

ISDN vs. DSL REVISITED

Rolf Taylor

Product Manager, Telephony
Telos Systems
Cleveland, Ohio

WHAT DOES IT TAKE TO MAKE IP WORK FOR REMOTES?

Nearly eight years ago I wrote a paper titled “ISDN versus DSL — The real truth about high-speed connections”. At the time, I was hesitant to compare these completely different technologies, however questions from our codec users persuaded me to do so. At that time my assessment of the prospect of real-time (low delay) audio over DSL was not very optimistic.

Since that time, many things have happened: We developed the Zephyr Xstream and included support for MPEG IP streaming; various forms of “xDSL” have proliferated and dropped in cost; we developed Livewire Audio-Over-IP technology and launched the Axia division to bring it to those who were “beating down the door” looking for a better mousetrap; and, we’ve gained considerable experience in the field from Xstream users streaming over IP links. ISDN is still a perfect fit for broadcasters, but with the proper technology IP is becoming more useful every day.

WHAT IS “IP”?

Despite all of the above, Telos has only cautiously advocated IP codec use, and indeed many have forgotten that this ability is included in our Zephyr Xstream. The reason for our caution is that while we are firm believers in Audio-over-IP, the term “IP” covers a lot of ground, and is often misunderstood. It is fairly simple to packetize the data from an MPEG codec, send it out over an “IP network”, and at the receiving end include a buffer. And indeed we’ve seen such offerings from most of the usual codec companies. Is this enough?

This simple approach is generally adequate for use on a switched Ethernet local area network, where each device can control the amount of data

on its port, and a non-blocking Ethernet switch can handle the rest. When connecting to larger scale IP networks that are *not shared with other applications* this may also be suitable, though the error rate of the link becomes an additional consideration.

New IP services are now available that include multiple classes of service with Quality of Service (QoS) guarantees for each. These offer a controlled environment that will generally work well with generic IP codecs and even allow them to coexist with other data. These services are usually based on Multi-Protocol-Label-Switching. Telcos are marketing these services as a replacement for Frame Relay in situations where companies’ desire fully meshed private virtual networks. This approach offers the advantages of a fully meshed network (e.g. data can be exchanged directly to any site on the network) while allowing a degree of control so that IP voice telephony services (VoIP) can operate despite the existence of other data on the link. To achieve this the network provider must engineer their network with this in mind, and must include active surveillance to dynamically manage each class of data such that the QoS guarantees are met. This is quite similar to how traffic engineering works on the dial-up telephone network. Luckily, the requirements for MPEG codecs mirror those of VoIP applications, so we can make use of these new networks.

However, when most people talk about “IP Codecs” they are thinking about using the Internet. While this is what people think of first, it is the worst case scenario, particularly if you consider the connectivity types most wanted, namely low cost xDSL and of course WiFi. To get back to our comparison between ISDN and DSL, this assumption holds true, since we have yet to see Telcos offer point-to-point xDSL connections. Instead xDSL lines connect to the Telco’s ISP, meaning that at best there is shared bandwidth at the ISP, and at worst the data travels the Internet itself.

CIRCUIT SWITCHING VS. PACKET SWITCHING

So just what is the difference between ISDN and the Internet for audio delivery? Both are networks with multiple users spanning the globe. What makes one better than the other? The biggest difference is that the telephone network (including ISDN) uses “circuit switched channels” whereas IP networks use “packet switched” technology. The way these two types of networks deal with congestion is one important difference. Let’s examine each in turn.

CIRCUIT SWITCHED NETWORKS

A circuit switched network consists of many bi-directional channel elements. Each of these elements (generally the term “trunk” applies) is either “idle” or “in use”. Since the mid 1980s, the channels and associated switching have been digital, making the deployment of ISDN possible. Each of these “DSO” channels has 64 kilobits per second (kbps) of digital data capacity. When establishing a “dialed” connection therefore, one either has an end-to-end connection at 64kbps, or one does not. The network also provides a highly stable clock that is used to synchronize the sending and receiving functions, thereby eliminating the need for all but the smallest amount of buffering. Furthermore, the standards call for a low error rate. So far so good, but what happens when the network is very busy (called “congestion”)? In this case, there may not be available channel elements to establish the requested connection. In this case one gets a message indicating “network unavailable try again later” (fast busy). Of course one can repeatedly dial the call until a connection is available - the big advantage is that once you get that connection, it is yours (with data in both directions traveling the same route, and every bit traveling this route) until you decide to “hang up”. The downside is that for this exclusive use of a channel you pay by the minute.

PACKET SWITCHED NETWORKS

This is in stark contrast to a packet switched network. With these the data stream is divided into discrete pieces called “packets,” and then these packets are sent into the network. As each packet traverses the network, second by second decisions are made by the network “routers” as to the best route to the final destination. It is not unusual for different packets to traverse the network via different paths. And packets in the return direction (if any) take their own independently determined paths as conditions allow. If there is insufficient downstream capacity, then a router may discard packets. In fact, occasional discarded packets are not considered unusual in packet switched

networks. The system does have a degree of self-regulation — a device that floods a connection will back down its speed, but this process is not instantaneous, and of course other devices are continually changing their bandwidth requirements as well. The term used for this sort of network is “best effort,” and for “bursty” data such as web pages, it is a highly effective way to share data resources.

The best known Internet Protocol, TCP/IP, allows for the sending device to resend lost packets. Other protocols do not necessarily support this, and waiting for the re-sent packets means longer delay in any case. Since the amount of traffic on any network fluctuates day-to-day as well as minute-to-minute, unless such a network is “managed” (and the Internet is not) the result is that the packet throughput varies minutes to minute. As the degree of sharing goes up (from your xDSL to your ISP’s shared bandwidth to the massive sharing of the Internet) the probability is that at *some point* an application that needs sustained uninterrupted bandwidth, such as IP audio streaming, will experience a problem.

Users are intuitively aware of how each of these networks function. A person making a lot of telephone calls might state “gee the telephone network is busy today, I keep getting fast busies” just as nearly every Internet users has observed at some point “the Internet is busy today, it is very slow”.

CODEC REQUIREMENTS FOR RELIABLE OPERATION ON IP

So perhaps a more useful question is “can ISDN be replaced by an IP offering such as xDSL?” The answer is strictly speaking “probably not”. For example, if you need the ability to call the more than 25,000 audio codecs that are currently on ISDN, then xDSL is not going do you much good because you cannot place “calls” from xDSL to ISDN. This is unlike the situation with voice telephony where there are numerous companies that provide gateways between the IP world and the Circuit Switched Telephony world. We do expect that we’ll see some private gateways between these two worlds, but do not expect to see commercial offerings of such a service.

Now, if you alter the question to “will it be possible to use IP over xDSL to replace many of the broadcast applications currently using ISDN,” the answer is much more positive. But, before we can answer this question we need to first look at more of the details of how IP audio operates, and the requirements for reliable operation over the Internet.

PACKETIZATION: DELAY VS. PACKET SIZE

An MPEG encoder produces a stream of data at a constant rate. To an ISDN network this simply looks like a constant serial bit-stream. When this data stream is to be sent over a packet network, the packetizer must accumulate sufficient data to fill each packet before it can be sent. Thus a buffer must be included between the MPEG encoder and the packetizer. This added delay puts a packet-based network such as IP at a fundamental disadvantage to a synchronous network such as ISDN.

The IP specifications allow a wide range of packet sizes. Since each IP packet must include the same “header” information (such as destination address) regardless of size, the packet size determines the actual throughput requirements for a given MPEG payload — the larger the packet size the less bandwidth is required after packetization. On the other hand, the longer the packetizer must wait to fill a packet the more delay is introduced — the smaller the packet size the lower the delay. This trade-off is not present with ISDN.

IP TRANSMISSION

To minimize delay, IP codecs typically use RTP (real-time transport protocol). This IP protocol is intended for delay-sensitive streams; it minimizes delay and therefore does not allow lost data to be re-sent. On a properly managed packet network (e.g. one that includes provisions for QoS) this approach is efficient. In the cases where shared networks are used, provisions must be made to accommodate the inevitable (but generally rare) loss of packets.

PACKET JITTER

The larger and more complex an IP network, the more variability there will be in packet arrival time at the far end. This is because the more complex the network the more possible routes a packet may take. Not only will the time between packets vary, but it is not at all uncommon for packets to arrive out of order. This variation in arrival time is referred to as “packet jitter”. RTP supports packet numbering, which allows the receiver to put out of order packets (due to “late” packets) back in order before sending the data to the MPEG decoder. This requires that some packets be held aside before being read into the decoder, so that a late packet can be put back in its proper place in time to use its data. This packet “waiting room” is referred to as a “jitter buffer”. The size of the jitter buffer represents another delay trade-off — if it is set to be very small to minimize delay, then late packets may be lost. Basic IP codecs allow users to adjust the jitter buffer, but since network conditions vary, it is difficult

to find the optimal setting, and therefore conservative (e.g. longer delay) settings must be used to avoid audio drop-outs.

ERROR RECOVERY

ISDN networks have low enough error rates that it is rare that any special technique is needed to deal with lost data (and since the error rates are guaranteed the solution is to fix the problem source, not to attempt to ameliorate the symptoms). This is not the case with IP networks that do not have QoS mechanisms in place. The simplest approach to dealing with lost or corrupted data is to add redundancy to the system. So-called “forward error correction” (FEC) systems take this approach. The problem is that redundancy means a higher bandwidth is required for transmission. And of course the higher the bandwidth used, the greater the odds are that some of it will be lost. Research by the Internet Streaming Media Alliance indicates that for this reason this approach is not particularly useful, though many codecs include provisions for it.

ERROR CONCEALMENT

A much better approach is to make the decoder smarter, so that it can recover from a lost packet or two. This approach relies on psycho-acoustic principles and is called error concealment. Error concealment works remarkably well and should be, as a minimum, included in codecs intended for use on IP networks without QoS such as xDSL. The Zephyr Xstream includes error concealment in our AAC decoder, and of course the Zephyr/IP includes this as well.

ADAPTIVE BUFFERING

As mentioned earlier, the larger and more complex an IP network, the greater the packet jitter. Therefore, over the Internet substantial packet jitter can be expected, and as discussed above, jitter buffers are essential. Advanced IP codecs have provisions to automatically and dynamically adjust buffer size to minimize delay while avoiding dropouts. This approach is complimentary with error concealment since an occasional buffer over or underflow will not be audible when error concealment is present.

ADVANCED CODECS

Traditional codecs operate at a fixed rate. The more advanced IP codecs, as well as some of the more sophisticated Voice-over-IP systems, allow the codec encoder bit rate to be varied dynamically, while the decoder is designed to be smart enough to follow these rate changes. Actually, the system requires a feedback loop to be effective — At the decode side the jitter

buffer is monitored. If network conditions deteriorate, the adaptive buffer increases the buffer size, and if the condition persists the encoder is notified to ratchet down the bit rate so that the adaptive buffer can then adjust to bring the delay back down again.

Of course using the latest codecs permits the lowest bit rate, and therefore increases the odds of success substantially when compared to older codecs such as G.722 and MPEG Layers 2 and 3.

CONCLUSIONS

While it is safe to say that the services collectively called xDSL cannot *replace* ISDN, it is also true that many broadcast applications formerly handled by dedicated synchronous lines (such as T1 or E1) and ISDN will be using audio-over IP in the near future. In some cases (such as an STL, for example) the best approach will be to ensure that IP links being shared with other applications have suitable QoS mechanisms in place. In the typical remote scenario, ISDN will continue to be useful since it is either working or broken, making troubleshooting relatively easy.

When using ad hoc IP connections that traverse the Internet, such as xDSL or a public WiFi hot spot, network performance will be less predictable and therefore I recommend that the best technology be used to “make the best of a poor situation”. In the real world, when using an advanced IP codec such as the Zephyr/IP, with the features discussed above, users will find that IP audio is indeed a useful alternative to ISDN, and typically better than POTS codecs.