TRANSPORT OF CBR TRAFFIC ON IP NETWORKS

Tom Carr, Bruce Roe
Wave7 Optics

Abstract

The success of future distribution networks will depend on their ability to support legacy services including committed bit rate traffic. Most of this traffic is transported by the PSTNs over T1 facilities.

This paper describes a technology, Time Division Multiplexing over IP, which is capable of providing T1 circuit emulation over IP networks.

INTRODUCTION

Industry visionaries foresee Next Generation Networks that offer hundreds of megabits of bandwidth to the consumer, extensively or exclusively using IP as the network transport protocol. This is an easy vision to believe in, and one that we are convinced will evolve into reality.

While we can envision our final destination, the question remains, how do we get there? Clearly the distribution network cannot, and will not, be replaced en masse with a new IP-based architecture. Pockets of new high-speed IP networks will be deployed and, over a very long period of time, finally consign the old copper telecommunications plant to the pages of history. This slow evolution means that some legacy services must be supported on the new architectures.

One likely architecture for future distribution networks is that shown in Figure 1. An Ethernet network of one or more gigabits extends over fiber from a Head-End to an active bandwidth management element. From there, the Ethernet is extended over multiple fibers to serve a pocket of customers. By using

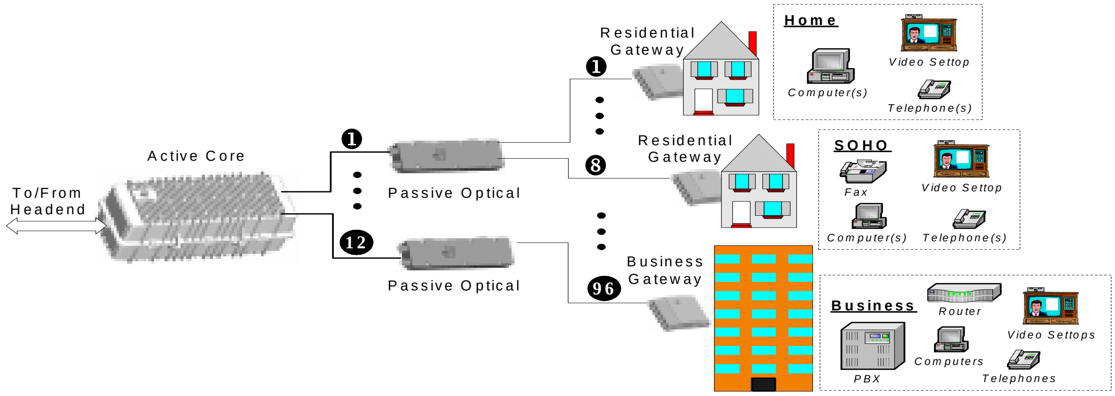

Figure 1: Next Generation IP-based Distribution Network

WDM, RF digital and analog video can be distributed over the same fibers. Plain old telephone service (POTS) would be supported by using voice over IP (VoIP) for transport. In this manner voice, video and very high-speed data could be offered over a single efficient network. Unfortunately, this alone will not provide for a significant and lucrative portion of traditional legacy telecommunication services, committed bit rate services.

Obituaries have been written for committed bit rate traffic such as private line service, international frame relay networks, and every other non-IP protocol, but the double-digit compounded growth rate for these services, particularly in the international markets, continues. Some forecasts[1] of international frame relay, for example, predict a compounded annual growth rate of 14-16% at least through 2004. Even X.25 networks still exist and continue to grow. The inertia of migrating these networks to IP will be fueled by sluggish economies, and the falling prices of both T1 service and old technology equipment. Any new distribution network, particularly those limited in geographical scope, must either accommodate these legacy services or exclude large, profitable markets. It is simply not economical for end-users to convert national or international non-IP networks to IP in a piecemeal fashion.

PBXs present another problem: signaling. Signaling consists of basic features such as recognizing that the phone is off-hook, or needs to ring; the more advanced properties required for reaching the proper destination and billing; and still more sophisticated characteristics, such as caller identification, call forwarding, and conference calls. There are literally thousands of such telephony features, with dozens of national and local variations. Available VoIP integrated circuits can handle some, but not all, of the PBX signaling in use in the U.S.. Converting PBX voice circuits to VoIP could require the end-user to give up some useful or much needed features.

The one common element of committed bit rate traffic carried by PSTNs is that they are primarily transported via T1s. Having the ability to transport T1s over IP, regardless of the data or signaling protocol, would be an ideal solution for supporting legacy services in a new IP environment. Such a technology does exist and is called TDMoIP, Time Division Multiplexing over IP. TDMoIP is a technology that combines features from Time Division Multiplexing and IP to deliver synchronous T1 circuits ***transparently*** over IP networks. An individual channel within the T1 stream is not changed in any way, nor is there any signaling conversion. This technology would be used point-to-point, from the customer's premises to the Head-End.

TDMoIP OVERVIEW

A T1 frame is composed of 24, single byte time slots plus a single synchronization bit, for a total of 193 bits. Frames are transmitted at a rate of 8000 per second, resulting in a data stream of 1.54 megabits per second. In principle, the simplest implementation of TDMoIP simply encapsulates a number of T1 frames in an IP packet by tacking on the appropriate IP header. At the destination, the stream is then recreated by stripping away the headers and reassembling the segments. *It is important to note that TDMoIP transports the T1 circuit without any attempt at interpreting the data.* This process is oblivious to signaling, time slots, or whether voice or data are being transmitted. This also implies that a data bit-stream using the entire 193 bit frame can be supported.

Standards

TDMoIP is essentially the IP counterpart of the same service in ATM referred to as "Circuit Emulation Service," (CES). While there are, as yet, no standards[2] for TDMoIP, such standards do exist for ATM-CES. Furthermore, since the performance requirements for TDM are independent of the method of transport, it is clear that the performance requirements for ATM-CES should be adhered to as closely as possible in TDMoIP. Nonetheless, *how* to achieve that has not been standardized, and will likely vary between TDMoIP platform vendors. Thus, the interoperability of equipment from different vendors should not be expected.

The TDM performance guidelines to follow primarily relate to clocking. The clock rate of the TDM stream should be stable to within +/- 32 ppm[3] and wander should not exceed 80µsec per day[4]. Performance standards directly related to IP networks, such as the maximum allowable packet loss, do not yet exist. These standards are far stricter than what is required when terminating TDM on end-user equipment and it is entirely possible that new standards may be formulated for this specific purpose.

Packetization

Primary issues to resolve include which IP protocol to use and how many T1 frames should be placed in each IP packet. Since there is no standard, any IP protocol could be used. Some, however, would clearly be inappropriate. The end-to-end reliability offered by TCP, for example, is not useful for voice packets, since re-transmitted voice packets will reach the receiving side out of order, only to be dropped anyway due to delay constraints. A good choice of protocol could be RTP and the associated RTCP, which in certain networks would offer better clocking functions. Ultimately, for Ethernet networks, only UDP is fundamentally needed.

There are tradeoffs to be considered with selecting the number of T1 frames per IP packet. The fewer the frames, the greater the IP overhead, which will increase the amount of bandwidth needed per T1. The greater the number of frames, the greater the end-to-end delay, packet loss becomes more onerous, and larger buffers are required. Larger number of frames per packet could also exacerbate adaptive clock wander. QoS demands that the number of frames per packet be kept small despite the overhead penalty. At four frames per packet, this penalty is about 50%, meaning a 1.54 Mb/s T1 would require more than 2.3 Mb/s bandwidth in IP. Even with this overhead, a 1 Gb Ethernet network is capable of supporting several hundred T1s.

Signaling

There are three primary types of signaling: in-band signaling, channel associated signaling, and common channel signaling. None of these are impacted by TDMoIP. In -band, as the name suggests, is signaling in the audio band of speech. The ubiquitous 'touch tone', or Dual Tone, Multiple Frequency (DTMF) is an example of in-band signaling. Since these tones are encoded in the T1 frame time slots, they are automatically carried over TDMoIP.

Channel associated signaling is also carried within the T1 frame time slots. Specific voice bits are 'robbed' and the signaling bits are substituted. TDMoIP does not distinguish between bits used for voice and data bits, thus this signaling is carried transparently.

Primary Rate ISDN signaling, PRI, is a popular type of common channel signaling. The twenty-fourth time slot of the T1 frame is used to carry the signaling data for the other twenty-three time slots. Again, since TDMoIP does not distinguish between voice and data, the signaling is carried transparently.

Clocking

Clocking is the most difficult problem to solve in deploying TDMoIP. There are several methods of clocking in any type network. These are: independent Stratum 1 clocks at each end-point; a synchronous network in which the primary reference source (PRS) clock is distributed throughout; an asynchronous network in which a network clock is distributed through out; and adaptive clocking in an asynchronous network with no distributed clock.

Timing is not provided in IP networks, thus synchronization must be achieved from an external source. This can be accomplished by: a Stratum 1 external master clock at each end of the TDMoIP circuit; clocking from an external clocking distribution network; or in-band clock recovery and regeneration, i.e., adaptive clocking. Stratum 1 clocks are so precise that T1 streams timed by separate Stratum 1s will be synchronized. This is a relatively costly solution, although the prices, particularly GPS-based Stratum 1s, have been declining recently.

An external distribution network for clocking is also an expensive solution, and severely compromises the entire concept of having a single network to maintain. In this scenario, a separate network would be maintained just to send clocking signals to every end-point.

In-band clock recovery and regeneration, or adaptive clocking, is the most cost effective and will meet the requirements of customer premise equipment such as PBXs. In adaptive clocking, the source TDMoIP unit, which is clocked to a Primary Reference Source, simply sends the data to the customer TDMoIP unit. The customer unit writes data to the segmentation and re-assembly (SAR) buffer and reads it with the local clock. The level of the SAR buffer controls the output frequency of the local clock by continuously measuring the fill level around the median position and feeding this measurement to drive a Phase Lock Loop (PLL), which in turn drives the local clock. Thus, the local clock frequency is modified to keep the re-assembly buffer depth constant. When the TDMoIP unit senses that its SAR buffer is filling up, it increases the clock rate. When the unit senses that the SAR buffer is emptying, it decreases the clock rate. Since the packet arrival rate is directly dependent upon the packet transmission rate established by the PRS at the head-end, synchronization of the TDM stream is maintained.

The proper choice of buffer size can prevent buffer overflow and underflow, and at the same time, control delay (greater buffer sized implies greater delay). The buffer size is proportional to the maximum packet delay variation. This variation should be determined by summing the delay variation of each network device in the circuit path. The sum of the measured delay variations that each piece of equipment introduces must be smaller than the maximum packet delay variation configured on the TDMoIP unit. If not, underflows and overflows will occur. This buffering will also remove any jitter encountered by packets arriving at slightly different time intervals.

In the event a TDMoIP packet is lost during transport, a dummy packet is transmitted by the customer TDMoIP unit in order to maintain clocking in the T1 output stream. Since the packet will contain no customer data, it must still be considered a frame-slip, however, timing problems will be minimized.

TDMoIP and Network Delay

A T1 frame represents .125 milliseconds of real time. Processing time for packetization and recreation of the T1 is less than five milliseconds. Delay is not an issue in TDMoIP unless each TDMoIP packet contains a great many T1 frames.

End-to-end round-trip network delays greater than 30 milliseconds could necessitate the need for echo cancellation on voice circuits. Round-trip delays of more that 300 milliseconds will result in unacceptable QoS for conversational speech. This is the same as for VoIP service. In both TDMoIP and VoIP, quality can only be provided and assured on tightly managed networks with well-executed prioritization procedures.

Packet Loss

Sequence bits are used to determine a lost packet condition. In the event of lost packets, timing is maintained through the insertion of dummy packets carrying appropriate framing bits. It is possible to mitigate voice quality impairments by repeating the frames that preceded the lost packet. However, the total loss of several contiguous T1 frames would not significantly degrade a voice circuit.

VoIP packet loss could be used as a guideline. Unfortunately, various studies show that while some deployments can sustain a 5% packet loss before realizing a significant degradation of QoS, other deployments can suffer less than .2% packet loss. Obviously packet loss must be minimized. As with network delay, this implies that well conceived prioritization methods must be used within the IP network, with TDMoIP given the highest possible priority.

Prioritization

Properly prioritizing the TDMoIP packets will minimize network delay and packet loss. This is critical for maintaining satisfactory QoS. By marking TDMoIP packets they may be easily identified and prioritized. This is done through proper marking of the Type of Service (ToS) bits, and VLAN tagging and priority labeling according to IEEE 802.1 p&q. Additionally, there is an assigned, IANA-registered UDP socket number for TDMoIP. These features simplify flow classification through switches and routers. In a tightly managed network, the QoS of TDMoIP should be equal to that of a traditional T1 circuit.

Supported Features

TDMoIP is capable of supporting unframed T1, Super Framing (SF), Extended Super Framing (ESF), as well as Channel Associated Signaling (CAS) and Common Channel Signaling (CCS), including Primary Rate ISDN (PRI).

There are two solutions for supporting fractional T1s. The first is to put a large multiple of the individual time-slots in the same TDMoIP packet. This would reduce the overhead penalty. However, as discussed earlier, this would also increase delay as well as create major QoS problems in the event of a lost packet. The second method is to transport the fractional T1 as though it were a full T1, filling the unused time-slots with idle code. This would require using the same amount of network resources for a fractional T1 as a full T1.

Deployment

A single, small end-user integrated access device can be configured to house a video port, Ethernet port, and multiple POTS and TDMoIP ports. The TDMoIP packets would be routed through the network to a Head-End TDMoIP unit. A T1 circuit, identical to the end-user's original T1, would be generated by the Head-End unit.

TDMoIP may be deployed in overbuilds by network service providers or in a greenfield environment. There are three primary methods of provisioning the T1 service. The first is to simply port the T1 directly to a TDM service provider, essentially just leasing T1 service from another carrier and reselling the service. If the Ethernet network provider has a Class 5 circuit switch, then the T1s would be terminated on that switch.

In a greenfield environment in which a softswitch and PSTN media gateway are deployed for VoIP service, the T1s may be terminated on the PSTN media gateway.

The Figure 2 illustrates how a Head-End configuration of such a network.

SUMMARY

Future distribution networks must support legacy services such as T1 to be viable in the marketplace. If these networks are packet based, as is expected, technology must be deployed that will emulate a TDM T1 circuit. We have presented a straightforward method of providing such a service through the use of Time Division Multiplexing over IP (TDMoIP). Since the T1 stream is carried transparently in TDMoIP, any PBX signaling protocol or data format could be accommodated.

Careful control of prioritization through the use of existing standards can minimize delay and packet loss. The resulting T1 circuit emulation service can equal that of existing TDM technology, and easily meet the requirements of end-users.

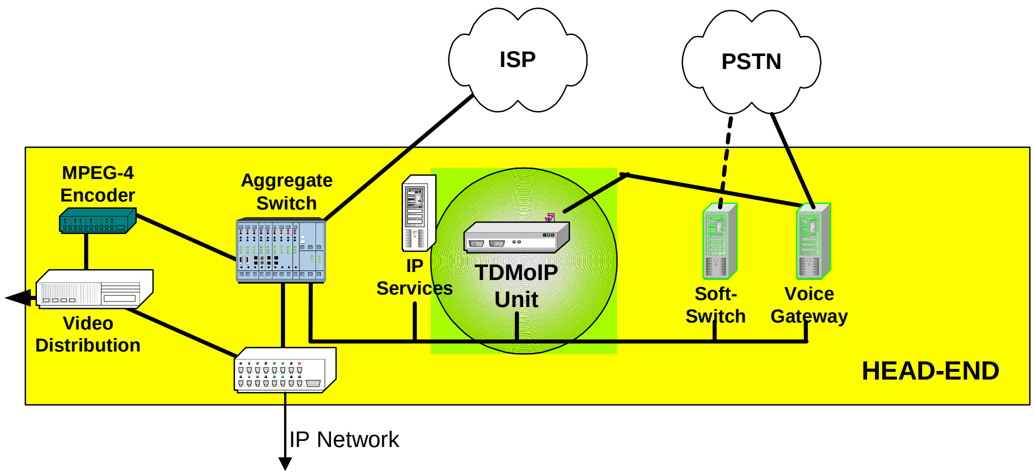

Figure 2: Greenfield Deployment Head-End Architecture

---

[1] 2000 Perspective on the Telecom Marketplace, Provisioning of Private Line and Frame Relay Services: A Global Perspective 1999-2004; Insight Corp.

[2] TDM over IP, Internet Draft August, 2001

[3] American National Standards Institute (ANSI) T1.403.1995

[4] ITU (CCITT) G.823 and G.824