length due to data loss.

### Longer Channel Change Times

During channel change, MPEG decoders wait until the next reference frame (I- or IDR-frame) before presenting the image to the viewer. Packet loss during this frame can cause the decoder to wait until the next good frame, thus significantly increasing the channel change time. In the case of MPEG-4 systems, these reference frames can be as far as 2 seconds apart (GOP of 60 for NTSC). Short channel change times have been one of the characteristic problems with digital video, specifically IPTV.

A study by Kamal Ahmed at TNO [2] provided the results shown in Figure 5. Different zapping (or channel change) times were evaluated and the Mean Opinion Score (MOS, where 5 means excellent and 3 fair) is recorded. One of the conclusions was that viewers perceived a service with a channel change time of 1 second as only “fair.” Therefore operators need to work on zapping delay and efficiency tradeoffs such as configuring shorter GOPs for improved zapping time, but at the added costs of higher bit rate, due to the larger sizes of reference frames.




Figure 5. Mean Opinion Score vs. channel change time study results

### New Equipment Purchases

Operators with existing data services could believe that their network is reliable enough for video. But data services are predominantly over TCP/IP, which provides reliability inherent to the protocol. Most video services are over UDP/IP, which provides no such reliability. Additionally, video and audio services are particularly loss-intolerant due to the high degree of compression and the sensitivity of viewers even to relatively infrequent video artifacts. So any sources of loss in an existing transport and access network need to be mitigated as much as possible. This could require new equipment in the form of newer-generation DSLAMs, different fiber runs, and new copper lines.

### **Potential Solutions**

Various solutions are currently available and implemented to varying degrees to reduce and/or eliminate packet loss.

### Physical/Link Layer Error Correction Codes

Various error correction schemes (e.g., convolutional and Reed-Solomon codes) at the physical layer provide for some protection against data loss on a single physical link, but such correction is primarily targeted at mitigating bit or byte errors and very short bursts. The typical Reed-Solomon scheme (188,204) of 16 bytes of repair data for a 188-byte source packet provides protection for only 8 bytes of loss, or a 0.03msec burst for a 2Mbps stream. Interleaving can improve this, but if the error rate exceeds the capacity of these protection mechanisms to correct, the packet is flagged as corrupt, e.g., an uncorrected Reed-Solomon error. Such packets are considered unusable or lost. Increasing the amount of protection at the physical layer can reduce packet loss to some extent, but the protection capability is limited due to the short time interleaving and the bit rate reductions required to increase the protection could become severe. Additionally any physical layer FEC, and Reed-Solomon codes in particular, are resource-intensive and usually require dedicated hardware for the computations.

### Application Layer Erasure Coding—ProMPEG COP3, DF Raptor™

Erasure coding attempts to restore lost packets missed by any physical layer error correction or lost for other reasons. Since it operates on whole packets, it can provide protection against significantly longer burst periods of seconds or longer (see Figure 6). Both ProMPEG COP3 and DF Raptor codes operate at this level, but each has distinct advantages depending on the losses experienced. Studies have shown that COP3 has significant inefficiencies above a packet loss ratio of  $\sim 10^{-4}$ , which makes it ill-suited for the consecutive packet loss found in reported bursts.




Figure 6. Erasure code repair format.

### Retransmission

Retransmission resolves packet loss by requesting lost packets. CPE decoders must determine which packets are missing and request them, usually from a distributed server that is receiving all eligible programs at once. Distributed servers are necessary to reduce the amount of buffering required on the STB due to round-trip latency, which can be significant if the STB is far from the headend. These servers must be sufficiently sized to be able to buffer all programming at once and simultaneously serve any retransmission requests. This serving can be unwieldy if multiple programs have issues at the same time, such as issues near the source causing multiple clients to experience lengthy burst loss. While a possible solution to resolve packet loss, this solution may not scale well in larger or geographically disperse deployments. Furthermore, the added delay required to allow retransmission adversely affects overall delay and therefore zapping times.

### Overengineering/Overbuilding

As packet loss is such a problem, operators typically scrutinize their transport networks to eliminate noise and reduce packet loss wherever possible. Fibers will be reanalyzed and/or reterminated, questionable connectors will be replaced, and problematic boxes and/or blades will be swapped out, all as part of an effort to minimize data loss in the system. This can manifest itself as unusually high allocated

bandwidths to prevent oversubscription (see Figure 7), extra physical layer correction to reduce the number of uncorrected Reed-Solomon errors, shorter loops than necessary to ensure sufficient bandwidth, underutilizing ports to ensure packet-handling limits are not exceeded, and extreme diligence in testing, cleaning, recutting, etc. fibers and optical equipment.

In spite of this effort, burstiness will occur, so operators must overallocate bandwidth to prevent exceeding maximum channel limits, simply to prevent occasional, or even highly infrequent, burstiness that could lead to packet loss. Additionally, reach will be reduced to account for intermittent impulse noise. In short, the system is overengineered, swapping efficiency and reach for a positive user experience. Usually this results in more conservatism than necessary. While operators must be careful to ensure sufficient quality throughout the system, excessive conservatism is not necessarily the best answer.




Figure 7. Green area shows unneeded bandwidth allocated for a VBR stream to prevent buffer overflows at 3Mbps.

### Tolerance

Despite the above, some minimum of packet loss may still reach the end-user. If such loss is sufficiently infrequent (e.g., 1/month) and sufficiently minor, operators may choose to tolerate the problem, but ultimately it is up to the customer base to determine what is and is not acceptable quality. If a sufficiently large number of subscribers complain forcefully enough (or drop the service), the operator's loss tolerance must be revised or churn will occur.

## Accurately Measuring Loss

As loss is an inherent part of IP networks and has such detrimental effects on the user experience, it is necessary to quantify the loss being experienced in a network. Key to this is accurately measuring the loss so as to be able to correctly mitigate or eliminate it, and a single loss number does not sufficiently characterize the problem.

Excess packets discarded in the manners above are relatively straightforward to find in a network: simply monitor for CC (continuity counter) errors over a sufficiently long period of time or use any of a wide range of MPEG/IP or IP measuring tools and software available, but what is vital with this measurement is to perform the testing properly, and time-averaging plays an important role in measuring results (see Figure 8 below). The temptation to measure long-term averages of loss, e.g., over hours or days, must be avoided, as this will hide many of the problems above, e.g., impulse noise, which tends to occur in bursts of ~10-50 milliseconds. Such a burst would be hidden in a one hour sample as merely a 0.001% packet loss, and the operator may believe there is no problem of concern. But to viewers, the video tiling resulting from a loss every hour could quickly lead to an undesirable viewing experience, regardless of the overall packet loss number. Operators must be careful to measure and analyze both the duration of any losses and the frequency of these versus simply capturing all losses over long periods of time for a tidy average. Only the avoidance of all packet loss can ensure predictable service quality. Errors always make the service quality unpredictable.




Figure 8. The effects of different time-averaging on packet loss.

## Benefits of Erasure Coding

One of the solutions mentioned above targets the loss problems still being experienced—a relatively new, promising loss recovery technique called erasure correction, also referred to as application layer forward error correction (AL-FEC). Erasure coding has won recent adoption in such standards as DVB-H and 3GPP in the mobile space and is soon to become included in the DVB-IP specification for IPTV. Erasure coding differs from classical bit-level error correction in that the unit blocks tend to be much larger (whole packets vs. single bits) and the entire unit (or packet) is assumed lost or unrecoverable. For these reasons, erasure correction can provide the following main benefits of protecting against burst loss and supplementing physical/link layer error correction. By doing so, many of the data loss issues mentioned above can be mitigated or even eliminated to provide such benefits as

- Eliminate packet loss
- Reduce overengineering
- Reduce QoS overheads

- Extend loop lengths
- Add service-specific protection

And while there are several erasure coding algorithms available, the most applicable for video over IP are DF Raptor and ProMPEG COP3, but COP3 has certain limitations that make it less advantageous, primarily

- Raptor is significantly more efficient with overhead than COP3 for loss rates above roughly  $10^{-4}$ , which would be required to protect against the typical burst losses of 10-30msec reported. This is in large part due to COP3's 2-D interleaving architecture.
- Raptor provides flexibility for changing loss environments, such as repaired links and replaced equipment, whereas COP3 must be customized to a particular loss scenario.
- Raptor provides extensibility to support expected upcoming IPTV services such as content download delivery services or firmware updates and is already integrated in such services, whereas COP3 lacks such support.
- Raptor has proven to be applicable and highly efficient in IPTV-mobile TV convergence, whereas COP3 is not designed for the expected convergence.

In the end, erasure correction can be used to ensure a positive customer experience against data loss.

## **Conclusion**

Historically, telcos have relied upon a vigilant discovery and elimination of data loss sources to maximize IPTV service quality, but even in the best of circumstances, packet loss will occur. If it occurs, even at relatively low rates, customers will be unsatisfied with the experience and switch content providers ("churn"). Operators must choose some or all of the solutions mentioned to recover from packet loss, preferably without overly constraining their customer market. In practice, use of the loss recovery techniques described here will be both more effective and more cost-effective than network engineering alone.

## **About Digital Fountain**

As experts in broadcast and real-time data transport, Digital Fountain software optimizes the delivery of digital media over any network. Our technology eliminates the common limitations associated with digital media distribution solutions over both private and public networks, maximizing our customers' existing infrastructure investments, and enabling expanded revenue opportunities. Our proprietary DF Raptor™ technology redefines the science of forward error correction (FEC) by offering a universal solution that can accommodate an unlimited range of network conditions. Digital Fountain technologies are used today by leading mobile carriers, IPTV providers, and national defense agencies throughout the world, and have been standardized by leading international standards bodies, including DVB, 3GPP, and IETF. Our partners and customers include leading global companies such as Cisco Systems, Sumitomo Electric Networks, Scientific Atlanta, Northrop Grumman, Pioneer, KDDI, Sirius Satellite Radio, XM Radio, Sony, Nokia, Adobe, and many more. For more information, please visit our website at [www.digitalfountain.com](http://www.digitalfountain.com)

## **REFERENCES**

- [1] Cavanaugh, J. D., "Protocol Overhead in IP/ATM Networks," Minnesota Supercomputer Technical Report, Minneapolis, MN, August 12, 1994.
- [2] Admed, Kamal, TNO, ITU-T IPTV Global Technical Workshop, Seoul, Korea, Oct. 13, 2006.