

George Scheets and Mark Allen. "Switched Network Carrying Capacities"
Handbook of Emerging Communications Technologies: The Next Decade.
Ed. Saba Zamir
Boca Raton: CRC Press LLC, 2000

7 Switched Network Carrying Capacities

George Scheets and Mark Allen

CONTENTS

- 7.1 Introduction
- 7.2 Measuring A Trunk's Ability to Carry Traffic
- 7.3 Circuit-Switched TDM Trunks
- 7.4 Hybrid Trunking
- 7.5 Packet-Switched Statistical Multiplexed Trunks
- 7.6 ATM Statistical Multiplexed Trunks
- 7.7 Head-to-Head Comparison
- 7.8 Conclusions
- Acknowledgments
- References
- Web Assistance

7.1 INTRODUCTION

Effectively delivering bandwidth today is more important than ever before. Networks that deliver bandwidth to users and their applications in the most efficient manner with the appropriate Quality of Service (QoS) will likely prevail over less effective bandwidth delivery schemes. Until recently, voice made up the bulk of a network's traffic. Today, data dominates. Tomorrow's traffic will likely be multimedia with varying levels of tolerance for loss, delay, and variation in bandwidth. Systems of the future must allow for this mixture. Clearly, the network that can most effectively satisfy the user's requirements will ultimately provide the service at the lowest price, a key factor being the ability not to over-engineer the system when a less expensive solution will satisfy the customer. For example, the success of the frame relay service from public carriers is due to the user's willingness to share a bandwidth pool with others and to tolerate occasional loss and delay in exchange for significantly lower data transport costs compared to DS1 (1.544 Mbps) and fractional DS1 private line service.

The increasing emphasis on mixed traffic requiring QoS guarantees and ever-expanding amounts of bandwidth means that protocols that rely on inefficient sharing of the physical media are less likely to be acceptable for end-to-end delivery. One consequence of these demands is that today's telecommunication networks are

increasingly switched in nature. WANs have relied on switches to consolidate and move traffic to their end destinations for years, but only recently have classical shared Ethernet and Token Ring LANs given way to switched Ethernet and Token Ring, and only recently have shared LAN backbones, such as FDDI, given way to switched Fast Ethernet, Gigabit Ethernet, and ATM systems.

This chapter examines the ability of switched networks to carry end users' application traffic, given a multiplexing choice and an offered load mix ranging from 100% voice or video to 100% bursty data traffic. The focus is on the backbone where fiber tends to dominate, but the concepts discussed are applicable in *any* environment where the effective use of bandwidth is required, including down to the desktop.

Today a network provider essentially has four trunking choices regarding what combination of switching and multiplexing schemes to use for delivering bandwidth to the customer: circuit-switched TDM trunking, hybrid TDM trunking, packet-switched StatMux trunking, and ATM StatMux trunking.

Circuit-Switched TDM Trunking — In this configuration, the network delivers bandwidth to the customer at a constant rate with no buffering or bursting capability and with a minimal fixed delay. [Figure 7.1](#) shows a simplified block diagram of an edge switch in this type of network. The input arrows represent *multiple* time-sensitive and data traffic sources. Traffic from these sources is multiplexed onto a trunk connection for transport to the next switch. Input voice or video, which will frequently be referred to in this chapter as *time sensitive traffic* (TST), could be either fixed, or variable rate in nature. Input from any data source is assumed to be bursty in nature. [Figure 7.2](#) shows what the traffic from a typical data source might look like. The input is either active, in which case data traffic is entering the switch from this particular source at the line speed, or inactive, in which case no data is entering the switch from this source. To successfully move all of the offered traffic, a TDM circuit switch, which does not include buffering and is unable to handle traffic bursts, must assign dedicated trunk bandwidth to each source based on the peak (line) rates of each input. A 64 Kbps fixed rate voice conversation must receive 64 Kbps of trunk bandwidth. Data traffic offered at an average rate of 154 Kbps, and a peak rate of 1.54 Mbps, must be assigned 1.54 Mbps of trunk bandwidth.

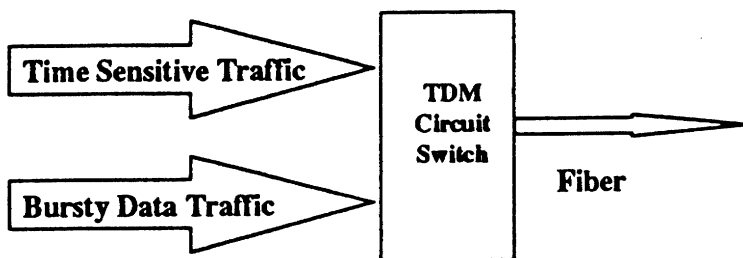


Figure 7.1 Circuit-Switched TDM Trunking

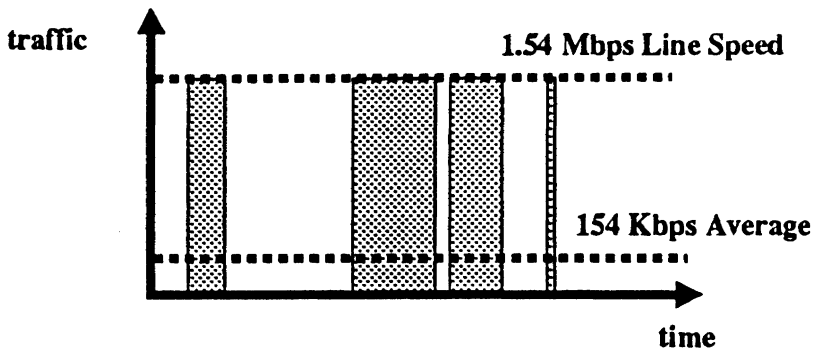


Figure 7.2 Bursty Data Traffic with a 10 to 1 Peak to Average Ratio

Hybrid TDM Trunking — [Figure 7.3](#) shows a simplified block diagram of a hybrid network edge switching node. It is similar to the circuit-switched TDM configuration of [Figure 7.1](#) except that *two separate networks are maintained*. All bursty data traffic is ideally groomed onto a packet-switched, statistically multiplexed (StatMux) network such as frame relay or the Internet, providing better utilization of network backbone resources. Fixed-rate time sensitive traffic remains on the circuit-switched TDM network. Variable rate voice and video could go either way, depending upon whether timely delivery or bandwidth efficiency is more important. Fiber bandwidth is assigned to the resulting packet-switched and circuit-switched bit streams on a dedicated circuit-switched TDM basis based on the peak rates of the resulting traffic.

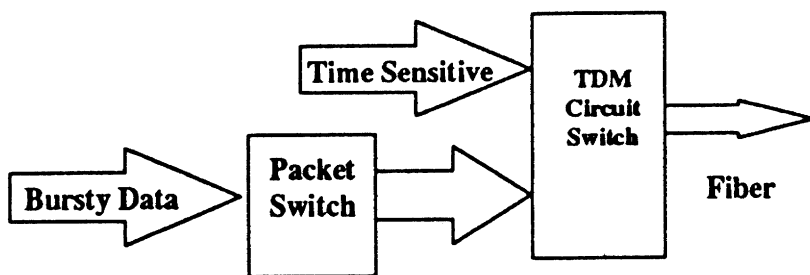


Figure 7.3 Hybrid TDM Trunking

Packet-Switched StatMux Trunking — All traffic, including that originating from fixed or variable rate time sensitive sources, is packetized and StatMuxed onto high speed trunks prior to insertion into the fiber, as illustrated in [Figure 7.4](#). Packet Switching and StatMux are the foundations upon which the Internet Protocol (IP) and the Internet, as well as frame relay, are based. The Internet model is claimed

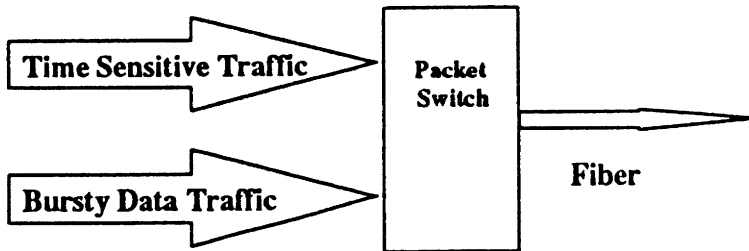


Figure 7.4 Packet Switch StatMux Trunking

by many to be the networking choice for the 21st century from which all telecommunications services will be delivered.

ATM StatMux Trunking — This is the technique perceived by others to be the network model for the 21st century. Traffic from all sources is segmented into fixed-size cells and StatMuxed onto high speed trunks prior to transmission over the fiber. Fixed-rate traffic is assigned constant bit rate (CBR) virtual circuits which are capable of providing TDM-like QoS. Bursty traffic is assigned variable bit rate (VBR), available bit rate (ABR), or unspecified bit rate (UBR) virtual circuits and is StatMuxed onto trunk bandwidth not reserved for CBR traffic. [Figure 7.5](#) shows a simplified block diagram of this configuration.

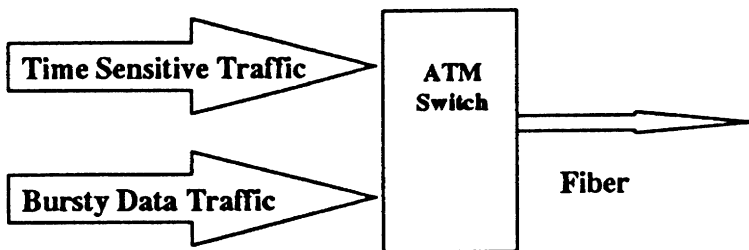


Figure 7.5 ATM StatMux Trunking

7.2 MEASURING A TRUNK'S ABILITY TO CARRY TRAFFIC

What parameter best measures a network trunk's carrying capacity? The trunk *Efficiency*, which is often defined as

$$\text{Efficiency} = \frac{\text{bits per second carried by trunk under heavy load conditions}}{\text{trunk line speed}} \quad (7.1)$$

is a parameter frequently touted. Figure 7.6 shows what a plot of trunk efficiency might look like. If the offered load consists of 100% time sensitive traffic, all of the previously mentioned configurations are able to completely load the trunk output lines, although it should be noted that the circuit switch TDM configuration can do so only if the voice and video are fixed rate. If any bursty traffic is offered, packet and ATM networks are more efficient as they are able to completely load the output trunk under heavy load conditions, but a circuit switch TDM backbone will have gaps in the traffic, as noted in Figure 7.2, and hence will have an efficiency less than 100%.

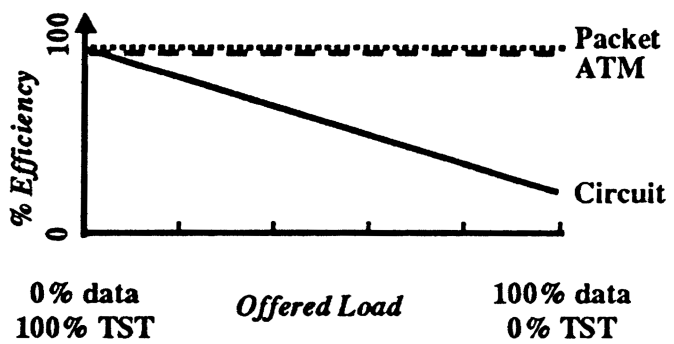


Figure 7.6 Switched Network Efficiency

However, the trunk efficiency does not tell the whole story. It does not account for the fact that a real-world StatMuxed trunk line carrying a 100% load is unusable as it either would have high queuing delays or would be dropping excessive amounts of offered traffic due to buffer overflows. As defined above, the efficiency also does not account for packet or cell overhead, although it should be noted that some definitions of efficiency do account for this overhead.

A more accurate measure would be the *carrying capacity* or *utilization*, which is defined here as

$$\text{Carrying Capacity} = \frac{\text{carriable end user application traffic in bits/second}}{\text{trunk line speed}} \tag{7.2}$$

The carrying capacity accounts for packet and cell overhead, and it accounts for the inability of StatMux switches to fully load output lines and have a usable system. Figure 7.7 shows what a plot of trunk utilization might be expected to look like. Note the differences between the packet switch and ATM utilization, and the packet switch and ATM efficiency.

The following sections provide details as to how the carrying capacity for each of the four different trunking options can be computed. They examine the issues that affect the amount of overhead consumed and how fully a trunk circuit can be loaded as the traffic mix changes between TST and data traffic. The overhead and

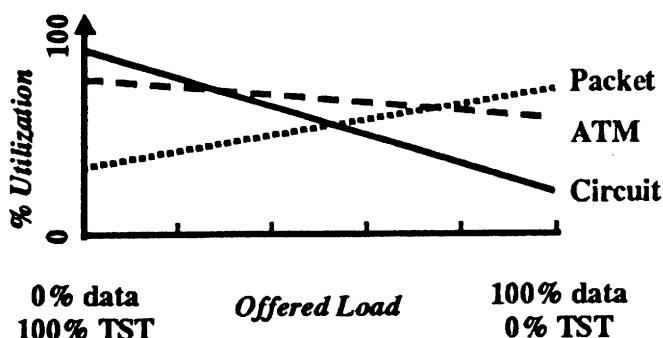


Figure 7.7 Switched Network Carrying Capacity

the StatMux queuing delays impose some severe penalties on a packet switch network's ability to carry time sensitive traffic, lowering the carrying capacity. ATM, which was originally designed to carry mixed traffic, not surprisingly shows high utilization when the offered traffic load consists of a combination of time sensitive and bursty data sources. ATM's ability to give CBR traffic TDM-like QoS gives it a high utilization when the offered load is all fixed-rate TST, and its ability to StatMux bursty traffic gives it high utilization when the offered load is all data.

The following discussion and examples focus somewhat on WANs, but the results can easily be extended to the MAN or LAN by appropriately adjusting the overhead and line speeds.

7.3 CIRCUIT-SWITCHED TDM TRUNKS

Traffic sources, be they fixed-rate voice or video, variable rate voice or video, or bursty data traffic, are all assigned trunk capacity based on the peak rates of each input circuit in a circuit switch TDM backbone network (see again Figure 7.1). The overall carrying capacity can be calculated based on knowledge of the *average* peak-to-average ratios of injected data traffic, the *average* peak-to-average ratios of the injected time sensitive traffic, traffic overhead, and knowledge of the ratio of data to TST being moved over the trunk, via the equation

$$\text{CapCSTDM} = \frac{(\% \text{ traffic to overhead})(\% \text{ usable line speed})}{(\text{peak-to-average ratio})} \quad (7.3)$$

An example of the calculations required is shown in Figure 7.8, which itemizes sources of bandwidth loss when the offered load is 100% bursty data traffic being carried over a SONET-based fiber system. On a typical 810 byte SONET frame, 36 bytes are set aside for operations, administration, and maintenance (OA&M) overhead purposes. Assuming the average packet size of data traffic is 300 bytes, as has recently been measured on the MCI Internet backbone,¹ data traffic originating

from routers would require 6 bytes of Level 2 overhead for High Level Data Link Control (HDLC), 20 bytes of Level 3 overhead for the Internet Protocol version 4 (IPv4), and 20 bytes of Level 4 & 5 overhead for Transmission Control Protocol. Hence, 46 out of 300 bytes (15%) are lost for overhead for each packet, on average. Assuming that a weighted average of all input circuits carrying packet traffic indicated that, on average, 83% of the time the input packet circuits have idle bandwidth, and 17% of the time traffic is actually moving, then a 6-1 peak-to-average ratio is indicated. The overall result would be a trunk utilization of

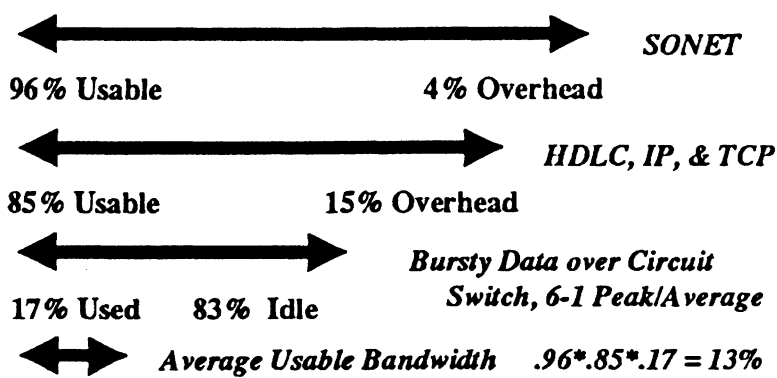


Figure 7.8 Usable Bandwidth: 100% Data over Circuit Switch TDM SONET

$$\text{CapCSTDM} = \frac{(254/300)(774/810)}{6} = 0.1348$$

in this situation. In other words, if the offered load to the switch is 100% bursty data being injected at an average rate of 100 million bits *of end user application traffic* each second with a 6-1 peak-to-average ratio, 100 Mbps/.1348 = 742 Mbps of trunk bandwidth would be required to carry this load. This is not a very effective way to haul data!

At the other extreme, if powerful add-drop multiplexers are available to multiplex individual 64 Kbps fixed-rate voice conversations onto SONET, the primary overhead would be the SONET OA&M traffic, allowing a carrying capacity near 96% to be achieved for TST.

Figure 7.9 shows several plots of circuit-switched TDM utilization as the switch offered load varies from 100% time sensitive to 100% bursty data traffic, for different data peak-to-average ratios. The TST is fixed rate for these graphs, as that is what this type of network most effectively transports.

7.4 HYBRID TRUNKING

In this configuration, the goal is to operate two distinct networks: a TDM-based network for transporting TST and a packet-based network for carrying bursty data

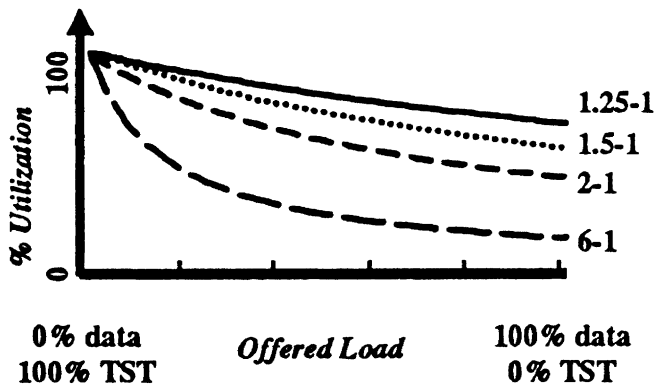


Figure 7.9 Circuit Switch TDM Trunk Utilization for various data peak-to-average ratios

traffic that lends itself to traffic shaping and StatMux. The key difference between this technique and the previous is that ideally *all* bursty data traffic is aggregated onto a packet-switched network (see Figure 7.3). StatMuxing many high peak-to-average ratio circuits together will generate fewer, more heavily utilized packet-switched output trunks, with lower peak-to-average ratios. Backbone capacity is again assigned on the basis of the peak traffic rates of the resulting circuits. As before, the overall carrying capacity can be calculated based on knowledge of the *average* peak-to-average ratios of injected data traffic, the *average* peak-to-average ratios of the injected time sensitive traffic (which ought to be 1-1 if *all* bursty traffic is shipped to the packet switch), traffic overhead, and knowledge of the ratio of data to TST being moved over the trunk.

Figure 7.9 may also be used to estimate the utilization for a hybrid network, as a key function of the hybrid system is to consolidate and shape the packet traffic, thereby reducing the peak-to-average ratio of bursty traffic injected onto the fiber. The consolidated traffic still utilizes dedicated circuit-switched TDM trunk connectivity to adjacent switches, so using the peak-to-average ratios as in Section 7.3 is appropriate for this discussion. It should be noted, however, that the techniques discussed for calculating the ATM carrying capacity in Section 7.6 could be modified to calculate the carrying capacity for hybrid networks, yielding slightly more accurate results.

As an example, if a circuit-switched TDM system with a mixture of fixed-rate voice and bursty data traffic with an *average* input peak-to-average ratio of 6-1 is replaced with a hybrid system capable of consolidating the data traffic onto a smaller number of high speed channels with an 80% load (a peak-to-average ratio of 1.25 to 1), the lowest line of Figure 7.9 would apply to the circuit-switched TDM system and the highest plotted line would apply to the hybrid system. A network that does not fully off-load all the data traffic onto the hybrid network packet switch would lie somewhere between these two extremes.

Examine this graph for an offered load mix of 70% data and 30% voice. The circuit switch system has a utilization of 18% and the hybrid system has a utilization

of 72%. This means that for this example, a circuit-switched TDM backbone would require $.72/.18 = 4$ times the trunk bandwidth and higher speed switches, than a hybrid system hauling the *same* offered load. Depending on the exact equipment costs associated with each network, the hybrid system is likely to offer considerable installation cost savings. The key problem faced here would be properly segregating the traffic so that the highest possible utilization is actually achieved.

Many of the established public carriers originally deployed circuit switch TDM networks in the seventies and eighties, as that was the most economical choice for the voice-dominated systems of the time. Increases in computing power accompanied by simultaneous decreases in the cost of that power resulted in a rise in data traffic and the realization that circuit-switched TDM backbones were not a good choice in an increasingly data intensive environment. Eventually carriers began deploying hybrid systems and made a concerted effort to move as much data traffic as possible onto packet networks, such as frame relay, in order to better utilize their trunk bandwidth and offer lower cost connectivity to their customers. Today, the older carriers commonly deploy some sort of hybrid network to satisfy the continually growing demand for voice and data transport, with varying degrees of success in moving bursty traffic onto the packet side of the house.

7.5 PACKET-SWITCHED STATISTICAL MULTIPLEXED TRUNKS

As shown in [Figure 7.4](#), in this technique traffic from *all* sources is packetized and StatMuxed onto trunks. Carrying capacity can be calculated based on knowledge of the average packet size of the injected data traffic, average packet size of the injected time sensitive traffic, tolerable delays through a typical network switch, ability of the network to prioritize traffic, knowledge of queuing theory and the recent discoveries of self-similarity in network traffic, and some knowledge of the processing limits associated with each switch or router.

In a manner analogous to what is shown in Section 7.3 and [Figure 7.8](#), the carrying capacity of a packet-switched StatMux network can be calculated via

$$\text{CapPSSM} = \frac{(\text{Average application traffic per package}) \times (\% \text{ Usable Line BW}) \times (\text{Trunk Load})}{(\text{Average Packet Size})} \quad (7.4)$$

Everything in this equation is relatively straightforward except for the trunk loading parameter, which is the inverse of the peak-to-average ratio. Determining the tolerable trunk loading requires a knowledge of queuing theory, a field which is currently somewhat unsettled due to discoveries in the last few years that data traffic has self-similar characteristics, meaning that many of the ‘old reliable’ (and inaccurate) queuing results have gone out the window. Some of the key results are briefly summarized here. The interested reader is referred to Stallings² for a very readable overview.

Queuing theory predicts that if the size of input packets is exponentially distributed and independent of the size of previous packets, and if the time between packet arrivals is also exponentially distributed and independent of the previous inter-arrival times, then the average queuing length in a switch is

$$\text{Average Queue Length (in packets)} = \frac{\text{Trunk Load}}{1 - \text{Trunk Load}} \quad (7.5)$$

Experience has shown that these assumptions are not quite true for real-world traffic, with the result that this equation tends to predict overly optimistic small queue sizes. More recent work indicates that under certain circumstances, the following equation provides a more accurate estimate of the average queue length

$$\text{Average Queue Length (in packets)} = \frac{(\text{Trunk Load})^{0.5/(1-H)}}{(1 - \text{Trunk Load})^{H/(1-H)}} \quad (7.6)$$

where H is the Hurst parameter, a value which lies between .5 and 1.0. A Hurst parameter of .5 implies that no self-similarity exists, and Equation 7.6 then simplifies to Equation 7.5. A Hurst parameter value of 1.0 implies that the traffic is completely self-similar, which essentially means that a traffic trace viewed on any time scale (any zoom factor) would look somewhat similar. Figure 7.10 shows a plot of Equation 1.6, for Hurst parameter values of .5 and .75. The key point to note here is that self-similar traffic (such as with $H=.75$), which has burstiness that is more ‘clumped’ than the ‘smooth’ burstiness associated with the exponentially distributed model ($H=.5$), has queues that tend to build more rapidly under smaller loads. This translates directly into higher queuing delays at a switch for packets that are not dropped, as the

$$\text{Average Queue Delay (in seconds)} = \frac{(\text{Average Queue Length}) \times (\text{Average Packet Length})}{\text{Trunk Line Speed}} \quad (7.7)$$

While the jury is not yet completely in, initial studies indicate that the Hurst parameter for typical packet and cell traffic is probably somewhere between .7 and .9.²⁻³

A StatMux network switch can be considered to be operating in one of two modes:

- (1) *low load*, where delay and not loss is a problem, or
- (2) *heavy load*, where loss and not delay is a problem.

The Hurst parameter of the offered traffic will impact both modes. Using Equations 7.6 and 7.7 the Hurst parameter can be used to estimate the average queuing delay for the low load instance. Of equal importance is the heavy load case. Here the Hurst parameter will impact the probability that a buffer overflows.

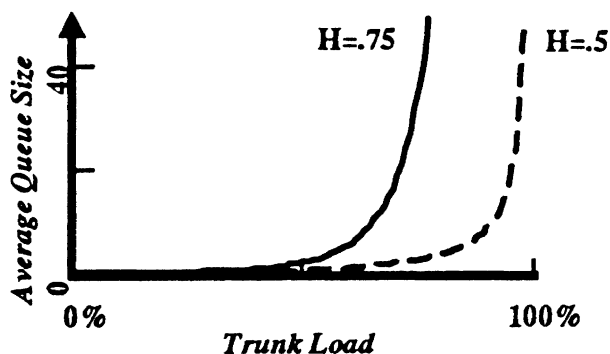


Figure 7.10 Queue Length vs. Trunk Load for $H = .75$ and $H = .5$

Figure 7.10 shows plots of the average queue lengths for switches with infinite length buffers. At any specific instant in time, the actual queue length is likely to be greater than or less than this average. To determine the probability that a switch with a finite length buffer overflows, which will impact the QoS hence the allowable load, what is needed is the distribution of the queue lengths as a function of the offered load traffic mix and the H parameter of that mix. Real-world distributions are generally *extremely* difficult, if not impossible, to find because they are directly impacted by the queue handling schemes of particular manufacturers and protocols, which are often quite complicated. Until research yields a simple and reasonably accurate solution, we suggest setting the maximum trunk load such that the average queue size predicted by Equation 7.6 is significantly less than the trunk queue size available in the switch. For comparison purposes, this chapter has standardized on an 80% maximum trunk load for all systems.

Considering the above information, estimates of a packet-switched network's carrying capacity can be obtained in the following manner:

1. Choose the target system-wide average end-to-end delays for *both* your time sensitive and data traffic, and estimate the average queuing delay allowable through a typical switch.
2. Estimate the average packet size and overhead associated with bursty data traffic and time sensitive traffic.
3. Estimate the Hurst parameters associated with your traffic. Doing this accurately may be somewhat difficult as determining the Hurst parameter from finite amounts of data is notoriously inaccurate.⁴ What is known is that a Hurst parameter of .5 (meaning no self-similarity) is *known to be inaccurate for data*. A Hurst parameter of 1.0 must also be inaccurate, because it would imply that traffic plots would look similar if plotted on any scale. This is clearly incorrect for real-world traffic, as different 'zooms' will yield nonsimilar plots. Consider Figure 7.2 if you've 'zoomed' down to a single bit. A value of .75 is tentatively suggested

for use as a compromise in the event that additional information is lacking, as this value lies in the middle of the extreme Hurst parameter values and is also near the middle of the ranges noted for actual traffic from preliminary studies.

4. Estimate the maximum load your switches can reliably place on the output trunk lines. Trunk loads exceeding this value are assumed to result in intolerable amounts of packets being dropped due to finite buffer sizes. This parameter will impact the carrying capacity under heavy load conditions, where the queuing delay is easily met but the fear of overflowing the switch buffer limits the trunk loading.
5. Use weighted averages of steps 1–3, above, to account for the appropriate traffic mix.
6. Then use Equation 7.7 to solve for the average queue lengths.
7. Use Equation 7.6 to solve for the trunk loads.
8. Bound the Trunk Load by the value in step 4 if necessary.
9. Use Equation 7.4 to compute the carrying capacity.

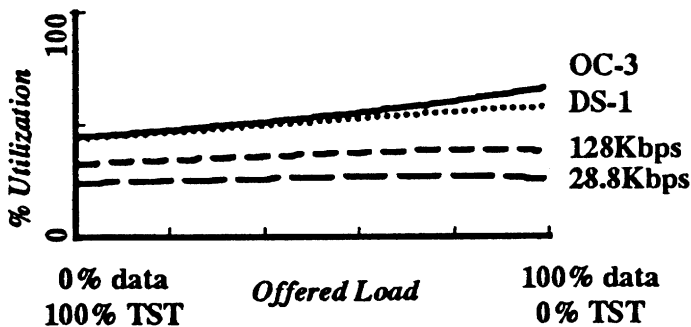


Figure 7.11 Packet Switch StatMux Trunk Utilization

Figure 7.11 shows some plots of packet-switched StatMux utilization as the switch offered load varies from 100% time sensitive to 100% bursty data traffic, for different trunk line speeds. These plots are based on the following assumptions:

- Average queuing delay through a network packet switch for time sensitive traffic is 20 msec, 40 msec for data. IPv4 is being used with no QoS provisions enabled, meaning all traffic must be moved through a switch with an average queuing delay of 20 msec in order to meet the tighter TST requirements.
- The Hurst parameter associated with both the data and time sensitive traffic is .75, a value believed to be a reasonable compromise based on some preliminary studies.
- Maximum reliable load that a packet switch can place on its output trunk is 80%.

- Average packet size of the data traffic is 300 bytes.¹ As mentioned earlier, the overhead would consist of 46 bytes, 6 bytes of Level 2 overhead for HDLC, 20 bytes of Level 3 overhead for IPv4, and 20 bytes of Level 4 & 5 overhead for TCP, leaving 254 bytes for the application.
- Time sensitive traffic is assumed to be mainly 8 Kbps compressed voice being moved at an average rate of 20 packets/second (50 bytes of voice + 8 bytes of user datagram protocol (UDP) overhead + 20 bytes of IPv4 overhead + 6 bytes of HDLC overhead).

Note that of the parameters listed above, the values that most affect the carrying capacity at broadband rates are the packet sizes (smaller packets have a larger percentage of overhead), and the maximum reliable load that the switches can support. With high speed trunks the carrying capacity will often not be limited by the allowable average switch queuing delays, but instead will be limited by switch buffer sizes, i.e., the switch will often be operating under heavy load conditions.

Figure 7.11, shows that with high speed trunks the small packet sizes required for timely delivery of digitized voice adversely impact the network's carrying capacity. Larger voice packets would improve the utilization, but at the same time they would drive down the quality perceived by the end user by increasing the end-to-end delivery delay. Broadband packet-switched StatMux networks offer the highest carrying capacities if they carry the type of traffic they were originally designed for, bursty data traffic.

Not evident from this plot is that increasing the trunk line speed to greater than OC-3 rates will not yield any additional utilization benefits, if the heaviest load that a switch can reliably place on the trunk line is 80%. Under this condition, a plot of OC-12 carrying capacity is virtually identical to that of OC-3. If a switch could handle a trunk load greater than 80%, which, depending upon the switch configuration, may very well be possible due to increased buffer sizes or the increased StatMux gains available using larger trunk sizes, these systems would show slight utilization increases per Equation 7.4.

At lower line speeds, the packet sizes, coupled with the choice of average switch queuing delay for this example, require that the trunks be lightly loaded, limiting the overall utilization.

7.6 ATM STATISTICAL MULTIPLEXED TRUNKS

As is noted in Figure 7.5, in this technique all traffic is inserted into fixed-size 53-byte cells and multiplexed onto a high speed trunk prior to insertion into fiber for transmission.

Fixed-rate traffic is best treated as a native ATM application hauled via CBR using ATM Adaptation Layer One (AAL1), which adds one byte of overhead per cell for sequencing purposes. As a result, 47 of the 53 bytes are available to carry traffic. ATM switches can offer TDM-like services to CBR traffic, reserving an appropriate number of cells at regular time intervals for this class of service.

Bursty traffic is normally carried via either VBR, ABR, or UBR classes of service, which are StatMuxed onto the remaining trunk bandwidth not reserved for

CBR traffic. In this chapter, bursty traffic is assumed to be passed down to AAL5 in the form of IP packets. AAL5 adds 16 bytes of overhead to each packet prior to segmentation.

Similar to what we saw in Section 7.5, the carrying capacity of ATM trunks can be calculated via

$$\text{CapATM} = \frac{(\text{Average application traffic per cell}) \times (\% \text{ Usable Line BW}) \times (\text{Trunk Loading})}{(53 \text{ bytes})} \quad (7.8)$$

The key difference between Equations 7.8 and 7.4 is how the trunk loading is treated. In ATM, since fixed rate sources can be given TDM-like service by reserving specific cells for CBR traffic, the trunk loading for CBR under heavy load conditions is 100%. Bursty traffic would be StatMuxed onto the remaining trunk bandwidth not reserved for CBR service. Note that for a trunk with a fixed-amount bandwidth, as the offered load is varied from 100% bursty traffic to 100% fixed-rate traffic, the bandwidth available for StatMux use will decrease as more and more will be reserved for the fixed-rate traffic. Otherwise, the same technique used in Section 7.5 is used to estimate the carrying capacities here.

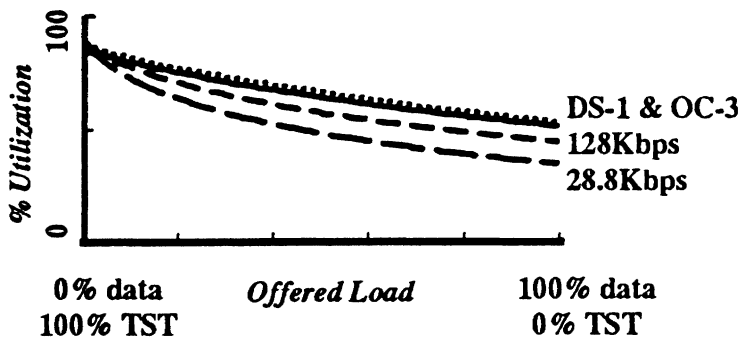


Figure 7.12 ATM Switch StatMux Trunk Utilization

Figure 7.12 shows a plot of ATM utilization as the switch offered load varies from 100% time sensitive to 100% bursty data traffic, for different trunk line speeds. This plot is based on the following choices:

- Average tolerable queuing delay through a network StatMuxed cell switch is 40 msec for data traffic, the same as in the previous section. These delays would be the average delay of all moved VBR, ABR, and UBR cells.
- The Hurst parameter associated with the bursty traffic is .75.
- The maximum reliable load that a cell switch can StatMux onto its output trunk is 80% of the line speed not reserved for CBR traffic.

- Average packet size of the data traffic offered to AAL5 is 300 bytes.¹ An ATM switch would first drop the overhead associated with HDLC and, as mentioned earlier, would then add 16 bytes of AAL5 overhead to each packet. The result would then be segmented into 48-byte chunks for insertion into ATM cells.
- Voice and video traffic is a fixed-rate native ATM application.

As with the packet-switched StatMux case, of the parameters listed above the values that most affect the carrying capacity at broadband rates are the packet sizes (smaller data packets offered for segmentation have a larger percentage of overhead) and the maximum reliable load that the switches can support. Note the ability of ATM to offer reasonably high utilization at low speeds. The smaller fixed-sized cells allow a higher load to be placed on the outgoing trunk while still meeting switch average delay specifications.

7.7 HEAD-TO-HEAD COMPARISON

It is illuminating to plot the carrying capacities of the four types of networks on a single graph for comparison purposes, similar to Figure 7.7. Figure 7.13 does so for OC-3 trunks. Note the following:

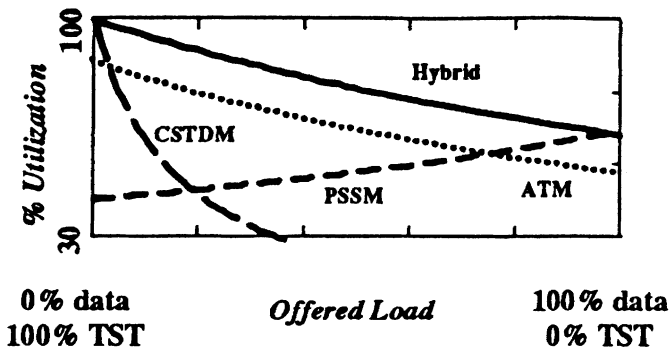


Figure 7.13 OC-3 IPv4 Head-to-Head Comparison

- The circuit-switched TDM backbone offers its highest carrying capacities if the offered load is almost 100% fixed rate. It *rapidly* falls off as bursty data becomes a larger percentage of the load, due to the well-known inability of this technique to efficiently carry bursty traffic. It is capable of hauling fixed-rate voice and video with a minimum amount of overhead.
- Packet switching and StatMuxing, which were originally designed to haul bursty data, not surprisingly haul this type of traffic best. However, when time sensitive traffic such as voice is offered, the overhead associated with packetizing this traffic seriously impacts the utilization. Given voice traffic with either fixed or variable bit rates, a packet-switched StatMuxed

network *cannot* match the utilization that a circuit-switched TDM network can achieve with fixed-rate voice traffic, provided that the average bit rates of the voice sources are the same.

- The hybrid backbone uses the best of both worlds, circuit switching and TDM for fixed-rate time sensitive traffic, and packet switching with Stat-Muxing for bursty data traffic. Provided that the load is segmented properly, this technique can *potentially* offer the highest possible overall utilization. Note, however, that if the traffic is not properly segmented, if some of the bursty data traffic is transmitted over the circuit-switched TDM network, then the *average* peak-to-average ratio will go up, reflecting the fact that more of the data traffic will not have been consolidated onto a small number of heavily used StatMuxed trunks. Depending on the degree of segmentation, the utilization of a hybrid backbone could lie between the plotted values (for 100% segmentation) and the circuit-switched TDM backbone (for 0% segmentation).
- ATM hauls no specific type of traffic best, but instead is clearly well suited for the mixed traffic environment for which it was designed. Its ability to offer different classes of service to different traffic sources allows it to follow in the shadow of the hybrid network in terms of carrying capacity. It cannot quite match the hybrid network's utilization due to the additional AAL and cell overhead, as well as the fact that it is a compromise, not tailored to a specific type of traffic, as are the circuit-switched TDM and packet-switched StatMux techniques. While ATM suffers a common problem with the hybrid network in that improperly segmented traffic will reduce the system carrying capacity, it is potentially far easier to properly classify the traffic because each flow can be assigned an appropriate class of service by the end user based on cost and desired quality, without constant carrier oversight.

Clearly, in terms of the network carrying capacity, different traffic mixes are best served by different trunking technologies.

Figure 7.14 shows essentially the same plot except that IPv4 has been replaced with IPv6. Shortcomings associated with IPv4 have resulted in the development of IPv6 which adds additional features at the cost of additional overhead — primarily larger source and destination address fields. IPv6 is expected by many to see significant deployment around the turn of the century. This change balloons the IP header from 20 to at least 40 bytes. Additionally, in the plots shown, IPv6's priorities are assumed to be enabled such that packet switches are able to meet a 20 msec average delay for the time sensitive traffic, and a 40 msec average delay for the bursty data. The overall result is that the utilization crossover point of ATM and packet switching moves from about 75% data to 85% data. In terms of utilization, even though the processing requirements at packet switches are relaxed to 40 msec for bursty data traffic, the additional IP overhead clearly makes the case worse for a 100% Internet backbone. Not evident from the plot, however, is that the use of these priorities would improve the quality of TST.

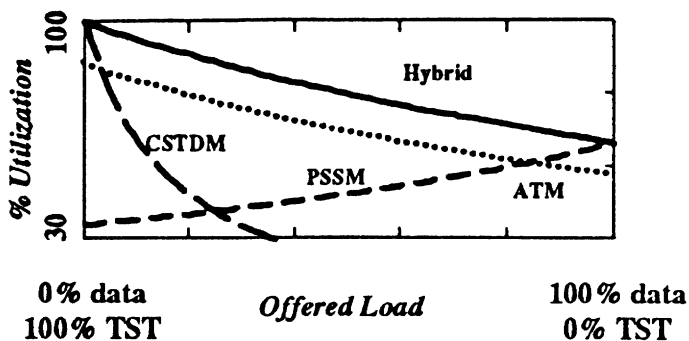


Figure 7.14 OC-3 IPv6 Head-to-Head Comparison

It was mentioned earlier that, at least on the WAN, carriers that have been around for awhile and have much capital sunk into older technology have tended to deploy hybrid-type networks. Recent buildouts of newer carriers, who don't have to worry about backward compatibility, have tended towards consolidated techniques whereby *all* traffic is carried on the backbone over a single, core technology. The two techniques most heavily touted have been 100% packet-switched StatMuxed backbones, specifically the IP-based Internet, and ATM. The carrying capacity provides useful insight into which technique has the potential to be the lowest cost solution, provided one can nail down the current and future offered traffic mix. Today, data traffic is clearly growing at a faster rate than that of time sensitive traffic, but will that remain the case in the immediate future? As the cost of bandwidth declines, how will real-time video traffic grow? Plenty of science fiction movies show high fidelity, interactive video conferencing as commonplace as the current telephone system. Can they all be wrong? Future communications will certainly require interactive video that is both time sensitive and possibly bursty in nature. Clearly, TDM-based networks will be left behind. The only remaining question is what technologies (IP, ATM, or both) are best suited to implement the necessary flow-based queuing and bandwidth reservation schemes.

Figure 7.15 shows a plot of the carrying capacities of circuit-switched, ATM, and packet-switched trunks running at DS-1, a speed commonly used for corporate enterprise WAN connectivity. Shown is the case when IPv6 and priorities are enabled. A plot for IPv4 looks almost identical; the inability of IPv4 to load trunk lines more heavily by prioritizing traffic (meaning that in order to meet time sensitive traffic delay criteria it must also whisk data through switches at TST rates) is compensated by the smaller percentage of packet overhead. Interestingly, at slow speeds ATM's utilization tends to dominate for almost any traffic mix. Given a target average queuing delay through a switch, slower speed connections are more likely to be delay-constrained than buffer-constrained. Equation 7.7 indicates that for identical target queuing delays, an ATM switch will have on average about six times the 53-byte cells queued up than the number of queued packets in a packet switch moving 300-byte packets. From Equation 7.6 and Figure 7.10 it can be seen

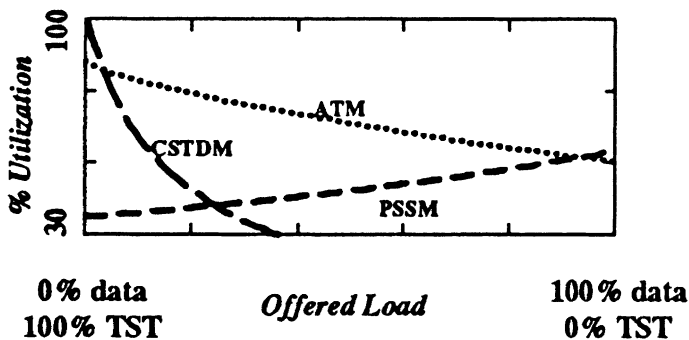


Figure 7.15 T-1 IPv6 Head-to-Head Comparison

that this allows an ATM trunk line to be more heavily loaded than the packet switch trunk line. The extra loading goes a long way towards canceling the extra overhead associated with chopping up the packet into smaller cells, resulting in ATM's carrying capacity almost matching that of packet switching for an offered load of 100% data. ATM warrants more consideration as the switching and multiplexing technique of choice for low speed connectivity than it is currently receiving.

7.8 CONCLUSIONS

The price a network technology pays to haul application traffic is clearly influenced by the value of the network's utilization, or carrying capacity. A network that requires a lot of overhead or is unable to load its trunk lines heavily is a network with a lowered carrying capacity. While this parameter, as discussed in this chapter, does not account for signaling overhead, the granularity associated with various trunking protocols such as SONET, or the impact of higher level protocols such as TCP, it nevertheless provides important insight into which techniques are best suited for hauling different traffic mixes. A technology with a higher utilization for a given traffic mix requires less trunk bandwidth and lower switch speeds to move a given amount of application traffic than a protocol with a lower utilization. This directly impacts the bottom line. Depending on the relative equipment costs, the network with the higher utilization potentially costs less to deploy.

Of the switching and multiplexing mixes discussed, the hybrid network clearly offers the highest *potential* utilization, but it suffers two key drawbacks. It requires twice the hardware of all-ATM, all-packet-switched, or all-circuit-switched networks, and it requires careful grooming to achieve these high utilization values.

Of the consolidated networks, a circuit-switched TDM backbone works best for fixed-rate time sensitive traffic, which makes it a *horrible* choice for today's traffic mix which is increasingly dominated by bursty data. The packet-switched StatMux network has the highest utilization for bursty traffic, and ATM is best for a mixture. Not surprisingly, each technology has the highest carrying capacity when used to haul the traffic mix for which it was originally designed.

The overall choice for deployment or upgrading of a network depends on at least three key issues: installation cost, maintenance costs, and reliability.

Many factors impact the installation cost of a network, including power requirements, bay size, backward compatibility needed, switch speeds and configuration, trunk sizes, and last-but-not-least, carrying capacity. Knowledge of the carrying capacity is vital here, as it enables an analyst to get a better idea of the trunk bandwidth and hardware speed requirements of networks hauling equivalent loads, factors that directly impact the price of installing the network.

Maintenance costs reflect the organizational, administrative, and day-to-day operating costs associated with running the network after it is installed. A backbone having the fewest *types* of equipment has significant advantages here, as less effort will be required to integrate disparate hardware with the network control center, and fewer engineers and technicians with expertise on specific pieces of equipment will be required. A network technology ready-for-prime-time is also likely to have reduced maintenance costs compared to a network still in the 'Bleeding Edge' stages.

Reliability refers to the network's ability to maintain maximum uptime and its susceptibility to a catastrophic failure. Barring that catastrophic failure, all four backbone techniques should be engineerable to equivalent 99.99% uptimes. Catastrophic failures which bring down the *entire* network are extremely rare but can have disastrous effects. One has only to look at recent Internet or frame relay events to wonder whether putting all one's eggs in a single basket is a good idea. For example, in July of 1997 a large chunk of the Internet was isolated when Network Solutions botched a top level domain name server update. In April of 1998, AT&T's entire frame relay network went down for over a day during a switch upgrade of improperly tested software. There *is* something to be said about maintaining a hybrid network which would be more difficult to bring down totally.

Despite some pundits' claims to the contrary, the choice of the 'best' network technology is not clear-cut, as that choice depends on intangibles, that are frequently hard to quantize, and the cost that one is willing to pay. That cost can be severely impacted by the parameter which is the focus of this chapter, trunk carrying capacity. Can 155 mbps of Internet bandwidth haul the same amount of customer application traffic as 155 mbps of ATM? The often overlooked carrying capacity will tell you.

ACKNOWLEDGMENTS

This work was supported in part by a contract from the Williams Communications Group, Tulsa, Oklahoma.

REFERENCES

1. Thompson, K., et al., 'Wide Area Internet Traffic Patterns and Characteristics,' *IEEE NETWORK*, November/December 1997, p. 10-22.
2. Stallings, W., *HIGH SPEED NETWORKS: TCP/IP AND ATM DESIGN PRINCIPLES*, Prentice-Hall, 1998.

3. Crovella, M. and Bestavros, A., 'Self Similarity in World Wide Web Traffic: Evidence and Possible Causes,' *IEEE/ACM TRANSACTIONS ON NETWORKING*, December 1997, p. 835-846.
4. Michiel, H. and Laevens, K., 'Teletraffic Engineering in a Broad-Band Era,' *IEEE Proceedings*, December 1997, p. 2007-2033.

WEB ASSISTANCE

Interested in evaluating the carrying capacity for your current or proposed network? Don't agree with the choice of parameters used here and want to examine the results with different selections? A downloadable MathCad® executable file can be obtained at http://www.mstm.okstate.edu/files/www/faculty/scheets/pub/wcg_cap4.html. A non-executable Word document is available at http://www.mstm.okstate.edu/faculty/scheets/pub/wcg_cap4.doc