Being Efficient with Bandwidth

Best practices for data interoperability and Efficient XML Interchange (EXI) compression can increase throughput in Navy tactical networks.

Motivation: Assured Command and Control

Naval Information Dominance (ID) hinges on three fundamental capabilities: assured command and control (C2), battlespace awareness, and integrated fires. None of these are possible without effective communications links. Networks, and more specifically, the information flowing through them, are now a center of gravity for the fleet.[1] Maritime tactics and operational plans rely on levels of synchronization only possible through high-bandwidth communications. Satellite communication (SATCOM) is the fleet's primary path for high-bandwidth C2. However, afloat units may be denied access due to equipment failure, technical problems, weather phenomenon, or enemy actions, forcing reliance on lower-bandwidth alternatives.

For afloat units, bandwidth has become a critical, but painfully finite, resource. Let's take a look at how we can conserve it.

For example, SATCOM carries data from a large number of disparate systems often referred to as "stovepipes." These systems vary in function from tactical to administrative, and the data formats for each application vary greatly. The result is communications only occurring vertically within a system, but not across the breadth of different systems. When many such stovepipes contend for access to the same ship-to-shore transport path, even the largest SATCOM channels can become congested. Future assured C2 requires interoperability between stovepipes and better prioritization of network traffic. Before identifying the solution we must understand the factors that impose constraints on the transmission path. These factors are bandwidth, latency, and throughput.

"Bandwidth" Is Not The Same As "Throughput"

Bandwidth is literally the "width" of the frequency band used to carry a data signal. It is more often described as the transmission capacity of the communications medium, measured in terms of bits per second.[2] In order to increase the capacity of an electromagnetic (EM) communications channel, modulation technologies and methods would need improvement or an additional antenna can be installed. Both approaches illustrate significant engineering and financial constraints associated with increasing bandwidth particularly in the shipboard

---

[1] CNO(ID). (2013). *Navy strategy for achieving information dominance: 2013-2017*. Washington, D.C. Retrieved from http://www.dtic.mil/docs/citations/ADA571217

[2] Gokhale, A. A. (2004). *Introduction to telecommunications* (p. 455). Cengage Learning. Retrieved from http://books.google.com/books?id=QowmxWAOEtYC&pgis=1

environment. SATCOM connections are often depicted as lightning bolts connecting deployed units with relay systems. These lightning bolts convey the impression that data are instantaneously transmitted from unit A to unit B through an optimally placed satellite node. Unfortunately SATCOM transmissions are far from instantaneous - they incur significant delays in comparison to terrestrial communications paths. The combined delay is known as latency.

Latency is an accumulated series of delays that can occur in each step of the communications path between sender and receiver. Such delays occur as part of propagation delay during signal transmission, network processing and interface delays, varying methods for buffering and queuing, and cumulative router and switch delays.[3] Latency from the perspective of network traffic is the delay from the time of the start of packet transmission at the sender host to the time of the end of packet reception at the receiver host. Unfortunately latency has significant effects on throughput. This is due to the primary networking protocols (TCP/IP) being designed to operate best with a latency less than 100 milliseconds. SATCOM channels routinely operate with latency between 500 to 800 milliseconds. Response "waiting time" is a particular problem for communications protocols that include frequent acknowledgement among participants. Increased latency ultimately results in decreased throughput.

Throughput is the rate at which new data – actual information – is transferred through a system. Like bandwidth, it is measured in bits per second, and can be considered the actual effective capacity of a channel or the "rate of successful message delivery" being achieved. A common misconception is that bandwidth and throughput are synonymous. Numerous additional constraints can limit the amount of data that can be transferred between two points, such as the overhead of communication protocols and latency delays, which may keep a channel idle. Thus bandwidth indicates the maximum possible data transfer capacity, while throughput is what capacity actually occurs. Throughput is often significantly lower than the communications channel's bandwidth capacity.

Common Practice: SATCOM

For Navy ships at sea, the only access to high bandwidth is through satellite communication (SATCOM) systems. In our increasingly connected world, the value placed on access to high bandwidth continues to rise. As bandwidth increases, the amount of data that can be transferred between two points also increases. As bandwidth is increased, additional capacity is quickly consumed by ever-more sophisticated sensors, unmanned vehicles, and other network-centric dependencies. Most high-bandwidth paths utilize the super- and extremely-high frequency (SHF/EHF) spectrum for SATCOM communications. Though data and voice circuits exist in other portions of the spectrum, SHF and EHF carry the brunt of the Navy traffic, with SHF (C/Ku/X band) ultimately providing the biggest "pipe" for data flows.

---

[3]Kay, R. (2009). *Pragmatic network latency engineering fundamental facts and analysis*. cPacket Networks. Retrieved from http://cpacket.com/wp-content/files_mf/introductiontonetworklatencyengineering.pdf

In the past, the solution to demand for increasing data transfer was to increase bandwidth, and thereby capacity. As the DoD throttles back spending, many areas must become more efficient in order to accomplish defense missions. Similar approaches for efficiency must be applied with respect to communication systems. The amount of information to be shared is not expected to decrease. Because constraints on SATCOM bandwidth make even marginal increases cost-prohibitive, the Navy must explore new tactics. Perhaps solutions lie not in the channel itself, but in the format of data transmitted? What if we can convey the same information using just a fraction of the original zeros and ones, while at the same time connecting stovepipes through data interoperability?

XML: The Language of Interoperability

Interoperability is essential to the key Information Dominance capabilities. Shipboard computers must talk to each other, to computers from other service branches, and to computers from partner nations. To facilitate interoperability, an open-standards approach is critical. The Department of the Navy Chief Information Officer (CIO) has designated the Extensible Markup Language (XML) as the data-definition language of choice for information standardization, and for good reason: it is the *de facto* standard format for systems talking across the web.[4] By design, XML adds structure to data, which in turn facilitates validation of correctness and system interoperability. XML is the *lingua franca* of the world's computers.

Though XML is a path to both technical and semantic interoperability, it has an Achilles heel: it was never intended to be compact.[5] In terrestrial networks with low latency contributing to massive throughput, this is usually unimportant. For the Navy, however, large messages mean slower connections and less information to forward-deployed units relying on SATCOM. Transmitting large messages also draws more power, so XML isn't ideal for mobile or unmanned devices running on batteries. Viewed in this light, XML is less attractive, but it doesn't have to be that way. Recent advances in data compression are providing new design options.

Efficient XML Shrinks Data, Broadens the Web

In 2004, the World Wide Web Consortium (W3C) began addressing this issue, and in 2014 released the Efficient XML Interchange (EXI) Format Recommendation.[6] EXI is an alternate encoding of XML data that leverages the inherent structure of XML to tightly compress it. Since it is designed specifically for XML, the results are superior to generic compression

---

[4] DoN CIO. (2012). *DON policy on the use of Extensible Markup Language (XML)*. Retrieved from http://xml.coverpages.org/DON-XMLPolicy200212.pdf

[5] Cokus, M., & Pericas-Geertsen, S. (2005). *XML binary characterization properties*. Retrieved from http://www.w3.org/TR/xbc-properties/#xml-design-goals

[6] Kamiya, T. (2014). *Efficient XML Interchange (EXI) working group*. Retrieved March 16, 2014, from http://www.w3.org/XML/EXI

methods. In some cases, EXI compression results in files that are less than 10% the size of the original XML file.[7] Perhaps even more surprising is that EXI decompresses faster, using fewer computations and therefore drawing less power than plain text-based ZIP and GZIP compression.

Given that XML enables interoperability, and that EXI shrinks it, fleet communications architects and program managers should be interested.[8] Systems could potentially convert and transmit information in XML format, and with EXI they could send more information in less time. By incorporating EXI, web-based architectures such as CANES and C4I systems using service-oriented architectures may be viable over constrained SATCOM links. Unmanned systems and remote sensors might use EXI to conserve batteries on extended missions. A single file cut to a tenth of its original size is useful in itself, but the aggregate impact over thousands of nodes in a cloud, each sending thousands of files, could be immense.

Other impacts pertain as well. For example, encryption is usually considered independent of compression. However, by randomizing a bit stream, encryption scrambles the structure necessary for effective compression. That means encrypted streams cannot be compressed! Compression must occur *before* encryption when transmitting, and decompression *after* decryption on the receiving end. This principle is so important that the order should be checked for all Navy communications channels.

Since message size is just one of many factors in network throughput, EXI is not a silver-bullet for Navy bandwidth woes, but it certainly can't hurt. It is not mutually exclusive of other attempts to address the issue. Navy communications designers need not choose between a new SATCOM constellation and EXI, or between commercial network accelerators and EXI; they can have both. Considering that EXI is an open standard, supports interoperability, and shrinks data the Navy is already sending over its networks, there is little to lose and much to gain. The Navy can be more efficient with a precious afloat resource: bandwidth.

Acknowledgements

Additional insights by Dr. Dan Boger, CAPT Louis (Tim) Unrein USN and other reviewers are gratefully acknowledged.

---

[7] Snyder, S., with advisors McGregor, D., & Brutzman, D. (2009). *Efficient XML Interchange(EXI): Compact, efficient and standards-based XML for modeling and simulation*. Retrieved from http://calhoun.nps.edu/public/handle/10945/5422

[8] Williams, J. (2009). *Document-based message-centric security using XML authentication and encryption for coalition and interagency operations*. Naval Postgraduate School. Retrieved from http://calhoun.nps.edu/public/handle/10945/4610

About the Authors

Lieutenant Commander Debich and Lieutenant Hill are Information Professional Officers studying Network Operations at NPS.

CAPT Scot Miller USN (Ret.) is a Research Associate with the NPS Cebrowski Institute and former commanding officer of the Navy Center for Tactical System Interoperability (NCTSI).

Dr. Don Brutzman is a retired submarine officer working in the Information Sciences Department, Undersea Warfare Academic Group and MOVES Institute at NPS.