HAIPE Compliant TCP Performance Enhancing Proxy for Bandwidth-On-Demand Environment

HAIPE Compliant TCP Performance Enhancing Proxy for Bandwidth-On-Demand Environment

F. Dell Kronewitter   Bo Ryu   Zhensheng Zhang   Liangping Ma
ArgonST Network Systems
6696 Mesa Ridge Road, Suite A
San Diego, CA 92121

***Abstract*— IP layer encryption introduces substantial challenges for Bandwidth on Demand satellite communication. Our solution, namely Broadband HAIPE-embeddable SATCOM Terminal (BHeST), utilizes novel network performance enhancement algorithms for high latency geosynchronous bandwidth-on-demand satellite links protected in the presence of High Assurance Internet Protocol Encryption (HAIPE). The problems experienced by TCP over geosynchronous satellites are well understood: while standard modems (on the BLACK side) employ TCP PEP's which have been shown to work well, the HAIPE encryption of TCP headers renders the onboard modem's PEP useless. This is attributed to the fact that under the bandwidth-on-demand environment, PEP must use traditional TCP mechanisms such as Van Jacobson to probe for the bandwidth of the link (which eliminates the usefulness of the PEP) or use the bandwidth signaling that does not violate RED/BLACK boundary. Modem vendors typically recommend disabling the PEP when a HAIPE device is used. By moving the PEP into the secure network (RED) and exploiting the bypass mechanisms allowed by the latest HAIPE standard, we have been able to regain the PEP's desired network enhancement that was lost due to HAIPE encryption. Our BHeST solution employs Direct Video Broadcast – Return Channel Service (DVB-RCS), an open standard chosen for Joint IP Modem (JIPM) initiative by Defense Information System Agency, as a means of providing bandwidth-on-demand satellite links as a placeholder for future Transformational Satcom (TSAT) terminals. Another issue we address is the estimation of current satellite bandwidth allocated to a remote terminal which is not readily available in DVB-RCS.**

***Index Terms*— TCP-PEP, Satellite communication, HAIPE, GIG, Joint IP Modem (JIPM), DVB-RCS.**

I. INTRODUCTION

TODAY'S threats present a wide array of asymmetric challenges to war fighting capability across a variety of missions for individual Services, Joint levels, and multi-national, coalition environments. These missions are ongoing around the world in support of military and civil organizations. The current information technology (IT) infrastructure no longer provides the best solution to meet the globally distributed information superiority needs of warfighters and sustainers within the increasingly important context of coalition operations. The Global Information Grid (GIG) will provide the Joint and Coalition war fighter with a single, end-to-end information system capability that includes a secure network environment, allowing users to access shared data and applications regardless of location, and is supported by a robust network/information-centric infrastructure.

Satellite communication is expected to be one of the key facilitators of GIG because of its ability to cover a very large area with a few assets and its ability to bypass terrestrial communication infrastructure whenever necessary. However, the strengths of satellite communication often come with the hefty cost of development (in terms of bandwidth per unit investment). These costs are due to relatively low user volume, low utilization attributed to static allocation and bursty traffic nature, and lack of widespread and affordable commercial deployment of broadband two-way terminals that the government could readily access.

The introduction of the Direct Video Broadcast – Return Channel Services (DVB-RCS) standard published and maintained by the European Telecommunications Standard Institute [3] and its growing adoption by satellite industry and Defense Information System Agency (DISA) [6][7][8][9] provide an array of hopes to tackle the aforementioned weaknesses. Because of its potential of supporting large user base with bandwidth-on-demand capability, the cost of terminals is becoming more affordable. In addition, the Multi-Frequency TDMA (MF-TDMA) frame structure, together with the dynamic slot allocation modes (rate based dynamic capacity, volume based dynamic capacity, etc.), provides an ample opportunity to achieve high uplink (reverse link) utilization for bursty multimedia traffic via statistical multiplexing. This will add to lower system cost because more users can be supported by a single system. It also supports Internet protocols via IP-over-ATM (AAL5 encapsulation). More recently, DVB-S2 standard [4] has been established as an expansion of the DVB-S standard (used by DVB-RCS for the downlink). DVB-S2 provides additional modulation and coding schemes that deliver higher link margins, and adds a "rain fade compensation" capability that allows the system to compensate for specific remote terminal link characteristics in unicast traffic directed to that terminal.

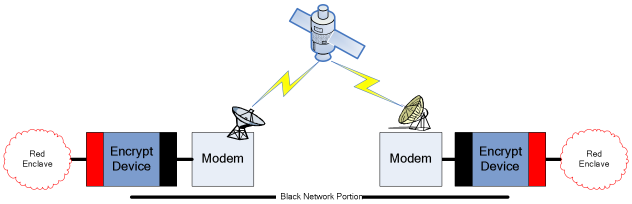

Fig. 1. Our communication scenario

978-1-4244-2677-5/08/$25.00 ©2008 IEEE

Despite these promising features and major push by DISA for wider adoption of DVB-based technologies, one of the key technical challenges still remain for DVB-RCS based systems (and for all GEO-based satellite systems in general) if they are to be seamlessly integrated into and widely deployed within the GIG infrastructure: *how to support TCP Performance Enhancing Proxy (PEP) in the presence of end-to-end HAIPE (High Assurance Internet Protocol Encryption), GIG's layer-3 security standard?* It is a well-documented problem that once HAIPE is enabled, the popular performance enhancing proxies (PEP) employed for boosting TCP performance over GEO satellite network such as TCP spoofing will not work because HAIPE's tunnel mode will hide TCP header information required by PEP agents. A technique called Multi-Level IPSEC was proposed to support both TCP PEP and HAIPE (which is based on IETF IPSEC standard) [14]. However, our past experience with both IETF IPSEC working group and HAIPE community indicates that they do not perceive the degraded performance of TCP over a certain type of network (satellite in this case) nor the problem or deficiency of HAIPE (or IPSEC) that it should address. In addition, they consider ML-IPSEC not as strong as original IPSEC in terms of end-to-end security (because it proposes trusted intermediate nodes to access TCP/IP header information). Consequently, it is unlikely that the HAIPE community will adopt this new change anytime soon, if at all, and therefore a new method of boosting TCP performance while fully supporting HAIPE and seamlessly interoperable with existing TCP-based applications must be developed, tested, and implemented.

While the primary focus of our research is a feasible solution towards the aforementioned problem, we also adopted a comprehensive system approach towards the development of a secure, efficient satellite communication system using readily available, commercially available communications equipment and software whenever possible by addressing the following two other challenges:

- Modern tactical, net-centric MANET communication requirements include Beyond-Line-of-Site needs. DVB-RCS systems must be closely integrated with terrestrial wireless technologies, such as 802.11 or 802.16, to serve as a backbone between out of reach subnets or back to the C3.
- Certification/accreditation authorities, including NSA, must approve the system architecture and packaging for deployment within the DoD information technology ecosystem.

The result is an inexpensive, compact, powerful system whose components are already well known by potential customers, certification authorities, and procurement personnel. A block diagram of our system appears in Figure 2.

II. SYSTEM ARCHITECTURE

The Department of Defense has mandated (TEMPEST) one-meter separation between a red processor and black equipment, wire lines, power lines, or conductors [1]. In order to comply with this mandate and at the same time to provide multiple options, we have chosen to provide our BHeST solution in two configurations, depending on whether HAIPE device is embedded within the BHeST box or external to it. Figure 2 illustrates these two options.

A major component of BHeST is the red side single board computer. Our prototype single board computer is Pentium M X86 based, and comes in the 3.5" PC/104 type of form factor. Features include a 10/100 BaseT Ethernet controller, IDE support for a hard drive, and two PCMCIA slots. We use the PCMCIA slot to host the embeddable HAIPE Talon card made by L-3. The HAIPE PCMCIA Talon in turn connects to a USB style adapter which then converts to Ethernet. Keys may be loaded via a proprietary dongle on the Talon Card which interfaces to the standard DS-101 Data Transfer Device key fill. Our single board computer runs Linux 2.6 kernel and hosts high-performance networking software, such as TCP-PEP, QoS queuing, and IPv4 to IPv6 translation. By introducing a powerful general purpose processor into the red environment we can offer better synchronization between red traffic and the black physical satellite modem environment.

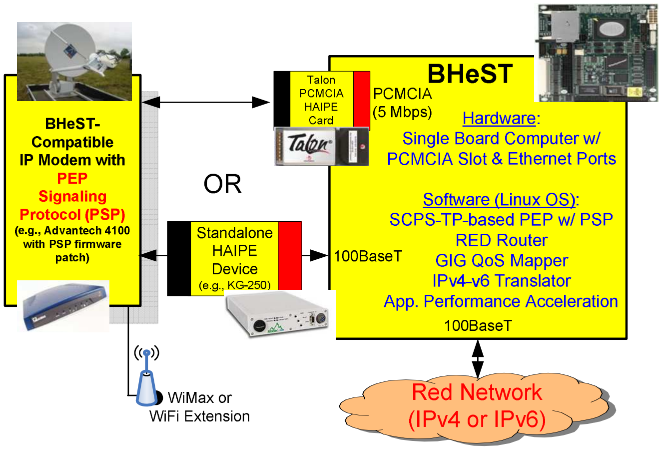

Fig. 2. Prototype integrated modem using COTS equipment.

III. HAIPE COMPLIANT TCP-PEP

The issues associated with TCP over high latency geosynchronous satellite links are well known. Due to the high bandwidth-delay product standard TCP slow start and congestion control algorithms behave very differently than over the terrestrial links for which TCP was designed.

Several solutions have been pursued to overcome the limitations of current transport protocols over satellite networks. These approaches generally fall in the categories of TCP interactions with lower layers, TCP enhancements, performance enhancing proxies (PEP), and totally novel architectures. Due to the ubiquitous nature of TCP our approach has adopted the PEP strategy which breaks the end to end integrity of the TCP connection, but accelerates the TCP performance of legacy hosts. The form of PEP we have adopted is called a split connection PEP where a transport protocol optimized for satellite links is used between the gateways which proxy for the standard TCP hosts. The Qualnet scenario which appears in Figure 3 demonstrates this split connection for the two-hop case where our TCP hosts are connected via hub or Network Channel Controller. Here, the connections between nodes 9 and 5 and nodes 6 and 8 are standard TCP, while the transport layer connection between nodes 5 and 6 is specially designed for satellite environments. In the case of this two-hop case the delay between hosts is typically 540ms and standard TCP protocol performance suffers dramatically.

As noted in the introduction, the splitting of TCP connections to perform acceleration requires the manipulation of TCP headers. The presence of HAIPE encrypts those headers so that they are not available to the PEP. While certain network architects have suggested that the PEP may be deployed in the host rather than at the gateway, extensive network research reported in [2] indicates that this solution is sub-optimal: “TCP gateways provide the maximum benefit when they are deployed right next to a satellite link, because they can use highly optimized protocols specifically designed for a SATCOM environment without affecting the rest of the network.” Other obvious advantages of incorporating the PEP in the SATCOM modem gateway are that legacy devices may enjoy the benefits of the PEP without upgrading software, PEP management is centralized, and the PEP has better access to information concerning the behavior of the problem satellite link whose performance is trying to be mitigated.

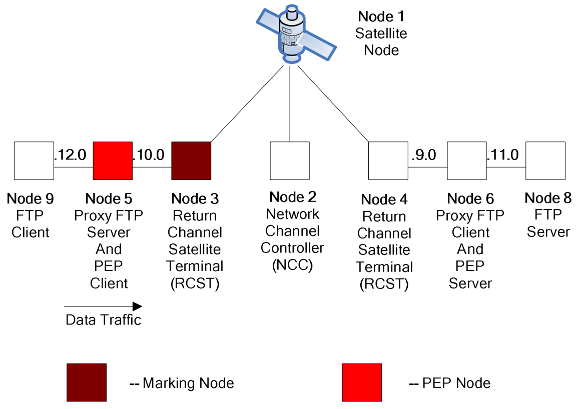

Fig. 3. PEP Signaling Protocol Qualnet simulation

While the solutions we propose are amenable to most PEP systems, we have chosen to use the Space Communication Protocol Standard-Transport Protocol (SCPS-TP) [12] to validate our approach due to its wide use and open availability [11]. Indeed SCPS-TP has exhibited performance equal to or surpassing commercial solutions. Reference [13] states that “Qualitative analysis of the protocols shows that SCPS would provide the greatest benefit to complex operational networks involving TCP-based communications, multiple network paths, and other redundancies. SCPS provides the flexibility that a complex redundant network would need. Additionally, SCPS-enabled gateways possess other advantages such as one-way acceleration and fail-over/fallback.”

A. PEP Signaling Protocol solution

While some PEPs, including a mode of SCPS-TP, operate over satellite links without explicit bandwidth information, they must resort to standard network “bandwidth probing” algorithms such as Van Jacobson’s original mechanisms or Reno to determine the network bandwidth currently available. More effective PEPs, including SCPS-TP open-loop rate control, make use of direct bandwidth information. When the PEP is placed in the satellite modem such bandwidth information is readily available especially with legacy fixed bandwidth satellite protocols such as SCPC. With modern statistically multiplexed protocols such as DVB-RCS bandwidth information is not typically directly available at the subscriber terminal, but can be inferred from lower layer signaling on a much faster time scale than transport layer methods.

Our scenario which includes a bandwidth on demand satellite protocol, DVB-RCS, and HAIPE encryption makes the problem especially challenging. As noted above the PEP must reside on the red side of the HAIPE encryptor while the SATCOM modem which converts the Ethernet data to satellite IF resides on the black side. Figure 4 demonstrates the position of our components: PEP, HAIPE, and modem.

Our solution is a bandwidth signaling method we call the PEP Signaling Protocol (PSP). We require a slight change to the software which runs on the SATCOM modem, by placing a PSP agent on board which is able to infer the available bandwidth and inform the PEP module of this via methods described below.

The new PSP protocol we have designed relies on marking two Explicit Congestion Notification (ECN) bits in encrypted packets to carry signaling information through the HAIPE encryptor. Modeled somewhat after Frame Relay’s FECNs and BECNs, IP ECN was first proposed in experimental draft RFC 2481 in 1999 and accepted as an Internet standard in 2001 as RFC 3168. This technique of signaling across cryptographic boundaries using the HAIPE ECN bypass feature has been exploited previously in INTSERV QoS research [15].

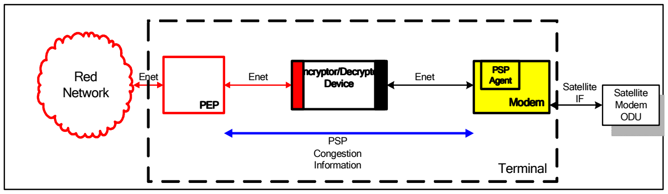

Fig.4. PEP Signaling Protocol architecture

The standard DVB-RCS traffic granularity rate is 2kbps or 32kbps depending on the scaling_factor bit defined in DVB-RCS satellite access control (SAC) messages. That is, capacity requests are scaled to 2 or 32 kbps. However, this seems too small for our purpose, and in PSP, we set the granularity to 50 kbps. Note that the most important signaling is rapid notification of large changes in available bandwidth, so our four signaling levels are roughly Additive and Multiplicative Increase and Decrease. With this in mind we have the following PSP black-to-red signaling format.

ECN bits	Description
00	Reduce current bandwidth 50% from current level

01	Reduce current bandwidth level 50kbps
10	Increase current bandwidth level 50kbps
11	Increase current bandwidth level 50% over current level or set current bandwidth to 512kbps

The PSP agent in the modem can then inform the PEP of the currently available bandwidth over the satellite link through numerous ECN signals. This signaling is at a primitive state and must be tuned with simulations and network traffic experiments.

As a practical matter we note that the PSP signaling is confined to the modem. The ECN bits on in-bound packets are cleared by the red side SBC. The ECN bits on outbound packets are cleared by the black side modem. By restricting our signaling to the ECN bits on outbound packets as well, we can support GIG DSCP preemptive class of service by applications in the red enclave which wish to make use of this GIG feature.

In order to evaluate the performance of the proposed PSP signaling scheme, we conducted extensive simulation. We have designed and implemented a Qualnet DVB-RCS environment to test our HAIPE compliant PSP. We have found that some of the features which make up SCPS-TP are standardized in RFC 1323 and have already been implemented in Qualnet TCP. We have had to add SNACK functionality and open-loop rate control as all of the Qualnet built in TCP variants use some form of rate estimation. Our Qualnet PSP simulation work contains one new "GATEWAY_PEP" application which serves as either PEP Server or PEP client. Our baseline Qualnet PSP simulation is pictured in Figure 4. The PEP client, Node 5, terminates the first (proxied) TCP connection and the PEP server, Node 6, initiates the final TCP connection. The protocol between these PEP nodes is SCPS-TP. The node 3 simulates the PSP agent in Figure 5 and marks packets which are going through over the DVB-RCS connection. The simulation includes a DVB-RCS satellite protocol implementation we have developed in compliance with Advantech's DVB-RCS implementation and the ETSI standard [3]. The resulting QualNet software is highly configurable allowing the system engineer to rapidly carry out experiments with many different network configurations and many different traffic models.In the simulation, we assumed two hosts can communicate with each other through a two-hop DVB-RCS architecture, shown in Figure 3. The PEP client, Node 5, terminates the first (proxied) TCP connection and the PEP server, Node 6, initiates the third or final TCP connection. The protocol between these PEP nodes is SCPS-TP with SNACK enabled. The node 3 simulates the PSP agent and marks packets which are going through over the DVB-RCS connection. The bandwidth of the satellite link is assumed to be 550kbps. The one-way propagation delay between node 3 and the satellite (or node 2, is about 540ms. The application is ftp to transmit a big file and the packet size is 50 bytes. There are 6 constant bit rates (CBR) background traffic sharing the same satellite link, each with rate of 480kbps. The TCP buffer size is set at 80kByte. The satellite link is shared using the multi frequency TDMA (MF-TDMA) scheme defined in the DVB-RCS standard. We compared the performance of the proposed PSP with a FIX PEP and the original TCP (without PEP). Here, FIX PEP refers to a SCPS-TP open loop rate control with a pre-determined fixed rate, while PSP is a SCPS-TP open loop rate control based on the available bandwidth of the satellite link. Based on the simulation results (see Figure 5), we observe that the improvement of PSP over FIX PEP is significant and could reach up to 100%. The improvement of PSP over the original TCP is even more (up to 500% for certain time periods).

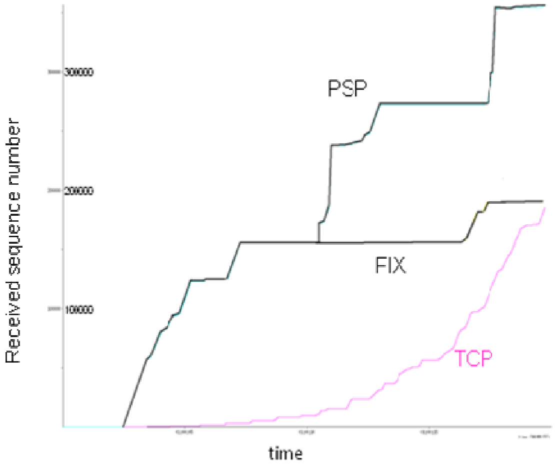

Fig 5. Simulation results of PSP with comparison against existing schemes.

B. *The Emulation Environment*

In order to further complement the simulation results, we have created an emulated satellite test bed to experiment with our PSP protocol. Our test bed consists of five computers as shown in Figure 6. The links around the outer perimeter of the figure represent the data path under test. The green links represent the satellite portion, the black links represent the hosted (secure) network, and the red segment demonstrates where the PSP signaling occurs in our emulation test bed. The computers at the upper left and lower left serve as traffic generators, sink and source respectively. For our initial experiments we have used the bandwidth measurement tools *iperf* and *scp*. The computer to the right in Figure 6 serves as the emulated two-hop satellite network. We have used the Linux based *netem* network emulation software to introduce a delay of 540ms and a .1% packet loss in both directions. This represents a two hop geosynchronous satellite. We also rely on netem to limit the bandwidth of the emulated satellite links to rates between 1Mbps and total blockage. The other two computers in Figure 6 host our modified SCPS-TP reference implementation.

The reference implementation SCPS-TP PEP does not support dynamic bandwidth satellite environments such as DVB-RCS. Rather the bandwidth of the satellite link is entered in a configuration file. This appears to be the first modification commercial PEP vendors make. Our dynamic bandwidth enhancement simply adjusts the rate used by the standard SCPS-TP code which has already implemented the token bucket filtering for open-loop rate control in a static satellite bandwidth environment. A similar situation occurs in systems which make use of external network compression equipment so that the data emitted from the SCPS-TP module does not reflect what actually ends up transmitted over the satellite link. In such a case vendors have implemented a feedback loop which informs the SCPS-TP module of the true traffic size.

Our PSP agent which in the final version runs on the modem as depicted in Figure 5 in our emulation scenario it runs on the netem computer in Figure 6 signaling the SCPS-TP. To emulate bandwidth on demand the emulation scenario modifies the rate of the emulated satellite link between 400kbps and 1Mbps. Changes occur every 500ms and happen in increments of 100kbps. Our PSP agent operates in two signaling modes: *FullBWInfo* which immediately signals the SCPS-TP the current rate, and *EstimatedBWInfo* which counts the packets which have traversed the link and provides one two bit signal for every packet. EstimatedBWInfo mode provides an accurate emulation for the fielded integrated HAIPE compliant DVB-RCS system.

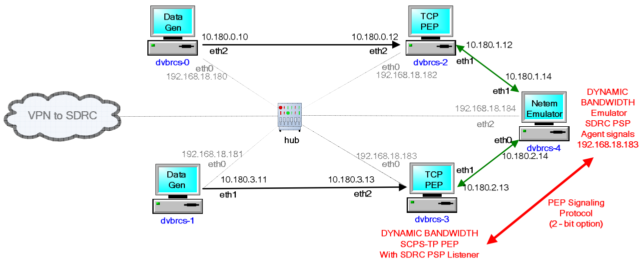

Fig. 6. Emulation test-bed setup

To validate our approach we performed experiments under three different PEP scenarios using the same traffic generation techniques under each PEP scenario. For each PEP scenario we used iperf to initiate a single TCP connection with client (source) dvbrcs-1 and server (sink) dvbrcs-0. The TCP connection attempted to transfer the maximum number of bytes over the emulated two-hop satellite link for 30 seconds. The PEP scenarios included "NoPEP" in which no PEP was used, "Estimated PEP" which used the PEP signaling protocol in EstimatedBWInfo mode, and "Actual PEP" which use the PEP signaling protocol in FullBWInfo mode. The "Actual PEP" scenario demonstrates how our dynamic bandwidth PEP would perform if the signaling channel were not restricted by HAIPE and the "Estimated PEP" demonstrates how our HAIPE compliant solution will perform. The "No PEP" scenario demonstrates the performance of a standard HAIPE and satellite modem with no PEP enabled.

Figure 7 shows the results of our experiment under the different PEP scenarios. Averaging throughput over the entire 30 second experiment gave us a "NoPEP" throughput of 215.9 Kbps, an "Estimated PEP" throughput of 704.45 Kbps, and an "Actual PEP" throughput of 711.4 Kbps. As expected the "Actual PEP" performs best and the "No PEP" is worst. Since the "Actual PEP" performance is close to the "Estimated PEP" performance we assert that our PSP is an effective method for improving performance in a HAIPE DVB-RCS environment.

We performed a similar experiment using scp (secure copy) as the traffic generation and measurement method. We transferred a 2.8 GB file from dvbrcs-1 to dvbrcs-0 under each of the PEP scenarios described above. The following table shows the results of our scp file transfer experiments in throughput and total time required. Again, the "Actual PEP" performs best and the "No PEP" is worst. Since the "Actual PEP" performance is similarly close to the "Estimated PEP" performance we have further evidence that our PSP is an effective method for improving performance in a HAIPE DVB-RCS environment. Since the upper layers of scp are rather verbose, the throughput is not equivalent to that exhibited by the raw iperf transfer. Still, this experiment demonstrates the effectiveness of our method on a real world TCP data transfer.

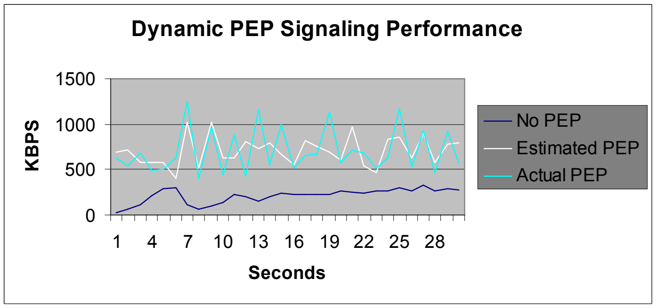

Fig. 7. TCP-PEP experimental results

Scenario	Throughput	Transfer Time
No HAIPE	46.41KBps (100%)	59.1sec (100%)
PSP HAIPE	38.32KBps (82%)	71.4sec (82%)
No PEP	29.7KBps (63%)	94.1sec (63%)

C. Bandwidth Estimation

DVB-RCS does not provide direct insight into the currently available bandwidth being provided by the NCC scheduler. The PSP agent in the modem must infer the current bandwidth from the currently allocated slots and the bandwidth requests coming from the modem to the network control center (NCC). The bandwidth requests coming from the modem could result in all of the requested bandwidth allocations during periods of low network usage or none of the requested bandwidth allocations during periods of high network usage. While the RCST has no direct insight into the NCC's Dynamic capacity allocation (DCA) scheme we suggest a method to infer future allocations based on (recent) past allocations. Mathematically we give the estimated bandwidth as:

$$EBW(i) = BW(i-1) + \frac{BG(i-\lambda)}{BR(i-\lambda)} * BR(i)$$

where $EBW(i)$ is the current bandwidth estimate, $BW(i)$ is the currently allocated bandwidth in the current TBTP, $BG(i-\lambda)$ is the bandwidth granted previously from that requested, $BR(i-\lambda)$ is the bandwidth requested previously, and $BR(i)$ is the currently requested bandwidth. An unburdened network will result in the granting of all requested bandwidth or $\frac{BG(i-\lambda)}{BR(i-\lambda)}=1$. In this case $EBW(i)=BW(i-1)+BR(i)$.

D. Quick start signaling

One difficult requirement for our approach is the existence of data flowing through the encryptor to use for signaling. Under load, in steady state, the TCP protocol guarantees the existence of TCP data flowing from black to red, the forward link, acknowledging the TCP data flowing from red to black, the reverse link. At the beginning of a large TCP data transfer to the NCC however no data is guaranteed to occupy the forward link to leverage for signaling. In a system with a non-zero Free Capacity Access (FCA) satellite link we can solve this problem with no reverse link bandwidth penalty by sending a periodic ping packet from the PEP over the link in the absence of TCP data. While the latency associated with each ping packet is very large over the satellite link (on the order of 550 ms), this is irrelevant since we require only the periodic existence of data to use for signaling. We expect to be able to provide 20 pings per second offering a signaling opportunity every 50 ms. This brings the PEP signaling latency to a level along the same magnitude as inter-frame DVB-RCS time which is typically 26.5ms for practical implementations.

The minimum size for a HAIPE tunnel encoded ICMP echo-request packet over the link is 88 bytes. (while 28 for the unencoded ICMP over IP echo-request) The 88 bytes translates to 2 ATM cells or 1 MPEG frame, typically a single DVB-RCS burst in either case. At the 20 packet per second rate this translates to 14.1 kbps which should be the minimum rate allocated to a RCST in standard DVB-RCS satellite network. As counterintuitive as this may be, the addition of these ICMP echo-request traffics results in no penalty on the network except a minimal transient penalty at the beginning of unsolicited TCP transfers over the forward (high-bandwidth) link. It is difficult to imagine such transfers since we would expect any traffic over the forward link to be solicited by the remote TCP node served by the RCST.

Studies have indicated that satellite users' experience is significantly improved by keeping a minimum non-zero bandwidth allocation for each user terminal. This allows small application layer request packets to traverse the satellite link without any delay associated with requesting a trivial amount of bandwidth. This fits in well with our quick start PSP signaling echo flow which requires this trivial amount of reverse link bandwidth (Figure 8). In the DVB-RCS QoS guide [10], ETSI makes a similar Free Capacity Assignment recommendation. Of course if the NCC MAC Scheduler is under extremely heavy load and no slots are available on the return link the PSP agent in the modem would signal the PEP and the PEP would refrain from sending the echo requests.

IV. WIRELESS EXTENSION

Recent tactical military communications approaches to NETCENTRIC warfare have focused almost exclusively on the development of MANET. Recent feedback from tactical military environments has shown that in order to maintain situation awareness for soldiers in the field it has become necessary to integrate beyond line-of-sight (BLoS) capabilities with line-of-sight (LoS) to enable effective information flow across the network. Recent experiments, such as those conducted under the DARPA FCS Communications program (the Network Centric Radio Systems instantiation), have effectively integrated satellite communications with MANET to maintain connectivity across the network. When nodes were beyond LoS from one another in the FCS Communications experiments the SATCOM system was used to provide connectivity. Given the incorporation of BLoS into military networks we can adjust the network paradigm to be inclusive and can change the basic network approach to utilize hybrid networking. Our integrated SATCOM modem provides for these types of deployment by offering an integrated (black) wireless extension which bypasses the HAIPE encryption as shown in Figure 2. IP Traffic entering the black port goes directly over the satellite link without incurring the overhead of HAIPE. We expect that the tactical MANET will not employ legacy TCP stacks and so currently have not investigated the coexistence of black and red PEPs. The PSP agent will however include the black traffic in its satellite bandwidth estimation algorithm. Our approach incorporates dynamic use of time, frequency and space for maximizing network performance and reliability using minimal overhead.

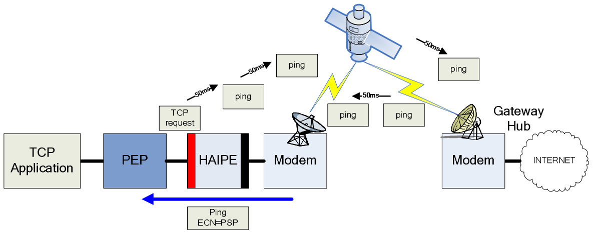

Fig. 8. Echo signaling in the absence of TCP data

V. QUALITY OF SERVICE

The prioritization and preemption of network traffic is a central feature in GIG architecture. The GIG NetCentric Implementation Document [5] states that traffic must be handled according to the priority associated with its DSCP value. Since our network device services a potential bottleneck link to the high speed backbone, differentiated traffic handling becomes especially important.

As mentioned earlier, the standard strategy we use is to keep the data in the unencrypted state as long as possible. This is a similar problem faced by the Linux enthusiasts accessing the

Internet over DSL lines. The Linux kernel contains built in queuing support including token bucket filters, classifiers, and conditioners as described in RFC 2475. We have modified the standard public script which implements such P&P, "wondershaper", to be compliant with GIG NCID directives and turn any modern Linux router into a GIG compliant traffic classification node. We call our new script "GIG-shaper".

VI. IP VERSION TRANSLATION

Despite the DoD-mandated transition to IPv6, it is expected that IPv4-only applications will continue to exist for many years due to practical aspects of transition. Also, those legacy systems that will be deprecated in the next five years are not likely to migrate to IPv6 at all. At the same time, the core network infrastructure of the DoD is being migrated to IPv6 because robust IPv6 support is a necessary part of the Net-Centric Operations and Warfare (NCOW) strategy. However, legacy IPv4 applications, hosts, and network infrastructure must all be included in the NCOW strategy and make a seamless use of the IPv6-enabled GIG.

Consequently, we are developing a robust, cost-effective, and versatile translation service (based on NAT/PT) within the BHeST to enable those legacy IPv4 applications to fully enjoy such IPv6 features as expanded addressing space, enhanced security, anycast, auto-configuration, and improved mobile IP support. By placing the translation service on the red SBC we are able to support the requirements, including the rare application level translation requirements, of IPv4 hosts accessing IPv6 infrastructure across the DVB-RCS link. Our solution leverages an extensive comparative study conducted under our OT-TES Communications Upgrade Army contract.

Furthermore, our approach eliminates the need for dual stack or tunneling middleware to be installed on each host offering "bump-in-the-stack" solutions, which allow IPv4 hosts to tunnel data over the IPv6 network, not really integrating the host into the IPv6 network. Administration is centralized at the BHeST, rather than distributed throughout the red network hosts, reducing O&M costs. IPv4 accelerated ASICs where the protocol lies in the hardware, small devices which can run only a single stack of IPv4 either because of memory constraints or cost constraints, or an IPv4 server which cannot be upgraded all may be supported in an IPv6 GIG with the BHeST.

VII. CONCLUSION

In this paper, we have presented a comprehensive description of BHeST system design which includes many novel and practical features that are essential for a GIG-compliant system. The biggest contribution of this solution is its ability to maintain TCP PEP function in the presence of HAIPE and bandwidth-on-demand satcom links. Preliminary results based on our detailed simulation testbed of DVB-RCS and PSP show promising trends. Further research on prototype design, demonstration and performance evaluation will further prove the usefulness and effectiveness of BHeST for GIG.

REFERENCES

[1] NSTISS TEMPEST, "Red/Black Installation Guide", December 12, 1995.

[2] T. Henderson, M. Duke, J. Megan, and S. Yi, "TCP Performance Enhancements for Tactical Networks", Boeing Tech. Rep., Sept. 2005.

[3] ETSI EN 301 790, "Digital Video Broadcasting (DVB); Interaction channel for satellite distribution systems," v1.3.1, European Telecommunication Standards Institute, March 2003.

[4] ETSI EN 302 307, Digital Video Broadcasting (DVB); Second Generation Framing Structure, Channel Coding and Modulation Systems for Broadcasting, Interactive Services, news Gathering and Other Broadband Satellite Applications, v 1.1.1, European Telecommunications Standards Institute, (2004-2006).

[5] DISA, "Global Information GridNet-Centric Implementation Document: Quality of Service (T300)", December 2005.

[6] B. Bennet, K. Quock, E. Summers, M. DiFrancisco, "DVB-S2 Technology Development for DoD IP Satcom", MILCOM 2006, Washington D.C., 2006.

[7] B. Bennett, K. Quock, M. Skowrunski, and M. Difrancisco, "Digital Video Broadcast Return Channel Satellite (DVB-RCS) Architectures and Applications for the Department of Defense (DoD)," MILCOM 2005, Atlantic City, NJ, Oct. 2005.

[8] B. Bennet, D. Hannan, J. Marshall, R. Gibbons, "WGS Capacity Using The DOD Joint IP Modem (DVB-S2, RCS)", MILCOM 2006, Washington D.C., 2006.

[9] DoD, "CIO John Grimes: DoD Policy for Transmission of Internet Protocol over DoD-Leased and DoD-Owned Transponded Satellite Communication Systems", Feb. 2006.

[10] ETSI, "Draft ETSI TS 102 462: Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM) Services and Architectures: QoS Functional Architecture," March 2006.

[11] http://www.openchannelsoftware.com/projects/SCPS

[12] "Space Communications Protocol Specification (SCPS) – Transport Protocol (SCPS-TP)", CCSDS 714.0-B-1, Blue Book, May 1999.

[13] J. Doffoh, R. Mereish, and M. Puckett, "Analysis and Comparisons of Acceleration Protocols for TCP Over Satellite", MILCOM 2005, Atlantic City, NJ, Oct. 2005.

[14] Y. Zhang, "A Multi-Layer IP Security Protocol for TCP Performance Enhancement in Wireless Networks," IEEE JSAC 22(4):767-776, May 2004.

[15] B. Bennett, B. Jaspal, and F. Yao, "The Next Generation IP SATCOM: Exploration of IPv6 Technology in Global Broadcast Service Technology Refresh," MILCOM 2005.

[16] O. Brewer, A. Ayyagari, and M. Foster, "Integrated Services Provisioning Across Cryptographic Boundaries", MILCOM 2006, Washington D.C., 2006.