An Intelligent Router for Improved Network Performance over Mobile Satellite Links

LCDR Gregory W. Johnson (gjohnson@exmail.uscga.edu)

Cadet 1/c Robin Kawamoto

Cadet 1/c Dave Bauer

USCG Academy, New London, CT

Abstract

The United States Coast Guard is nearing the end of a 5-year migration to Windows NT-based computers connected Coast Guard-wide via a private Transmission Control Protocol / Internet Protocol (TCP/IP) communications network. This network is known as the Coast Guard Data Network Plus (CGDN+). Current Coast Guard messaging and information access systems are being rewritten to run on the CGDN+ using web browser and e-mail applications as the front ends. As the Coast Guard moves into the 21st Century it is becoming increasingly important for the units in the field (the ships and the aircraft) to be able to access the information stored on this network in near real-time. Unfortunately, at the present time they do not have reliable dedicated access to this network. However, plans are in progress to link these mobile assets into the network using commercial mobile satellite systems. The systems available now, or in the near future (Inmarsat, AMSC, and Globalstar) provide low-speed and/or high-delay connections. Over the next year or so, all of the large cutters will be migrated from the current analog Inmarsat-A technology to the digital Inmarsat-B service, including the high-speed data option. However, even this is only at 64 kbps, and the latency in the geostationary satellite link is still high. Also, "online" time using these systems is very expensive. In order to improve performance and reduce the operating cost of using such a system, it is necessary to increase the efficiency of the applications used over the satellite link. Previous work at the Coast Guard Academy has focused on improving the performance of Windows NT-based TCP/IP stacks over very slow speed links (specifically the 2400 bps Inmarsat-M service). Current research has extended this to the Inmarsat-B high speed link (64 kbps). In addition, a system architecture to improve the performance on both of these links is being developed. This includes the use of an "intelligent router" which is being developed for the Coast Guard under an SBIR (Small Business Innovative Research Contract) by Architecture Technology Corp. (ATC) in Minneapolis, MN. The intelligent router allows the satellite links to be managed more efficiently by setting TCP/IP parameters to match the known characteristics of the link. To save on airtime, virtual connections are maintained between two intelligent routers while the physical link is taken down temporarily when there is no traffic. The intelligent routers also incorporate other performance enhancing features such as e-mail and web proxies.

Introduction

The U.S. Coast Guard is nearing the end of a 5-year migration of its computer systems from CTOS systems to Windows-NT based systems. In addition, the enterprise network has been completely replaced; upgrading from 9,600 and 56,000 bps X.25 lines to a high-speed TCP/IP (Transmission Control Protocol/Internet Protocol)-based WAN (Wide Area Network). Currently there is a large amount of informal traffic sent using e-mail on this network and there are Coast Guard specific applications such as LEIS II (Law Enforcement Information System) being implemented at this time that take advantage of online databases. In the future, all of the Coast Guard's messaging and information access systems will be based on this network as the Coast Guard moves forward into the information age. Most of the Coast Guard's messaging and information access systems are being transferred to e-mail and web browsing applications that run over this TCP/IP network.

At the end of this computer and network migration are the operational assets: the ships and aircraft of the Coast Guard. These assets in the past have not had access to the terrestrial network while deployed (sometimes not even while in port!). However, it is increasingly becoming an operational necessity for these assets to access the network while deployed. Cutters and aircraft on patrol need to be able to tap into this network to get up-to-the-minute operational intelligence, communicate with operational commanders, as well as conduct administrative and logistic arrangements. There are currently several ongoing projects within the Coast Guard to enable this network access; primarily using commercial satellite systems. Large cutters (greater than 210 ft) are being equipped with Inmarsat-B terminals with

The opinions or assertions contained herein are the private ones of the author and are not meant to be construed as official or reflecting the views of the Commandant or the Coast Guard at large.

the high-speed data option (64 kbps). Smaller cutters will be equipped with HF e-mail capability and Inmarsat Mini-M terminals (2.4 kbps). Some aircraft will also be equipped with HF e-mail capability. Unfortunately, time spent using a satellite link is very expensive; ranging from two to ten dollars per minute depending upon the system. In order to reduce the cost of using such a system it is necessary to increase the efficiency of the TCP/IP applications that are used over the satellite link.

TCP and IP protocols [1,2] were designed to be used on high-speed/low-delay terrestrial networks, not low-speed/high-delay satellite links. TCP/IP-based applications work over satellite links; however, they do not run efficiently. The primary applications to be used over this network are e-mail (messaging), which uses SMTP [3], POP3 [4], IMAP [5], or MAPI [6] protocols and web-based data access, which uses HTTP [7,8]. These protocols also were not optimized for low-speed/high-delay connections. Any improvements that can be made to optimize any or all of these protocols for the specific link characteristics will improve performance and reduce costs.

Approach

Our approach is to attack this problem on multiple levels of the OSI (Open Systems Interconnection) model. In the OSI model, network services are split between different layers; each layer providing services to the layer above and utilizing services from the layer beneath, while being responsible for certain functions. The commonly used TCP/IP suite of protocols maps to the OSI model as in Figure 1. The end goal is to improve the performance from the user perspective (at the application layer) and this can be done by making improvements at all layers. There are two main ways of increasing performance: reduce the amount of data being transmitted and make the transmission more efficient; we will make use of both of these methods. The amount of data can be reduced through compression and caching at multiple layers of the OSI model. Transmission is made more efficient by improving the efficiency of the protocols at various layers and maximizing link utilization. In many cases this efficiency is achieved by reducing the back-and-forth nature of many of the standard protocols which increases transmission time greatly when used over long latency links.

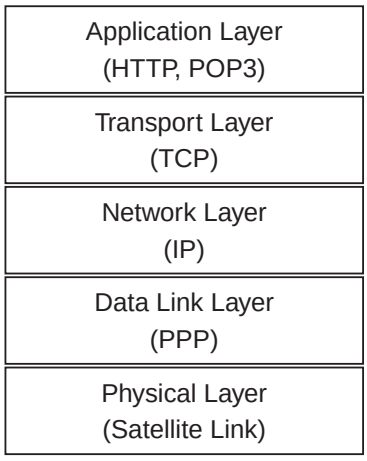

Figure 1 – TCP/IP Protocol Suite

The means to accomplish this is through intelligent routers and application layer proxies. In doing this we are making use of the fact that for this communications scenario, we control both ends of the link: all communications to/from a ship or aircraft can be forced to go through a common gateway. Also, we know a priori what type of communications link will be used and can make use of that knowledge to optimize the transmission over that link. These factors are not true for general transmission across the Internet, but can be used to our advantage here.

At the Application Layer data reductions can be made by using e-mail and web proxies that incorporate smart caching to eliminate the transmission of repetitive data. This is especially true for HTTP where HTTP GET and RESPONSE commands include 200-300 bytes of repetitive data about the capabilities of the browser and server. Data compression can also be done at this layer, using algorithms that achieve more compression by making use of knowledge of what type of information the data is. Also, some protocols can be implemented in a more efficient manner; for example SMTP. If both ends of the link are controlled by an e-mail proxy, SMTP commands can be batched and streamlined.

At the Transport and Network Layers compression can be implemented in the TCP and IP headers as well as the IP payload. The biggest area of improvement for these layers however, is in improving the efficiency of the transport protocol. This is where knowledge of the link characteristics can be used to "tune" the protocol for greatest efficiency: Maximum Transmission Unit (MTU) size, initial Round-Trip timeout (RTO) values, Selective Acknowledgements, Window Size, and Slow Start and other congestion control mechanisms are just a few of the areas that can be tuned to improve efficiency.

Compression can also be implemented in the lower layers as part of an intelligent router. In addition, the intelligent router manages the satellite link and improves link utilization. First of all, the intelligent router can maintain a virtual connection for the upper layers while physically tearing down the connection during idle periods to save money. The router also works in conjunction with the proxies; the proxies can accumulate and batch data to more efficiently use the connection. In addition, the proxies can communicate with the intelligent router their state so that the router can make better decisions on when to bring the link up and down. Furthermore, since one of the major sources of overhead on the links is the network connect time, the intelligent router uses a streamlined connect protocol to reduce the connect phase. Much of this capability is being developed under a Small Business Innovative Research (SBIR) contract, though some existing commercial products are also being evaluated.

Figure 2 –110 ft Patrol Boat

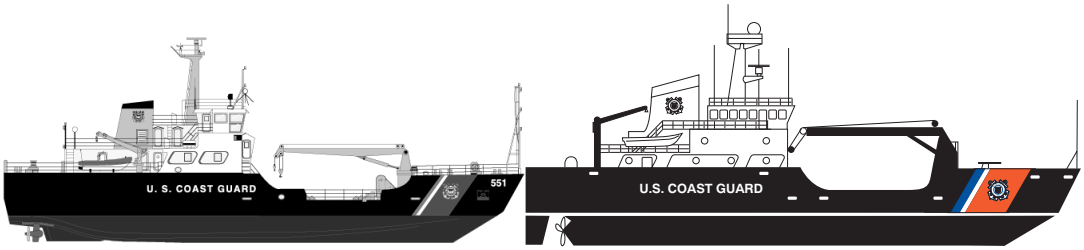

Figure 3 –175 and 225 ft Buoy Tenders

Satellite Links Used

The project breaks into two main subdivisions based upon the satellite system to be used since the link characteristics for the two systems are different. The two classes of cutter (small and large) will each be receiving different satellite systems; primarily due to space constraints on the small cutters. The small cutters: 87 and 110 ft Patrol Boats (Figure 2) and 175 and 225 ft buoy tenders (Figure 3) will each be receiving an Inmarsat Mini-M terminal. The large cutters: 210 and 270 ft Medium Endurance Cutters (Figure 4) and 378 ft High Endurance cutters (Figure 5) will be receiving the Inmarsat-B high-speed data terminals.

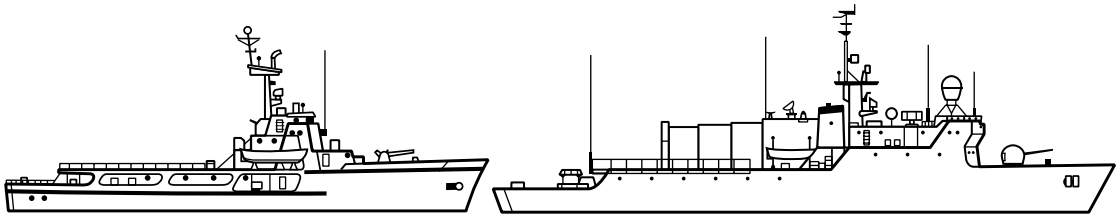

Figure 4 –Medium Endurance Cutters

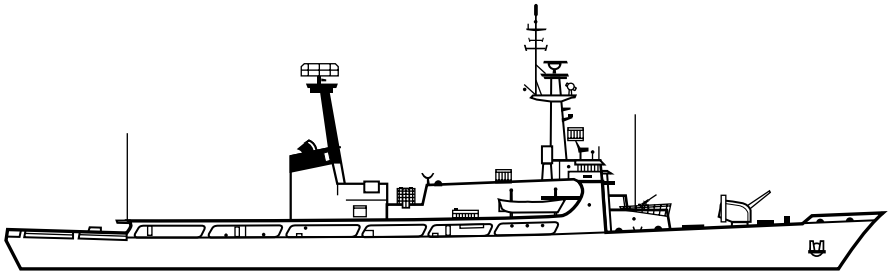

Figure 5 –High Endurance Cutter

Inmarsat Mini-M is one of the newer, digital, Inmarsat services that is provided using a small, satellite terminal that can be used for both voice and data at low speeds (2400 bps data, 4800 bps voice). The characteristics of this system are described in depth in [9]. The parameters of the system have been investigated and reported on in [10] as well as in our previous research [11].

Inmarsat B is also a digital service that can be used for both voice and data. It requires a larger satellite terminal, but can provide high-speed data at 64,000 bps. The characteristics of this system are also described in [9]. The operational characteristics of this system were reported in [11] and are summarized in the following table.

Table 1 --Inmarsat B Parameters

Data Rate	64,000 bps
Channel Delay	442 ms (one way)
Bit Error Rate	10^{-6} errors/bit
Set-up time	17 sec

Current Research

The Coast Guard Academy, in conjunction with the Coast Guard Research & Development Center (R&DC), has been conducting research into this problem for the past two years [11,12,13] with the goal of achieving the architecture outlined above. The focus has been on the small cutter scenario; in particular, last year the Inmarsat Mini-M link was investigated closely [11,13]. Improvements to both e-mail and web applications were found by changing the MTU (maximum transmission unit) from 1500 bytes to 500 and the initial round-trip timeout (RTO) from 3 seconds to 5. This reduced or eliminated the time-outs and unnecessary retransmissions on the link thus reducing total connect time. In addition, it was found that for a link of this speed, a single connection will totally use up the available bandwidth so that opening additional socket connections across the same physical link has an adverse impact on performance. Care must thus be taken to ensure that only a single logical connection is used. Also, POP3 and SMTP protocols were compared to MAPI (used between Microsoft e-mail clients and Exchange Server). It was found that MAPI required significantly more time to transfer the same amount of information. Thus for cost constrained links, standard Internet protocols such as SMTP and POP3 should be used in preference to MAPI.

Research has continued this year on the small cutter problem. In addition, the scope of the effort was expanded this year to include investigating the large cutter problem in conjunction with the Coast Guard Telecommunications and Information System Center (TISCOM). The satellite test bed was modified to allow for simultaneous testing of both scenarios. In addition, the SBIR contract that was initiated by the R&D Center in 1999 has progressed to Phase II this year so the results of this effort have been tested as well as some commercial products on both the large and small cutter scenarios.

SBIR

The SBIR contract with Architecture Technology Corporation will provide the Coast Guard with several capabilities needed for the architecture outlined above and described in [15] and [16]. All of these will be tested, and feedback from our testing used to improve the software. First to be developed is a software, dial-on-demand router. This will function similar to a typical router with the added

ability to more efficiently manage the link. The router will intelligently set up and tear down the link according to software programmable parameters. The router will utilize a priori information about the link characteristics as well as information from the proxies to more intelligently manage the connection. In addition, the router will utilize a new streamlined connect protocol (Optimized Link Setup Protocol or OLSP) to enable the connection between two routers to be established much more quickly than using standard Point-to-Point Protocol (PPP) [17]. The second piece of software will be an e-mail proxy. This software will allow messages to be batched and prioritized so that the link is only established when there is sufficient data to send so that the link will be fully used. The e-mail proxy will also attempt to utilize SMTP more efficiently as well as compressing the data to reduce the transmission time further. The third piece of software will be the web proxy. This software will provide intelligent caching of web traffic to reduce the repetitive data that is transmitted in HTTP GET and RESPONSE messages. The final piece of software is the tunable TCP/IP stack, which is a modified TCP/IP protocol stack to replace the standard Microsoft stack. The new stack will have parameters accessible to the system administrator so that it can be "tuned" for the link being used.

Mentat SkyX

Mentat Inc. [18] has developed a product to accelerate TCP/IP performance over satellite links called SkyX™ Gateway. This product is inserted into the communications path, 1 Gateway on either side of the satellite link, as shown in Figure 8. It breaks the TCP end-to-end connection and runs SkyX Protocol as the transport protocol between the two Gateways. This protocol has been optimized for satellite links using such mechanisms as an efficient acknowledgment algorithm that uses negative acknowledgements, unlimited window sizes, and a streamlined handshake algorithm. It also eliminates congestion avoidance mechanisms such as Slow Start and uses an explicit knowledge of the link speed to match the transmission rate to the available bandwidth. This protocol translation is only performed on TCP traffic; IP and IP-based protocols such as UDP, ICMP, and IPsec are not affected; they are passed through unmodified. The SkyX Gateway can also be used to compress all TCP traffic flowing through it. Additional detail can be found in [18] and [19].

The Mentat SkyX Gateway is being tested to see if it can improve performance on these low-speed satellite links (2.4 and 64 kbps). The manufacturer and others have tested it primarily on high-speed satellite links (1.544 Mbps and greater). If the system does improve performance, a determination can be made as to whether it should be used in conjunction with the SBIR efforts or in place of those efforts.

Small Cutter

The original small cutter test bed was modified slightly from that described in [11]. In order to test the new dial-on-demand router, computers running the router software had to be added to both the small cutter side and the shore side. In addition, the satellite delay simulator was replaced with an actual Mini-M unit to allow the dial-up performance to be tested (the simulator only modeled performance after the link connection was made). The e-mail proxy software is installed on the same machine as the one running the router software. See Figure 6.

To date, of the SBIR software, only the dial-on-demand router (DDR) has been delivered and tested. Since the proxies have not been delivered yet the only aspect of the DDR that could be tested was the improvement in link set-up time and since this has the most potential benefit to the small cutter problem it has been tested there only. In order to measure the improvement of using the DDR, the link was first tested in a standard dial-up configuration, using Windows NT RAS services and Dial-up Networking to establish a TCP/IP connection over PPP. This has been done before by the R&D Center, but was repeated just to verify the baseline performance.

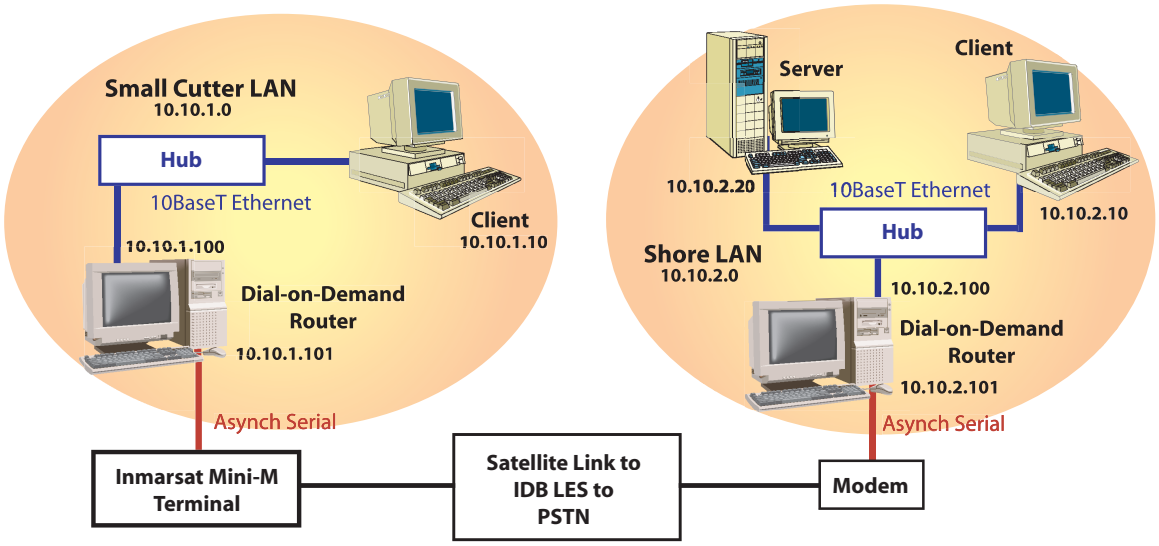

Figure 6 –Small Cutter Test bed

The link establishment times are reported in Table 2 . Timing measurements were done using a stopwatch and multiple trials were performed and averaged to arrive at the values in the table. The first three time marks for a total of 28.28 seconds are part of the Mini-M system and not subject to our optimization; the first time is for the terminal to get a channel and connect to the Land Earth Station (LES), the second time mark is when the terminal has established the connection on the PSTN, and the third time mark is when the modems have connected. The final time is for the network connection to be established (LCP and PPP negotiation).

Table 2 -- Inmarsat Mini-M Standard Network Connect Times: times are averages, and are the elapsed time from the previous mark

LES	Data	Modems Connected	Total Physical layer	On Network	Total
6.8	11.01	10.28	28.28	22.64	50.92

Once the network connection was established, the latency and throughput were retested and verified as well (Table 3 and Figure 7). Latency and Throughput measurements were done using WS_Ping [20]. The latency values were calculated using the same method as described in [10] and are similar to that measured then but longer than those measured in [11]. The connection was made using the Atlantic Ocean Region East (AORE) satellite and the IDB LES. The connection in [11] was through the COMSAT LES which perhaps explains the difference in the delay. In Figure 7 you can see the ramp-up in raw throughput as more packets are transmitted to a maximum value of just under 1600 bps which is comparable to that measured before.

Table 3 -- Mini-M Latency

Number of Bytes in Packet	56 bytes
Total Overhead (IP,ICMP,PPP)	31 bytes
Total Bits (Asynch, 8N1)	870 bits
Transmit time at 2400 bps (r/t)	725 ms
Ping Time (measured)	2720 ms
Latency (1 way)	**997.5 ms**

Once the dial-on-demand router (DDR) is running reliably, the same tests will be repeated using the dial-on-demand router to set up the connection. The initial parts of the connect phase (total physical layer) will be identical to the previous testing as this part is all done by the Mini-M unit. Once the Mini-M establishes the physical connection with the modem at the other end, then the DDR takes over to establish the network connection. Initial testing by ATC indicates that this phase is indeed much quicker than with standard PPP, on the order of several seconds vice 22 seconds. However, this has not

been able to be verified in our testing due to problems with the software. The dial-on-demand router is still a work in progress and at times unstable; as the bugs are worked out, the latency and throughput tests will be done again to see if the DDR has any impact on these results.

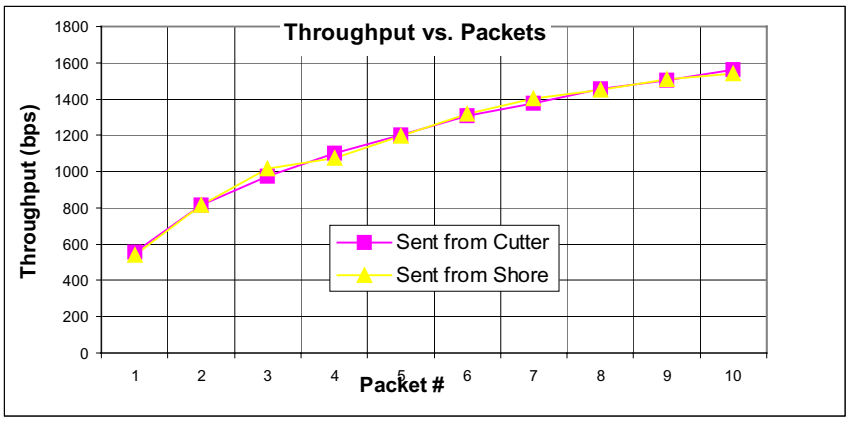

Figure 7 – Mini-M Throughput

Once the e-mail and web proxies are delivered, they will be integrated into the test bed and evaluated. The goal will be to see if they actually improve performance and how well they can integrate with the DDR to improve link management. In addition, once the testing is completed with the SkyX on the large cutter, it will be tested with the DDR and proxies on the small cutter test bed. All testing is done with the MTU set to 500 bytes and RTO set to 5 seconds, which were previously found to be better values than the defaults of 1500 bytes and 3 seconds.

Large Cutter

The large cutters operate in a slightly different manner than the small cutters. The large cutters all have their own Microsoft Exchange mail server onboard in order to provide e-mail services to multiple users and computers onboard the cutter via the shipboard Local Area Network (LAN). This is different than the small cutters which rely on a shore-side mail server. The test bed was configured as shown in Figure 8. On the large cutter side, there is a client terminal as well as the Exchange Server. This shipboard LAN connects to the outside world through a Cisco 2600 series router whether inport or underway. While underway, the connection from the Cisco router will be made through the Inmarsat-B high-speed data channel.

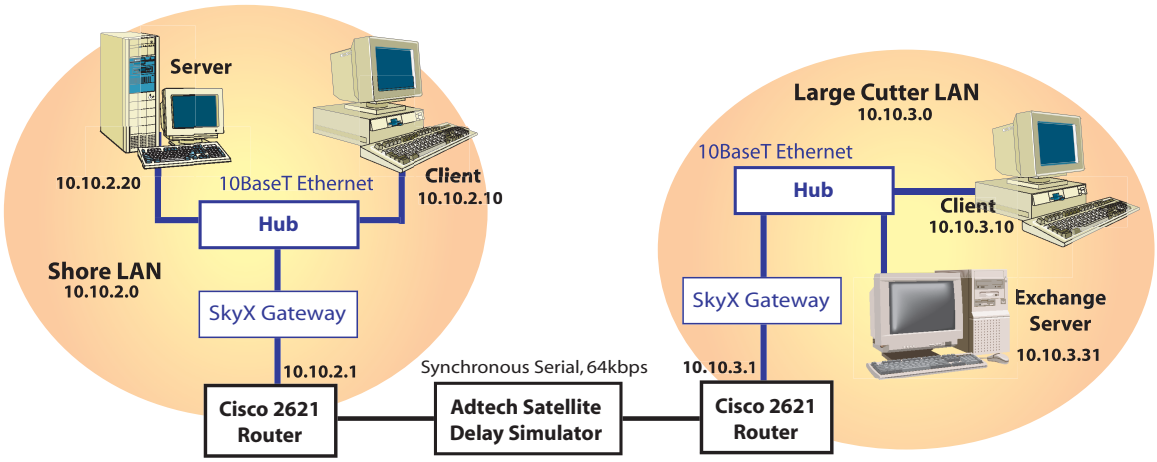

Figure 8 – Large Cutter testbed

The first step on the new large cutter test bed was to investigate the best MTU and RTO settings. Due to the speed of this link however, the default values are not as troublesome as they are on the small

cutter side. The speed of the link is 64,000 bits per second and is a synchronous serial connection (vice asynchronous for the Mini-M), which means that only 8 bits per character are used. This means that the raw speed of the link is 8000 bytes per second and in 200 ms (the suggested maximum time per packet in [21]) 1600 bytes can be transmitted. Thus a 1500 byte MTU is not a problem. Also, since the latency on this system is less (442 ms vice 998 ms) in addition to the faster transmission speed, the total roundtrip time for a 1500 byte packet to be sent and an acknowledgement to be received is only 442 ms + 187.5 ms (1500 byte packet) + 442 ms + 5 ms (40 byte ACK) = 1076.5 ms which is much less than the default 3000 ms RTO value. Initial testing with HTTP indicates that this is indeed true. Downloading a fairly large web page across the link with the default MTU and RTO values did not result in any retransmissions except for two that were probably the result of errors (the link is set for a BER of 10^{-6} which is roughly 1 error every 15 sec). This is illustrated in Figure 9, which graphs the sequence number as a function of time. In a plot such as this the sequence number should be always increasing so retransmissions are visible as dips in the line.

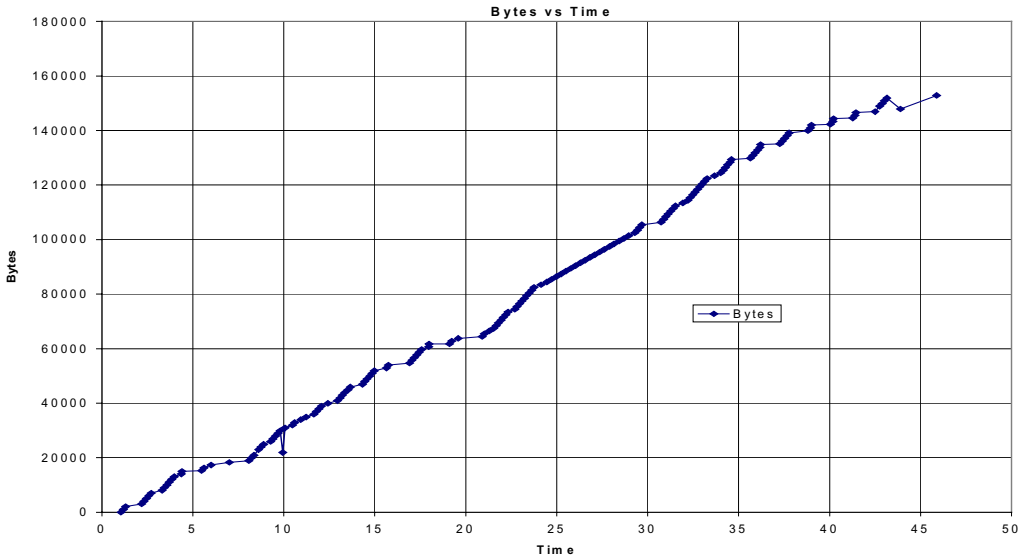

Figure 9 – Web test on Standard Inmarsat-B Connection

Testing similar to that described above was done for the large cutter to determine the throughput of this system (the latency was set at 442 ms one-way). The throughput was measured at both the IP layer (Figure 10) and the TCP layer (Figure 11).

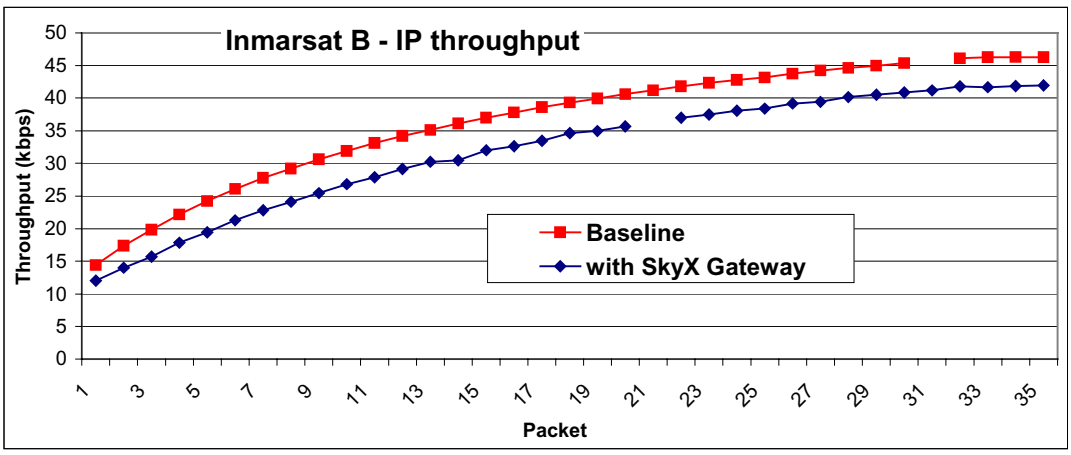

Figure 10 – Inmarsat B IP Throughput

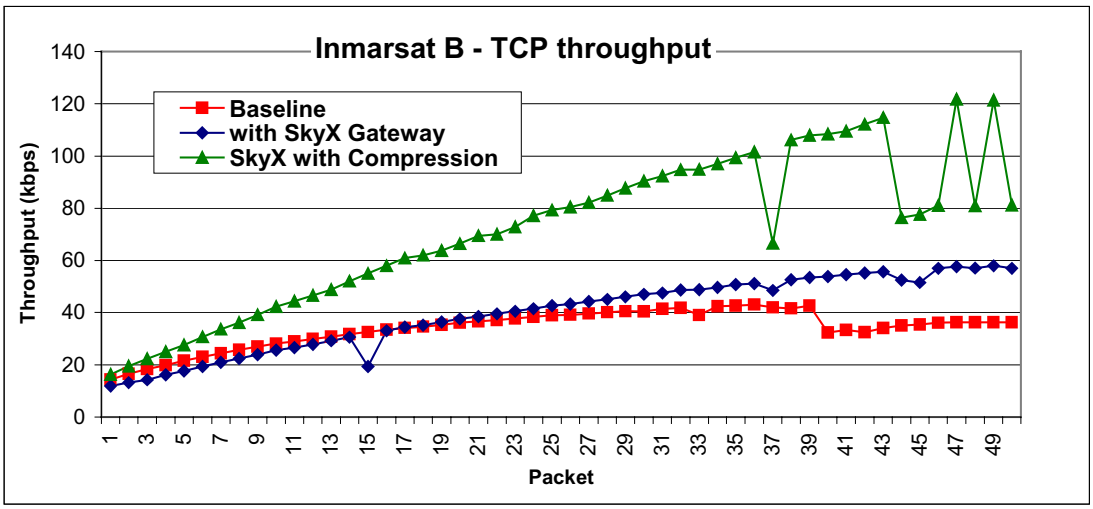

Figure 11 – Inmarsat B TCP Throughput

The Mentat SkyX Gateways were then installed into the link between the LAN and the routers on both sides as shown in Figure 8. The same tests were then rerun to determine if there was any improvement. As expected since the SkyX Gateways only work on TCP, at the IP layer there was no improvement. In fact, there was a slight degradation in performance. At the TCP layer however, there was a fairly substantial improvement. When the SkyX compression was enabled, there was an even greater improvement. Testing in the future will look at the application layer performance of e-mail and web protocols to see if they see any speed-up by using the SkyX Gateways.

Conclusions

Work has not progressed as rapidly as desired towards reaching the desired end-state architecture but some progress has been made. Initial assessments of both the large and small cutter systems have been completed to form a baseline for comparison. Progress continues to be made on some of the necessary software pieces; the intelligent router and the proxies. As these pieces of software are completed and improvements are made, they will continue to be evaluated to assess our progress towards the desired goal. The SkyX Gateway shows promise for the large cutter scenario, though it remains to be seen how much improvement there is at the application layer. It will also be tested on the slower Mini-M link. Additional commercial "link accelerators" may be looked at depending upon resources. Improvements in the protocols such as those being researched in the Performance Implications of Link Characteristics (PILC) Working Group of the IETF (Internet Engineering Task Force) [22,23] will continue to be monitored so that when the "tunable TCP/IP" stack is available they can be implemented if they appear to be worthwhile.

References

[1] Postel, J. "Internet Protocol." RFC 791, September 1981.

[2] Postel, J. "Transmission Control Protocol." RFC 793, September 1981.

[3] Postel. J. "Simple Mail Transfer Protocol." RFC 821, August 1982.

[4] Rose, M. "Post Office Protocol – Version 3." RFC 1939, May 1996.

[5] Crispin, M. "Internet Message Access Protocol - Version 4." RFC 1730, December 1994.

[6] Cox, Troy V. "Microsoft Messaging Application Program Interface (MAPI)." Microsoft White Paper, October 1993.

[7] Berners-Lee, T., Fielding, R., and Frystyk, H. "Hypertext Transfer Protocol Version 1.0." RFC 1945, May 1996.

[8] Fielding, R. et al. "Hypertext Transfer Protocol Version 1.1." RFC 2616, June 1999.

[9] Johnson, G., Turban, J., and Erickson, R. "Technology Assessment of Mobile Satellite System Alternatives." USCG R&D Center Report, April 1998.

[10] Johnson, G., and Fiorentino, P. "Inmarsat Mini-M Test Report." USCG R&D Center Report, May 1998.

[11] Johnson, G., and Reid, T. "The Performance of TCP/IP Applications Over Commercial Mobile Satellite Links." *Proceedings of the International Communication Satellite Systems Conference*. Oakland, CA, April 2000.

[12] Norton, R. "CG Long Range Communications Beyond 2000." USCG Academy Senior Design Project paper, May 1999.

[13] Reid, T. and Johnson, G. "Improving the Performance of Coast Guard TCP/IP Applications Over Mobile Satellite Links." Technical Report 06-00, Center for Advanced Studies, USCGA, May 2000.

[14] ATC web site at <http://www.atcorp.com/>.

[15] Ramanujan, R. "Optimized Transmission Control Protocol/Internet Protocol (TCP/IP) Stack." SBIR Phase I Final Report. Architecture Technology Corporation, March 17, 1999.

[16] Ramanujan, R., Takkella, S. and Thurber, K. "MobiWeb: An Approach for Mobile Web Access over Satellite Links." *Aerospace Conference Proceedings Volume I*. IEEE, March 2000.

[17] Simpson, W. "The Point-to-Point Protocol (PPP)." RFC 1661, July 1994.

[18] Mentat Inc. web site at <http://www.metat.com>.

[19] Mentat. *Mentat SkyX Gateway User's Manual 3rd Edition.* Mentat Inc. Los Angeles CA: March 2000.

[20] Ipswitch Inc. WS_Ping web site <http://www.ipswitch.com/Products/WS_Ping/>.

[21] Allman, M. and Falk, A. "On the Effective Evaluation of TCP." *ACM Computer Communication Review*. October 1999.

[22] Dawkins, S et al. "End-to-end Performance Implications of Slow Links." Internet-Draft. draft-ietf-pilc-slow-03.txt available at <http://www.ietf.org/html.charters/pilc-charter.html>.

[23] Border, J. et al. "Performance Enhancing Proxies." Internet-Draft. draft-ietf-pilc-pep-05.txt available at <http://www.ietf.org/html.charters/pilc-charter.html>.