

Strategies for Optimizing Bandwidth Efficiency

Todd Landers
SSC San Diego

DATA CHARACTERISTICS

Loss-Sensitive Information

The data type envisioned when discussing data communications is usually loss-sensitive data. Many network engineers mistakenly assume that loss-sensitive data are the predominant data type on the Navy's general-purpose wide-area networks (WANs). This data type must be faithfully reproduced with 100% accuracy at the distant end of the link before it can be used. The Transmission Control Protocol (TCP)/Internet Protocol (IP) usually transmits the data type as computer data/software. Any differences in the data reproduced on the distant end of a link make the entire communication unusable. If data are corrupt, the application of the data will be invalid. The time it takes the data to reach the distant end of the link has little effect on data usability. This data type is loss intolerant, but latency tolerant.

As mentioned, the bulk of this data type is transferred using the IP family of protocols. IP is inherently connectionless. It uses a 32-bit address scheme to identify hosts on the network. The User Datagram Protocol (UDP) and TCP use IP as its transport layer.

For broadcast applications, UDP uses the connectionless properties of IP to its advantage. It uses a checksum to verify data integrity and discards corrupt data. UDP is ideal for applications that are loss and latency tolerant, and can be used in those instances to help minimize unneeded router chatter over bandwidth-constrained links.

TCP adds a connection-aware element on top of IP. TCP has embedded mechanisms that check the content and sequence of arriving packets. TCP also allows hosts to set timers. The host may "time out" a connection, enabling the host to free up system resources that would otherwise be tied up with a suspected dead connection. TCP will automatically request a re-send from the originating host if loss or corruption is detected. TCP uses these and other tools to keep the network working smoothly as long as latency is managed at lower levels of the network.

Unfortunately, the same features that help TCP work well in situations where bandwidth is ample can be disastrous when bandwidth becomes constrained. Once a link in the network becomes bandwidth-constrained, the applications start asking for retransmission of data assumed lost (in this case, just delayed). This unnecessary request for information is the

ABSTRACT

To optimize bandwidth efficiency, the natural limitations of each network-supported data type must be overcome or mitigated. This paper discusses issues affecting bandwidth efficiency through the U.S. Navy's bandwidth-constrained wide-area network (WAN). The paper details the prevalent data types found in the naval environment and describes the characteristics associated with each data type. Commercial, standards-based link layer protocols that have widespread application in Navy networks are also described. Finally, forward error correction and issues surrounding bandwidth efficiency are discussed.

beginning of the end of data transfer across the link. If Ethernet is used as the link layer protocol, connection timeouts caused by long round-trip time can cause a router to start seeking alternate paths to the desired host. The additional router chatter contributes to the congestion of the already congested link. This congestion is a death spiral for a TCP connection. The connection is terminated, and if the host is looking for the original information, it attempts to reconnect. Note that no useful information is exchanged, though bandwidth is consumed. The bandwidth consumption prevents other worthy circuits from exchanging useful information.

The way applications use IP may cause other inefficiencies. If the payload of the IP packet is not appropriately sized for the data type conveyed, huge amounts of bandwidth could be consumed because bit stuffing is needed to make complete packets. The data type usually originates at hosts that provide sensor inputs (like voice) to another application. If a sensor needs to transmit a sample containing a few bytes to a remote host through TCP, the originator usually stuffs filler bytes into a packet with a length that is probably several hundred bytes. The efficiency of this connection quickly approaches zero, which is not a problem until bandwidth becomes limited. File compression can reduce the size of an application data file or sensor output before transmission through the WAN. File compression engines must be deployed to all source and user sites, which causes a logistics problem, but the overall gain in bandwidth efficiency is worth the trouble. File compression can reduce the actual data transmitted across the link by 80%.

Header compression techniques can reduce bandwidth consumed over a point-to-point link through various network protocols. For header compression to add value to the WAN, it must add minimal latency due to processing overhead and be completely symmetrical. Compression abbreviates redundant header information in a data stream before transmission over the WAN and then restores the header to its original state after it reaches its final destination. The higher the compression engine is in the protocol stack, the more opportunity to save bandwidth. The more aggressive the compression engine, however, the more latency it adds to the circuit. Header compression at the transport and network layers depends on some error checking at the link layer to be effective.

Time-Sensitive, Loss-Tolerant Information (Voice)

The most common type of time-sensitive, loss-tolerant data is plain old telephone system (POTS). Unlike computer-oriented information, the interpretive device for POTS is the human ear. Studies show that over 50% of a speech signal can be removed and the human ear can still assemble the required information to extract the audible message.

However, as little as a 0.5-second delay can cause severe degradation of the intended communication. A system designed to handle large quantities of this data type can lose a lot of data, but if data are delayed or delivered out of order, it is useless, and interpreted as noise.

Networks specializing in this data type are quite different from those that handle large quantities of loss-sensitive data. In terrestrial networks where bandwidth is ample, a typical voice call is digitized and transmitted using a G.711 protocol through the public switched telephone network (PSTN) at 64 kbps. Toll-quality voice has an upper latency limit of a 200-ms delay across the network. These networks are composed of various

sizes of public branch exchange (PBX) switches. The interconnections between PBXs generally scale in 64-Kbps chunks.

PSTNs are connection-oriented. When a call initiates, the originator transmits a setup preamble that negotiates for a connection at each intermediate switch along the way. If the connection cannot be supported at any point along the way, the entire connection is denied. If all attempts to establish the end-to-end connection are denied, the originator gets a busy signal. For most PSTN users, this busy signal only happens on Mother's Day or after a natural disaster. There is no such thing as a lower grade of service; a connection exists or it does not. During the call, a near-real-time connection for $N \times 64$ kbps allows the user to talk, send a facsimile (FAX), or use a modem, secure telephone unit (STU), or secure telephone equipment (STE), etc. After the call is completed, a teardown sequence allows each switch in the circuit to release the resources reserved for that connection.

This type of network has many sources of inefficiency. First, voice is the most common type of connection supported through this network. Voice typically has less than a 50% duty cycle. Generally, only one person talks at a time, and usually there is silence between words. Everything else is dead air (wasted bandwidth). Silence-suppression techniques reduce the dead-air bandwidth consumption to help solve this problem.

Another source of inefficiency is that 64 kbps is not really needed to digitize and communicate using voice. The 64-kbps convention was adopted because it was an easily implemented solution, not because it was the most efficient. Several voice compression algorithms can drastically reduce the amount of bandwidth used for each voice call. Toll-quality voice has been compressed to 8-kbps or one-eighth of the bandwidth allotted for a typical voice call. Good-quality voice has been transmitted using less than 800 bps. Unfortunately, other applications using the PSTN do not respond well when compressed with some of the more aggressive compression techniques, so compression must be applied selectively.

Modems and FAX machines use the PSTN to transmit analog-modulated digital signals. This inefficient means of digital data transfer was developed years ago to overcome noisy analog transmission lines that were once used to interconnect PBXs and end-users. Connections have improved, but the format is outdated. The most efficient way to accommodate these types of connections is to convert the modulated digital signal back into ones and zeros and transmit them through the network using much less bandwidth. A FAX machine can be supported at 9.6 to 14.4 kbps instead of consuming the full 64 kbps allocated to each connection by the PSTN. STU-III can also be supported using this type of compression technique, but the modulation scheme must be implemented in accordance with National Security Agency (NSA) policy.

Time and Loss-Sensitive Information

The synchronous serial data type is traditionally used where the system designer had creative control of the entire system. These links are susceptible to loss of content and fluctuations in the end-to-end timing. Each communication link was usually built to support one application set. Interoperability and flexibility were not considered in the design. These systems are probably the single largest source of wasted bandwidth. After the communication link initiates, it remains active regardless of use.

Circuits had to be provisioned to support worst-case bandwidth needs, and as applications became more bandwidth-efficient, their bandwidth usage remained high and constant.

As networked applications became more popular, synchronous serial communications became known as communications "stovepipes." The Department of Defense has invested huge amounts of resources into developing and refining stovepipe systems over the past 30 years. Although new systems focus more on networked solutions, stovepipes are still with us today primarily because of the cryptography developed to support legacy applications. The slow development of network cryptography has hindered application development and subsequent migration away from stovepipes.

Video

Video is generally more tolerant of timing than synchronous serial connections, but jitter is deadly. The type of video compression used should vary depending on the video content. Compressed video usually transmits all information needed to paint the screen the first time, and transmits only the changes to the initial image. Regardless of the resolution or quality of the video, this approach allows video to be supported using variable bit-rate service contracts through the network. Talking-head videoteleconference (VTC) video should use the most aggressive video compression techniques. This type of application can operate well on less than 64 kbps.

One of the worst misuses of bandwidth for video traffic occurs when the host or the network provisioning creates a fixed-bandwidth pipe for the video call. H.320 is a common video compression format used with Integrated Services Digital Network (ISDN) networks. Each video call allocates $N \times 64$ kbps to support the call resolution selected by the user. The bandwidth within the fixed allocation continues to fluctuate; however, even though the bandwidth need reduces when the picture stabilizes, no bandwidth is available for other applications.

Although still both time and loss sensitive, the H.323 protocol is much more tolerant in both areas. H.323 works with IP networks and has progressed in overcoming some of the inherent obstacles for supporting voice and video on an IP network.

Unfortunately, H.323 is susceptible to many problems that plague data transmission over IP-based WANs. While some jitter or delay can be tolerated, excessive congestion can cause the H.323 session to freeze. Though there is some time-sensitivity, packets are mixed with and sometimes delayed because of packets that have no time-sensitivity. To overcome this deficiency, priority queuing can reduce the likelihood that the time-sensitive video will incur fatal transmission delay.

The screen capture in Figure 1 shows the bandwidth consumption of the H.323 protocol generated using Microsoft Netmeeting. The link information shown represents one-half of a bidirectional link. The video resolution for this example was set up to run at best-fidelity voice and video. H.323 is highly variable in its bandwidth requirement (red trace in Figure 1). The

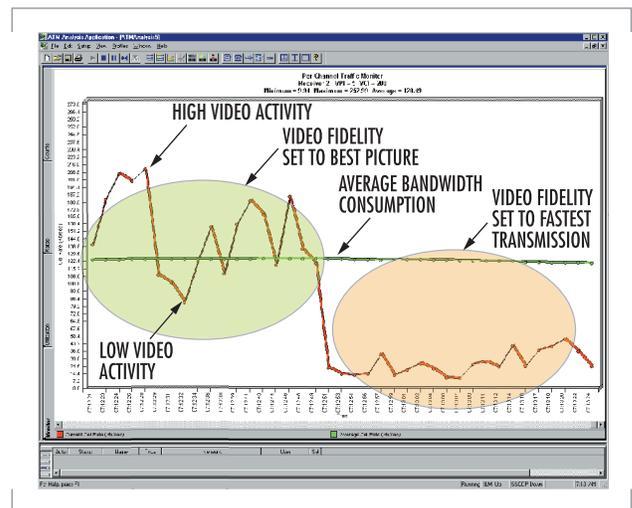


FIGURE 1. Bandwidth consumption of the H.323 protocol.

surges in bandwidth usage occur when the video compression engine must transmit updates to large portions of the image. With video quality set to its highest setting, the average bandwidth usage settles out near 130 kbps, with surges up to nearly 190 kbps. During periods when the video does not change, bandwidth usage drops to below 80 kbps. Netmeeting allows the user to reduce the fidelity of the picture to accommodate bandwidth-constrained connections. At the lowest resolution settings, the average bandwidth usage has been observed below 20 kbps for a full-motion VTC.

TYPICAL NAVY WAN DATA LOADING

The U.S. Navy operates in a truly converged WAN environment. Figure 2 shows a sample of the circuits assigned to USS *Coronado* (AGF 11), a command ship, during a typical deployment. The bandwidth available on *Coronado* should be considered the best possible case because she has been outfitted with the best communications available in the U.S. Navy to support various developmental enterprises.

Coronado runs multiple T-1s using various super high-frequency (SHF) and Challenge Athena configurations.

Voice, video, and data must all co-exist on the WAN (Figure 2). The major users include a Joint Service Imagery Processing System (JSIPS), which is an intelligence circuit currently using a synchronous serial EIA-530-based system. This circuit was one of the primary reasons for the procurement of the Challenge Athena system, so when this circuit becomes active, other lower priority circuits are manually disconnected. The JSIPS circuit is a prime example of current bandwidth management practices. Other large data users include secure and non-secure voice (both circuits are listed in the figure as POTS LINE). POTS lines are supported using compressed voice cards in various time-division multiplexers (TDMs). The compression cards reduce the bandwidth required to support each voice call from 64 kbps to between 8 and 16 kbps. The compressed voice signals are aggregated as synchronous serial circuits before porting to the satellite communications (SATCOM) modems.

In the bandwidth management approach, fixed-bandwidth synchronous serial circuits constrain circuits that use protocols that dynamically consume bandwidth such as IP. The circuits marked "ADNS" (Automated Digital Network System) represent the wide-area IP-based traffic, and typically are assigned up to 384 kbps, supporting a mixture of classified and unclassified data.

Figure 3 shows how the circuits on the Y-axis might consume bandwidth when provisioned through a TDM using today's provisioning approach. The white space represents provisioned, but unused, bandwidth. A few points to notice in this figure are as follows:

- Wasted bandwidth by low-usage, high-bandwidth systems such as the Video Information Exchange Subsystem (VIXS), which is a H.320-based VTC using synchronous serial cryptography for security

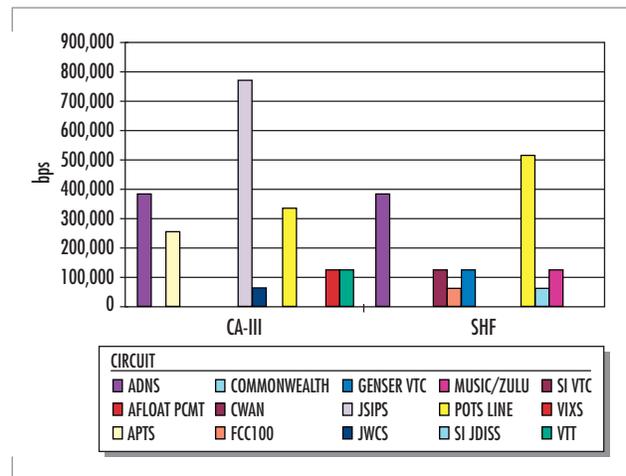


FIGURE 2. Circuits assigned to *Coronado* during typical deployment.

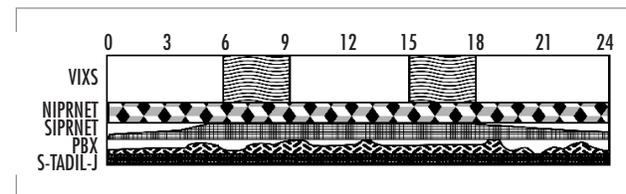


FIGURE 3. TDM circuit.

- Congestion in the low-priority, high-usage circuits such as NIPRNET (unclassified but sensitive IP router network)
- "Nailed up" circuits consuming bandwidth to keep the circuit timing alive, such as Satellite Tactical Data Information Link-Joint (S-TADIL-J)
- More bandwidth available to support additional voice and data circuits than are allowed under current provisioning policy

LINK-LAYER CONSIDERATIONS

Ethernet

Ethernet is commonly used as the link-layer protocol for TCP/IP and UDP/IP. Ethernet is a very cost-effective way to deliver data to end-users. The network equipment is inexpensive and mature; there is a large application base with drivers supporting Ethernet; and there is a large pool of competent network administrators who understand the technology. Ethernet is inexpensive to deploy and administer. It scales very easily to 10 Gbps on the backbone. Ethernet uses the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) approach to sharing bandwidth among the network users. If there is a collision, it simply re-sends the information. Ethernet works best on generously provisioned networks; for the most part, the assumption of few collisions is true.

Ethernet starts having difficulty when congestion occurs. At approximately 40% of the rated network throughput, Ethernet begins to bog down. At 60% link saturation, the link becomes nearly unusable because much of the traffic is re-sent from prior collisions.

Switched Ethernet technology provides answers to some of these problems by explicitly controlling traffic destined for users. An Ethernet switch can support single dedicated or multiple users on a switched segment. The switch logically separates the segments to eliminate collisions between segments. Multiple user segments continue to have the contention problem among subnetwork users.

Ethernet is problematic in the wide area because the organization deploying the local network will not have control over congestion in the wide area. Any single link in the wide area will slow performance experienced by the user. Network performance is only as good as the slowest link in the WAN.

Point-to-Point Protocol

Point-to-point protocol (PPP) is an encapsulation approach to transmitting IP over serial point-to-point links. PPP is a very flexible approach to transmitting IP datagrams through a serial link. The only real limitation PPP imposes on the link is that it has to be a full-duplex link. It supports synchronous and asynchronous transmission. It works fine over a number of common physical media including EIA 530, RS-232, V.35, etc. However, PPP links can cause unwanted latency and jitter because of the variable nature of the IP datagram contained in the data payload of the PPP frame.

ATM

Asynchronous transfer mode (ATM) effectively transmits a wide variety of data across a network. The size of the ATM cell (53 bytes) was developed as a compromise between the voice camp (small, prompt data

delivery) and the IP camp (large, continuous streams of guaranteed data delivery). While ATM was originally envisioned to work on high-speed networks (OC-3 and above), it has been adapted for the WAN because it works through congested links.

ATM statistically multiplexes fixed-size cells through a link. The protocol organizes cells into logical or virtual point-to-point circuits through an interface. At the time of circuit setup, each interface in the circuit establishes a service contract with its neighbors. Each switch has a unique address to ease automated connection setup. Once all of the interfaces in the path have established the required service contracts, the data transfer begins. Cells with that circuit identifier are automatically switched along its path to the end-user. Once the data transmission is complete, the contracts are canceled and the circuit is disestablished.

The service contracts have built-in quality-of-service features. At the top layer, there are ATM Adaptation Layers (AALs). Each AAL has some predetermined characteristics and some preconceived notions of what applications that adaptation layer would support. For example, AAL-1 supports synchronous serial connections and looks to users like a static TDM. AAL-2 supports voice, and while it maintains the time relationships between cells, it can take advantage of the other characteristics of voice discussed earlier. AAL-5 supports IP and has many features to take advantage of the characteristics of IP data transfer.

Forward Error Correction

Forward error correction (FEC) is commonly applied on noisy links to improve error performance and, thus, the performance of the link. The different FEC algorithms include block, convolution, and viterbi codes. For the purposes of this paper, FEC can be applied in varying degrees to reduce the error rate of a link; however, the more rigorously FEC is applied, the more bandwidth overhead and processing latency increases.

As discussed, different data types have varying degrees of error tolerance. FEC should be tailored to the data type passing through the link. For example, voice is loss tolerant and will perform well even when the link has some errors. Some synchronous data streams are very loss sensitive and may cause the end-user equipment to malfunction if there are too many errors on the link. TCP/IP may start flooding the link with re-sends if the error rate is too high, thus causing data congestion.

Recently, many products have been shipping with adaptive FEC. This approach samples the noise on the link and adjusts the FEC algorithm to keep the link error rate nearly constant. As the FEC is applied more aggressively, the effective throughput drops. Applying adaptive FEC at the link layer provides better performance by reducing the amount of re-sends by the hosted applications.

Link Design Decisions

As with all aspects of an engineered solution, engineers must choose the best tools to confront each aspect of the link design. Because of the economics and maturity of the technology, Ethernet is a clear choice in the local area networks, but falls short in the wide area. PPP is a good choice if all of the applications supported by the network are IP-based, but PPP falls short for a general-purpose network that supports various data types. With the technology currently available, ATM is the only technology

discussed that can support a truly converged network supporting voice, video, data, and legacy applications. Figure 4 shows the dynamic bandwidth allocation achieved using ATM for the WAN.

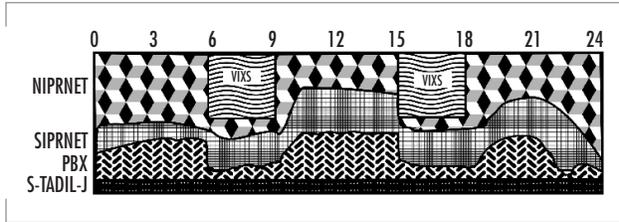


FIGURE 4. ATM circuit.

CONCLUSION

We can objectively maximize the information payload, optimizing for communication channel size, reducing non-productive information transfer, and using aggressive compression and forward error correction techniques. To get the greatest use from a SATCOM system, we must maximize the gross bandwidth efficiency. Additionally, user application data must be optimized. TCP/IP or UDP/IP solutions should be implemented wherever possible. Finally, the correct link layer technology must be selected for the environment in which it will be deployed.

Successful information transfer occurs only when enough data are transferred from the source through a data link to the distant end of the link to facilitate reassembly of the information suitable for end-device perception. Various types of information will be transmitted through our communications links. Some information is loss sensitive; other types of data are time sensitive; and still others are time and loss sensitive. They are all vital to the operation of the Fleet.

There are opportunities at every layer to optimize. The fiscal cost of not optimizing is tremendous. The operational cost could be devastating.



Todd Landers

BS in Electrical Engineering,
San Diego State University,
1985

Current Research: Wide-band
ADNS architecture develop-
ment; Tactical Switch System
development; DoD Teleport
requirements analysis and system
definition; design/test of the
IT-21 block one baseband
system.