

The ISP Column

An occasional column on things Internet

October 2005

Geoff Huston

A Little Carrier Nostalgia

We've come a long way in just a decade or so in terms of the way in which packets are passed around the network. These days Ethernet framing is getting close to ubiquitous, both in the long haul carriage markets and in various forms of last mile access technologies. These days carriage technologies are being defined in terms of ease of use for data and attempting to avoid gratuitous packet re-framing and various forms of packet shredding and the associated complexity of adaptation layers, rather than attempting to create a byzantine carriage system that could carry various diverse payloads using a lowest common denominator approach. As a simple example of the changes here, while many ADSL systems use a carriage technology of ATM as their base data payload format, the wireless access 802.11 systems present an Ethernet packet frame format to the IP protocol layer. Life is certainly far simpler for many network devices, not to mention for many network operators.

But it wasn't always this way of course.

For many years, indeed right up until the end of the twentieth century the relatively low value, and relatively small volume of data traffic on carrier networks implied that there was no dedicated data carriage technology, and much of the Internet was originally built on top of the carriers' telephone carriage plant. While there is still some operating examples left to look at, maybe its time to take a nostalgic look at the carrier hierarchy, on the understanding that without it its not clear how the data industry would ever have managed to grow large enough to start to define its own carriage services. So lets have a nostalgic look at what was the carrier service portfolio.

The Carrier Hierarchy

All of the world's communications system for telephony was constructed on the basic premise that the human spoken word uses a limited range of frequencies and has limited dynamic range. The world's telephone network was attuned to being able to reproduce the spoken word with acceptable clarity, and to do so used a system that could carry analogue signals of between 350Hz and 3,400Hz

To convert this analogue signal to a digital signal the technique used is Pulse Code Modulation (PCM). The first step is to transform a continuous analogue signal into a sequence of pulses. To undertake this transformation without information loss requires the application of sampling theory. Nyquist's theorem asserts that to undertake a discrete sampling of an analogue signal without information loss, the sampling rate must be no less that twice the highest frequency contained within the analogue signal. For voice signals, the Nyquist sampling rate should be no less than 6,800 samples per second, assuming that the highest frequency is 3,400Hz. The voice carrier industry standardized on a sampling rate of 8,000 samples per second to allow for intelligibility of voice reproduction, effectively choosing a network clock base of 125 μ s.

The number of bits used to encode the amplitude of each sample is the next analogue to digital conversion issue. The smaller the number of bits used for this encoding, the greater the distortion of the signal (quantization distortion), while a higher number of bits per sample imposes a greater load on the underlying digital carriage system. The industry standardized on an amplitude encoding which used 256 discrete levels, called quantization levels. Conveniently, this maps to an 8 bit encoding value. Voice signals have a limited dynamic range, and to reduce the distortion of voice the encoding mechanism from the analogue sample to an 8 bit encoded value was non-linear. The greatest concentration of encoding levels is within the dynamic range of the human voice. After worldwide adoption of a single sample time base, and a single quantization factor of these samples, it was probably too much to expect worldwide standardization of the encoding algorithm. The choice of a quantization algorithm is not uniform worldwide, and μ -law encoding systems are in use in the United States and Japan while A-law systems are common elsewhere. In either case, a voice call is mapped to a 64Kbps datastream, and this 64Kbps digital stream is the base building block of the voice carrier hierarchy.

To take these individual 64kbps streams and allow them to be carried within larger bearers across the network requires the use of a multiplexing technology. The most common multiplexing technology used in today's carrier networks is Time Division Multiplexing (TDM).

Time Division Multiplexing

Multiplexing takes a number of discrete inputs and multiplexes these inputs into a single higher capacity datastream. This stream can be transmitted over a higher capacity link and then demultiplexed back into the original discrete channels. Where the inputs are constant rate signals, the multiplexing creates a single signal whose rate is no less than the sum of the component rates.

Time division multiplexing (TDM) effectively takes a frame of data from each input in turn and transmits these frames across a common link. As the speed of the common link is the sum of the component links, no data is lost in this transmission model. The way in which this is achieved within the multiplexor (MUX) is by using a frame buffer for each input line. Incoming bits are loaded into the frame buffer. At every scan interval, the common line scheduler empties the frame buffer and loads it into the link driver. The common line scheduler scans each frame buffer in turn, where the complete scan of all input lines takes one scan interval. The output operation is similar, in that each frame is assembled in the common input driver and then placed in the link driver buffer for the output line.

The frame buffer can be of any length, although the longer the frame buffer, the greater the latency of the multiplexing operation. The scan interval of the carrier hierarchy is typically based at the sample interval of PCM encoding of voice circuits, which is a scan time interval of 125 μ s, or 8,000 scan intervals per second.

TDM is inherently very simple in its operation, but such simplicity is not without cost. TDM allocates a fixed amount of capacity to each input channel, whether the channel is active in transmitting a data element or not. A simple TDM MUX cannot operate on an adaptive basis where the common channel capacity is less than the sum of the input capacities and each channel is allocated resources on the basis of data activity. Nor can a TDM allow uncontrolled clock slippage of any of the input channels. The assumption in a simple TDM model is that each input line operates within a synchronized clocked mode, so that packing the common multiplexed datastream with control information is not required.

TDM systems commonly use a very simple framing protocol, in which each scan is terminated with a single frame control bit. This bit alternates on each scan, creating a recognizable bit pattern '01010...', which can be used to synchronize the demultiplexing unit to the input unit, so that the demultiplexing unit can recognize the alignment of each scan frame within the bit stream.

The Plesiochronous Digital Hierarchy

Of course the assumption that all clocks are highly stable and conveniently synchronized is one that does not always apply in practice. This TDM system is complicated by allowing the various clock sources to drift within certain parameters. There is no single 125 μ s master clock driving the entire voice network, and the system was constructed to allow some level of longer term clock drift between the various component PCM clocks. This is accommodated within the TDM system by providing overflow space within each scan frame, so that a clock operating at a slightly faster rate can insert additional bits into the overflow space, with an associated timeslot label to indicate which input line has generated the additional bits. This allows the TDM MUX to correct the overrunning source at periodic intervals.

In a TDM digital switched hierarchy, the 64Kbps PCM encoded datastreams are termed DS-0 circuits. From this base point, the carrier hierarchy is constructed (see Table 1).

- In North America and Japan, the first level of carrier multiplexing is to take 24 of these DS-0 circuits to create a DS-1 circuit, operating with the inclusion of an 8Kbps framing signal stream (remember those alternating 1's and 0's at the end of each scan interval?) at a data rate of 1.544Mbps (commonly termed a T-1 circuit). Elsewhere, the first level of the hierarchy uses 30 DS-0 circuits to create the CEPT-1 circuit, clocked at 2.048Mbps, here allowing for 128Kbps framing and control signal streams in addition to the 30 PCM streams (commonly termed a E-1 circuit).
- The next level of the hierarchy, DS-2, is built from four multiplexed DS1 streams. This point in the hierarchy is not used in most carrier systems.
- The next level is that of a DS-3 bearer. In the North American system, this takes 28 DS-1 groups and maps them into a 44.736Mbps bearer, termed a T-3 circuit. In the CCITT bearer system, a CEPT-3 is a mapping of 16 E1 circuits, which is a 34.368 Mbps bearer, termed an E-3 circuit.

Hierarchical Level	North American	European	Japanese	International
	DS-x	CEPT-x		
	Kbps	Kbps	Kbps	Kbps
0	64	64	64	64
1	1544	2048	1544	2048
2	6312	8448	6312	6312
3	44736	34368	32064	44736
4	139264	139624	97728	139264

Table 1 The Plesiochronous Digital Hierarchy

The additional space allocated in these higher order points of the carrier hierarchy are used to allow for slightly different clocking speeds of the individual streams, so that space is allocated for additional bits in each frame scan to allow for clock alignment at the start of every frame. This is why this hierarchy is termed a plesiochronous digital hierarchy (PDH), (where plesiochronous means "almost synchronous") because the different DS-1 bearers are not tightly time synchronized with each other. If one stream is running slightly faster than the other multiplexed streams, it can overrun into the defined overflow bits within each scan frame. This technique spreads each component stream over the entire aggregate frame, and the task of extracting or inserting a single component stream into a multiplexed group without disturbing the remainder of the multiplexed streams is not commonly undertaken. Every voice switch within the PDH carrier hierarchy must demultiplex the trunk signal down to the level of DS-0 streams before a single channel can be peeled out of the aggregate carrier bearer.

The PDH carrier system allows for point-to-point DS-0 data links operating at speeds of 56Kbps in North America, or at 64K elsewhere, to be used for private point-to-point circuits. The reason for the lower bit speed in the North American system was the use of 1 bit in 7 for clock integrity. Typically, carriers also allow for a number of DS-0 circuits to be provided in a composite bundle, normally by provisioning a logical DS-1 bearer between the two circuit end-points and then marking out a number of timeslots that are available for use within the logical data circuit. The PDH also allows for the provisioning of DS-1 point-to-point circuits, at 1.544Mbps or E1 at 2.048Mbps, depending on the locally used hierarchy, and where available, from the carriage operator, DS-3 circuits, which operate at 44.736Mbps, or E3 circuits at 34.368Mbps. No composite DS-1 circuits are available from the carriage operators, due to the operation of PDH framing, although the use of inverse multiplexors can take a number of DS-1 circuits and create composite data rates between DS-1 and DS-3 speeds

The Synchronous Digital Hierarchy

These issues of making the hierarchy fit a set of imprecisely synchronized clocks can be eliminated if all component datastreams are synchronized. In a precisely synchronized environment, every component stream will occupy a fixed area of the multiplexed data frame, allowing streams to be added or removed without a complete demultiplexing operation as a prerequisite. This is used in the Synchronous Optical Network (SONET) to define a carrier hierarchy. SONET is the North American standard, and the Synchronous Digital Hierarchy (SDH) standard is used in other parts of the world.

Hierarchical Level	North American Designation	International Designation	Presented Rate Mbps	Data
1	OC-1		51.84	
2	OC-3	STM-1	155.52	
3	OC-9	STM-3	466.56	
4	OC-12	STM-4	622.08	
5	OC-18	STM-6	933.12	
6	OC-24	STM-8	1244.16	
7	OC-36	STM-12	1866.24	
8	OC-48	STM-16	2488.32	
9	OC-192	STM-64	9953.28	

Table 2 The Synchronous Digital Hierarchy

The defined speeds within the hierarchy are indicated in Table 2. As these speeds are now synchronous, there is no need to add overflow bits into the multiplexed frames, so that the data rates are now exact multiples of the data rates of the lower speed trunk systems. The frame format remains locked to the constant 8,000 frames per second base, or 125µs per frame. The base frame is an OC-1 frame, which is 810 bytes. Each byte within an OC-1 frame corresponds to one 64K DS-0 stream. Higher order rates are constructed by interleaving OC-1 frames using byte-by-byte interleaving. These OC-1 frames could be treated as a single aggregate bit stream through the use of concatenation. A 'c' following the designation denoted a concatenated data stream that presented to the customer at the aggregate data rate.

Within the carrier industry, four points on this hierarchy were commonly used as carriage options, namely OC-3c and OC-12c as a starting position, and later OC-48c circuits became more prevalent.

Where the SDH is still used in the data industry there is now widespread use of OC-192c circuits, as they map easily into 10G Ethernet systems.

Digital Modems

Much of the early access systems for the Internet were based on modem calls. The low cost of local modem calls suited the funding nature of such networks, and, as long as total volumes were low, the modem-based systems could keep pace with the total amounts of traffic to be transported.

The task of the modem is to map a digital stream into an analogue signal that can be passed over a voice line. The voice band telephone channel is a bandpass channel, operating in a band from 300 Hz to 3,400 Hz.. Early modems used discrete tones that fell within this frequency band for communicating data, with a data rate of 300bps. Quadrature amplitude modulation (QAM) was a significant improvement for modems, modulating a carrier sine wave signal in both phase and amplitude, allowing a number of discrete signals, offering information densities of multiple bits per hertz. Depending on the line quality, the number of symbols per second can be increased to fill the available voice channel spectrum. The V.34 rate, for example, used a carrier of 1959 Hz, and a symbol rate of 3429 symbols per second, giving a bandwidth from 244 Hz to 3674 Hz. This bandwidth is effectively the theoretical maximum spectrum space available on a voice system..

However, digital telephony, such as ISDN, allows the use of 56Kbps modems to operate, using pulse amplitude modulation (PAM) from the network to the client modem, while using QAM on the return path. This allows a 56Kbps data rate from the network to the modem and a 33.6Kbps rate in the reverse direction. The standard for this form of modem operation is V.90, adopted as a communications standard in February, 1998.

For some millions of Internet access subscribers, the V.90 standard played a central role in connecting to the Internet for some years. This was in spite of early experiences of variable connection speeds, variable connection success rates, and long coding latencies within the modem.

Digital Circuits

A number of point-to-point digital carriage services were introduced to the market following the introduction of digital infrastructure into the telephone network. The earliest of these services is the Digital Data Service (or DDS). DDS services are point-to-point 56Kbps circuits (in some networks it is a 48Kbps circuit). DDS services are provisioned as 64K digital circuits within the digital transmission network. On the copper loop, the loop is groomed to remove loading coils, taps, and other sources of unwanted noise and distortion. The encoding used on DDS systems is alternate mark inversion (AMI), in which each 1 bit is represented alternatively by a positive and a negative voltage pulse. The properties of this encoding include a net zero DC voltage, and relatively easy detection of impulse noise. However, no inherent clocking of 0 bits takes place, and the two end points may lose clock synchronization with a long sequence of 0 bits. The clock signal uses 1 bit in every 8 as a 1 bit to ensure the maintenance of clock synchronization. The DDS data rate is therefore 7 bits of every 8 bit frame, or 56Kbps.

A variant of the encoding system addresses the potential loss of clock synchronization on long strings of 0 bits by the use of bipolar 8th zero substitution (B8ZS) line encoding, to offer clear channel capacity at 64Kbps. The encoding system requires slightly greater complexity, in which every sequence of 8 0 bits is encoded with a pair of bipolar pulses on the 8th 0. This preserves net zero DC voltage and preserves sufficient pulse density to maintain synchronous clocking at both ends.

Higher digital speeds were based on a T-1 (or E-1) circuit. Commonly available carrier services used a number of 64K frames from a framed T1 (or E1) and use a Customer Service Unit (CSU) to aggregate these framed segments and present them to the customer as a single-clocked service. Carrier services also typically included the provision of clear channel T-1 (or E-1) circuits, which allow the customer to use CSU equipment to clock at the effective data rate of the circuit.

Yet higher speeds were available using a number of hybrid methods. Inverse multiplexing was used to bond together a number of parallel T-1 (or E-1 circuits) and present an aggregate data rate to the customer termination equipment.

Within typical carrier pricing structures when these circuits were in widespread use, the unit cost of data circuits decreased for each 64K increment from 64K to the T-1 or E-1 point, as the common provisioning model was to actual provision an entire bearer between the two end points and then enable the selected number of DS-0 circuits on the bearer.

The next price point in many carrier's portfolios were at the T-3 or E-3 service level, where lower unit costs were often realizable. Such higher speed services were not universally available, but where provisioned, they were normally of a form of clear channel. Framed T-3 or E-3 services were not a common feature of the typical carrier service portfolio.

These circuits were all based on the PDH carrier hierarchy. They offered an end-to-end synchronized data clock, which preserved bit-level and packet-level clocking between the sender and receiver.

ISDN

These digital circuit services shared a second common attribute, as well as synchronized data clocking. The services were statically configured on an end-to-end basis. If the customer wanted to change the location of either end of the circuit, the carrier had to reconfigure the digital circuit service network to relocate the circuit. The architecture of the Integrated Services Digital Network (ISDN) represented a carrier effort to combine the utility of digital end-to-end circuits with the switching systems used in the PSTN environment, allowing digital calls to be dynamically created and torn down by the customer.

The ISDN system was part of the switched telephone network and uses the same switching services. Telephone and data services could be accessed via ISDN, although, confusingly, often at different tariff rates. In the ISDN architecture the local loop was not an analog signal; the local loop was a collection of 64K data channels, which were effectively an extension of the internal digital channels used to carry voice circuits within the PSTN.

Two access services were defined within the ISDN architecture; a Basic Rate Interface (BRI), and a Primary Rate Interface (PRI). A BRI used three separate channels into the network; a 16Kbps signaling channel, the D channel, and two independent 64Kbps data channels, B channels, that could be used for data or voice. The D channel was used to control the B channel connections, using a set of call control messages (the format of these messages is defined in the ITU-T standard Q.931), to initiate, and terminate, B channel calls. Each B channel call is a clear channel 64Kbps circuit. In this architecture, the customer equipment was more complex than the equipment used to terminate digital end-to-end circuits, as the customer equipment now had to manage the three circuits and use Q.931 to manage calls on the two B channels.

A PRI used a configuration of 23 B channels (in North America, or 30 B channels elsewhere) and a 64Kbps D channel. Again, the D channel was used to control the calls made on the B channels, and the operation was similar to that of the BRI.

The ISDN service network treated each 64K circuit as an independent call, and it was the responsibility of the customer equipment to use inverse multiplexing (or bonding) to group together bundles of B channel calls and create the functional equivalent of a larger capacity aggregate channel.

The use of V.90 high-speed modems for Internet access was predicated on the assumption that the service provider end of the call used a digital interface to the PSTN network. In those markets where digital T-1 services with signaling interfaces were not a customer-accessible service this was often implemented using an ISDN PRI, and the dial-in Network Access Server terminated the incoming access calls using onboard digital signal processors.

SONET and SDH

The SDH represents the use of a tightly synchronized clocking environment to create a digital carriage hierarchy that does not distribute framed data across the frame in order to compensate for clock drift within the various constituent circuits. This allows for the construction of high-speed digital carriage hierarchies that can be easily combined or split. Individual circuit groups can be peeled off the aggregate circuit or readily inserted into a vacant circuit slot.

The SDH Add/Drop Multiplexor (ADM) undertakes the insertion of a data flow into a SDH stream and can perform the removal of a data flow. This does not remove the requirement to set up a channel within an SDH bearer system, but it does allow the use of individual channel manipulation without the need to demultiplex the entire circuit hierarchy. In addition, SDH systems are typically configured in a dual ring structure, using the architecture of a working and a protection ring. The ADM automatically switches data into the protect ring in the event of failure of the working ring.

The synchronized clocking also allows the carrier speed to be increased well beyond the 34Mbps or 45Mbps speeds, which were the typical ceiling of carriage services provided from the PDH hierarchy. The early services emerging from the SDH bearer system were STM-1c at 155Mbps and STM-4c at 622Mbps. With the continuing growth of the Internet, use of STM-16c circuits became more common, and these days where SDH services are used STM-64c, or 10G SDH is not uncommon.

Did the Voice Carriage Network subsidize the Internet?

The telephone industry had a pretty sophisticated operating model by the 1980's. The way they tended to work was to build capacity in bursts, rather than attempt to undertake continual incremental expansion. They had figured out by then that "just in time" capacity provisioning was making life incredibly hard, and the way forward was to attempt to build in infrequent intervals. When the carrier built capacity the intention was to do so at such a scale that they not have to return and rebuild for, ideally, some decades to come. So the model was one of over-provisioning at a relatively dramatic scale.

So when data wanted "circuits" what we received was circuits drawn from this over-provisioning regime that was intended ultimately to support voice traffic. Now since this capacity had been already installed and capitalized some years previously, the marginal cost of provisioning data circuits was in fact marginally small. It had effectively been paid for years before. But what the phone companies were nervous about was the leakage of their core business, and they were all too aware that it was easy to put PABXs at either end of a data circuit and have the data customer perform the toll bypass game. So, quite deliberately the data circuits were tarified at levels that were typically well above actual cost. These data carriage services were often tarified at prices that were based on the voice toll charges. That way if you did your sums, attempts to get away with toll bypass were frustrated by the high data circuit prices.

So in effect the data customers were paying much the same prices as the voice folk. Now they were not paying the full price of dedicated provisioning of capacity, but they were paying their 'fair' share of this infrastructure cost (on the assumption that the prevailing voice toll charges at the time represented a 'fair' price, of course). The voice carriers had no interest in underpricing the data market as they were all too aware that if they did so they would stimulate a cannibalistic toll-bypass reaction from their favourite large customers.

While the data demands were low enough to fit into the carrier's digital hierarchy of the day, things were just fine - 56K, 64K, n x 64K, T-1, E-1, T-3, and E-3, services all fitted relatively comfortably into this model of provisioning data services from the margins of over-supply of voice infrastructure, and tariffing the service at the equivalent of voice prices.

Things broke down in the mid-90's. When data demands started to want individual circuits at the 155Mbps level and higher all of a sudden there was no margin of oversupply in the common switched network to sign over to data. All of a sudden the carrier industry was forced to pass over sets of base bearers to the data folk. And as the data demands continued the carriers found themselves provisioning additional bearer capacity on the strength of data forecasts alone, and the entire structure of providing data services from the margins of oversupply of the voice network vaporized. The carrier's bearer forecasters were unused to the data demands and it appears that they managed to continually under-estimate the Internet's demands. The supply shortfall was at its most acute in the mid to late 90's and the carrier folk were clearly not reacting quickly enough in terms of getting the carrier business to make rapid and large capital investment in expanded infrastructure services. This was the case in both national and international markets where there was the emergence of a mass consumer and corporate market in Internet services with a seemingly rapacious demand for carriage capacity.

Where the traditional carrier industry had stalled, other industry players leapt in to fill the gap, and the careful control of the data infrastructure market fragmented. The revolution of Dense Wave Division Multiplexing that created massive capacity on fibre trunk circuits came in to play, and infrastructure build out in a highly competitive market ensured in an effort by other players to position themselves into this data carriage market. And from there dedicated data provisioning in the carrier space became the normal operational practice. From that point the carrier hierarchy started to retreat back to its basic role of service to the voice industry, while the data world has started on along path of dedicated data framing services.

Or at least that's the way I see it!

Disclaimer

The views expressed are the author's and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

About the Author

Geoff Huston B.Sc., M.Sc., has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and has been active in the Internet Engineering Task Force for many years.

www.potaroo.net