A PERFORMANCE EVALUATION OF TRANSPORT MECHANISMS IN HYBRID NETWORKS

N. Schult, R Wade, G. Comparetto, M. Mirhakkak
The MITRE Corporation
McLean, VA

ABSTRACT

This paper evaluates the performance of several alternative reliable unicast transport mechanisms in a hybrid network. Options investigated include end-to-end TCP (different flavors), end-to-end Space Communications Protocol Standards-Transport Protocol (SCPS-TP), and Performance Enhancing Proxies (PEPs) (also called Transport Layer Proxies). Our approach is to analyze these options in a specific scenario using Modeling and Simulation (M&S). We describe this scenario and the corresponding OPNET Network Model, our experiment plan, and the results obtained. Finally, we identify several areas for further analyses.

INTRODUCTION

The DoD is evolving towards end-to-end, seamless, network-centric communications using multiple networks with very different characteristics (e.g., wireless and wired links, fixed and mobile network components). Such a network is typically described as "heterogeneous" or "hybrid". One example is a network composed of a high-speed wired backbone, tactical radio networks and satellite communication (SATCOM) links connecting the radio networks to the backbone.

Different types of information will need to traverse these networks, with varying requirements for reliability and timeliness. One type of information is unicast data that needs to be sent reliably; the focus of this paper is on how to support this requirement of (unicast) data that needs to be sent reliably, end-to-end, in a heterogeneous IP network.

Reliable end-to-end communication in an IP network implies the use of a transport layer protocol. TCP is typically used in "traditional" (wired) networks, but there are concerns about its effectiveness in hybrid networks.

TCP has been the predominant transport protocol for reliable end-to-end delivery of data and has evolved for use in the "wired" network, where links are relatively error-free and packet loss is usually due to congestion. Different "flavors" of TCP, including Reno [7, 8], New Reno [5], Vegas [14], and Westwood [15], each contain slightly different congestion control and congestion avoidance mechanisms. These flavors can also differ by the assumed source of the packet loss.

Two predominant flavors of TCP in use in the Internet today include TCP Reno with Selective Acknowledgements (SACK) and TCP New Reno. Both of these flavors address the problem of multiple packets being lost from a transmission window (which previous flavors, e.g., Tahoe, did not).

Space Communications Protocol Standards (SCPS) – Transport Protocol (TP) defines a set of TCP options and behaviors (congestion control, rate control, assumed source of loss, ACK frequency, etc.) that can be used to extend and/or move the "domain' in which TCP performs well [2]. Thus SCPS-TP is well-suited to stressed communications paths characterized by long delays, limited bandwidth, asymmetric bandwidth, and high error rates. SCPS-TP is part of the SCPS suite of protocols developed to support communication with nodes in space. SCPS exists as ISO standards, the Consultative Committee for Space Data Systems (CCSDS) standards, and U.S. Military Standards.

PEPs are in-network devices that attempt to improve end-to-end throughput, generally by interacting with the transport protocol in some way. Split-connection transport layer PEPs, or Transport Layer proxies, take a transport protocol connection and divide it into multiple connections. PEPs are transparent to the end-users or applications [4]. The idea with using PEPs is to have different transport connections, each tuned to the environment it is operating in, such as wired, tactical radio, or SATCOM. These different connections can use TCP and/or SCPS-TP. More information about PEPs can be found in [4].

Figure 1 illustrates these alternative transport mechanisms over a sample hybrid network path composed of a fixed wired network, a satellite link, and a radio or wireless network.

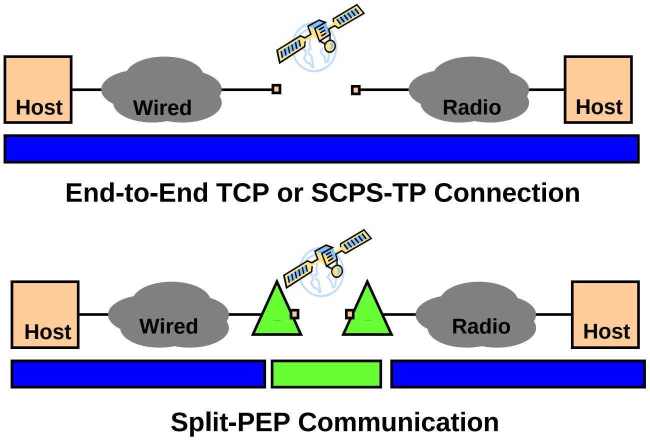

Figure 1: Examples of Transport Mechanisms

The top example in Figure 1 assumes the “traditional” approach for providing reliable end-to-end communication – a single connection, typically TCP, between the source and destination hosts. A SCPS-TP connection can also be used to support end-to-end communication.

The bottom example in Figure 1 incorporates the use of split-connection PEPs, denoted by triangles. In this example, there is a TCP connection between each host and its nearest PEP, with the idea that each connection is tailored for the given network. Between the PEPs, another transport connection, e.g., SCPS-TP is utilized. As a result, the communication between the two hosts actually consists of three connections. Note that PEPs are most effective when they “bracket” a network like the SATCOM link above, and there is no alternative path. If there is an alternate path between the two hosts that data could flow, once a connection is established using the PEPs, packets that do not use the PEPs will be rejected by the destination.

For our study, we are interested in comparing the performance impact of using TCP-only with that of other alternative transport layer protocols and/or mechanisms, to support reliable end-to-end communications.

APPROACH

Our approach is to quantitatively evaluate these alternatives in a representative heterogeneous network under “realistic” conditions and traffic loads. The network utilized includes wired components, tactical radio components, and satellite assets. Some of the tactical radio components are mobile, and realistic links are assumed. The traffic is described by Threads and Information Exchange Requirements (IERs), which were derived from actual activities and traffic.

To assess the performance of the various approaches, we used the OPNET Modeler discrete-event simulation tool (which we will refer to as “OPNET”). OPNET is part of a M&S environment developed by MITRE called the End-to-end Modeling And Simulation Testbed (EMAST), to analyze issues in large-scale heterogeneous networks [1].

EMAST consists of three main components: a Model repository, Scenarios repository, and Engine. A goal of EMAST is Model and Scenario re-use, which is achieved with the repositories above. The Model repository contains models developed by the MITRE team and others, developed primarily as part of previous work. The Scenarios repository includes both scenarios that have been developed from previous work, as well as those developed for this project. By “scenarios” we mean the node laydown, the movement of the nodes over time, and network traffic generated by the nodes; this information is conveyed in a set of scenario description files. The Engine is MITRE’s Modeling and Simulation Environment (MSE), which includes OPNET, a Scenario Generation tool, and Parser. Using well-defined files, scenario description files created by the Scenario Generation tool, and models in the Model repository, the Parser generates a program that, when executed, builds an OPNET network model, which can then be evaluated using OPNET .

The scenario used for our assessment is the EMAST Proof of Concept (POC) Stryker Force [11]. The Stryker Force Brigade Thread served as the basis for developing this scenario, which utilized assets in both the strategic (wired) and tactical (wireless) networks. It also included unicast data IERs that need to be delivered reliably and that traverse over both heterogeneous (wired-SATCOM-wireless) links and wireless links. (For this study, these are the only IERs in the scenario identified as needing to be delivered reliably). A previous scenario, the DARPA Future Combat System – Communications (FCS-C) Program’s Demo 3 Boise Scenario (e.g., node laydown, mobility, and traffic) was used to represent the behavior/activity of tactical radio nodes in the network. Figure 2 illustrates the scenario utilized for our assessment of alternative transport mechanisms.

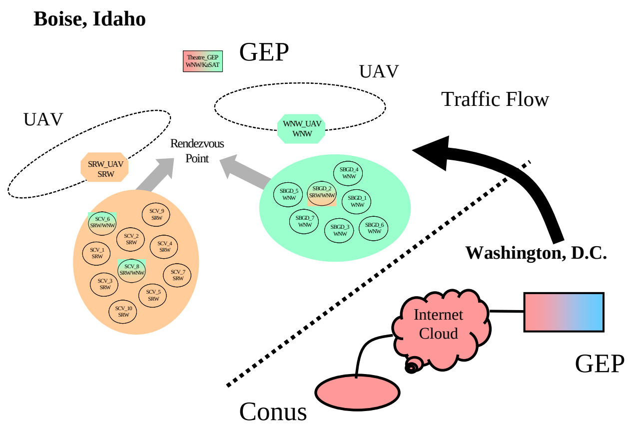

Figure 2: EMAST POC Stryker Force Scenario

In this figure, the left side is situated in Boise, Idaho and is based on the actual DARPA FCS-C Demo 3 Boise scenario. Two groups of nodes, each with at least one tactical radio, are moving towards a rendezvous point over the course of the scenario, with a mountain separating them from direct communication for much of the time. Two UAVs above the two groups provide the connectivity between them when the groups are not in direct contact.

There are two Ground Entry Points (GEPs) in this figure: one in Theatre (Boise) and one in Washington, D.C.

The right side of the figure includes fixed assets. The node labeled "Conus" is a host connected to the Internet Cloud, and is the initiator of and participant in the thread of interest. Note that communication from this node to a radio node in Boise will traverse wired, SATCOM, and radio network assets. This is the path we are particularly interested in evaluating in this study (i.e., a heterogeneous path), which is traversed 3 times, in three of the 10 steps that make up the Stryker Force Brigade Thread. In each of these three steps, unicast data needs to be sent reliably, and packet sizes are assumed to be either 8 Kilobyes (KB) or 100 bytes. The duration of the scenario is 7800 seconds.

Using EMAST, an OPNET network model of this scenario was generated, with each platform represented by an OPNET Mobile Subnet Model. Each platform consisted of at least a Host and Router. Each GEP platform also includes a KA SATCOM (KASAT) terminal. The Host, Router, and KASAT terminals are each represented by an OPNET Node Model. A Host generates/receives IERs and includes a protocol stack, with UDP and either TCP or SCPS-TP incorporated as transport protocol options. A router includes two or more interfaces, e.g., Point to Point Protocol (PPP), Solder Radio Waveform (SRW), Wideband Networking Waveform (WNW), Ethernet (for the KASAT terminal). The GEP router model also includes TCP, SCPS-TP, and a PEP.

The IP-based routing protocols used include:

- In the wired network and KASAT network, OSPF v2
- In the WNW network, OSPF/ROSPF
- In the SRW network, DS Routing (ITT)

For multicast routing, PIM-SM was used.

Finally, the satellite link had a 281.09 millisecond propagation delay and a bit error rate (BER) of 10^{-6}.

EXPERIMENTAL PLAN

For our investigation, thirteen test cases were defined, which are distinguished by the transport mechanism used by the unicast data. These cases included UDP end-to-end (as a point of reference), TCP end-to-end (two variations), SCPS-TP end-to-end (5 variations), and split-connection PEPs with TCP and SCPS-TP (5 variations). Split-connection PEPs were only used on the wired-SATCOM-radio connection.

For end-to-end TCP, two cases were defined: TCP Reno with SACK and TCP New Reno. Additionally, the Timestamp (TS) and Window Scaling (WS) options were enabled. A window size of 250,000 bytes was used.

For SCPS-TP end-to-end, five variations/cases were investigated, which included Selective Negative Acknowledgements (SNACK), rate control with SNACK, SACK, rate control with SACK, and TCP-Vegas with SNACK. These also had the TS and WS options enabled. A window size of 250,000 bytes was used.

Finally, five split-connection PEP cases were investigated, with each using two PEPs bracketing the SATCOM link. Between the PEPs, the five variations of SCPS-TP mentioned above were studied. For the connections between the host and PEP, TCP New Reno was used. For these PEP cases, the same configuration of

TCP was used for both the wired and wireless paths, and the TS and WS options were disabled.

Metrics collected included the following:

- For unicast data: end-to-end delay, IER completion probability, and goodput. IER completion probability is the number of IERs completed divided by the number initiated. We compute Goodput as the amount of "good" data received (e.g., excluding duplicate data) divided by its latency. These metrics were collected for wired-SATCOM-radio connections, for radio connections, and summarized over all connections.
- For all network traffic, POC thread, POC thread steps (i.e., IERs): Completion probability.

Each test case was executed or replicated 8 times, each time with a different (prime-number) random number seed.

RESULTS

Highlights of results obtained are found below, where the focus is on successful completions. The results presented are the averages of the output metrics from each replication.

Figure 3 illustrates overall unicast data Sent versus Received, in bytes, for four representative cases: UDP, TCP with New Reno end-to-end, SCPS with SNACK end-to-end, and a split-PEP Configuration that used TCP-New Reno for the wired and radio networks and SCPS with SNACK across the satellite link (referred to as Config5). When comparing these cases, both a higher number of bytes sent and received are considered "better" (explained further below).

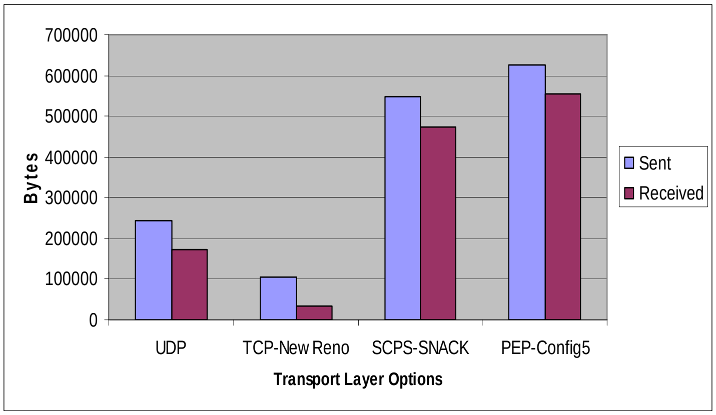

Figure 3: Average Unicast Data Sent vs. Received

In general, the cases that used TCP end-to-end over the heterogeneous link performed the worst, while those that used PEPs performed the best, for the cases considered. In Figure 3, more data was sent in the SCPS-TP and PEP cases, because more steps of the POC Thread were completed. When investigated TCP's poor performance, we found that the queues at the router located at the Boise Ground Entry Point (GEP) and radio network were very congested, resulting in large queueing delays. As a result, the TCP connections were not getting established. (It should be noted that there is no active queue management or QoS mechanisms active in the model.)

We also ran a side-case with TCP New Reno end-to-end, but with larger time out values, i.e. the same ones used for SCPS-TP. When the Initial, Maximum, and Minimum Retransmission TimeOut (RTO) values were set to 6, 240, and 0.5 seconds, respectively, some data managed to reach its destination, implying improved TCP performance

The differences in performance between the SCPS-SNACK and PEP-Config5 cases in Figure 3 reflect the "general" observation that the split-PEP configurations exhibited better performance, but not significantly better, in our scenario. However, this depended on the tuning of SCPS-TP, either end-to-end or between the PEPs. The two cases depicted in Figure 3, SCPS-SNACK end-to-end and PEP-Config5, were two of the better performing cases in our scenario. The error rate assumed for the SATCOM channel was 10^{-6}; we anticipate that with a higher BER, the performance differences would be more pronounced.

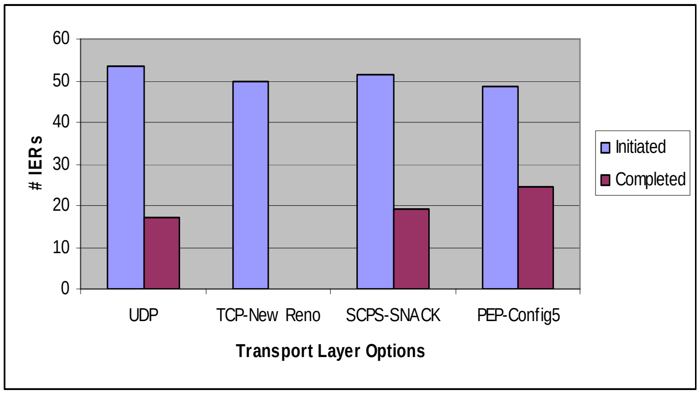

Figure 4: Avg Number of IERs Sent vs. Received, Step 0

For all cases, we found that the first step of the thread had the lowest probability of completion. Figure 4 illustrates the number of IERs sent versus received for the first step of the 10-step thread – Step 0 – for several representative options.

This is also illustrated another way, for one test case, PEP Config5, in Figure 5. This is indicative for all SCPS-TP and split-PEP cases we investigated, with PEP-Config5 having the highest IER completion probability values.

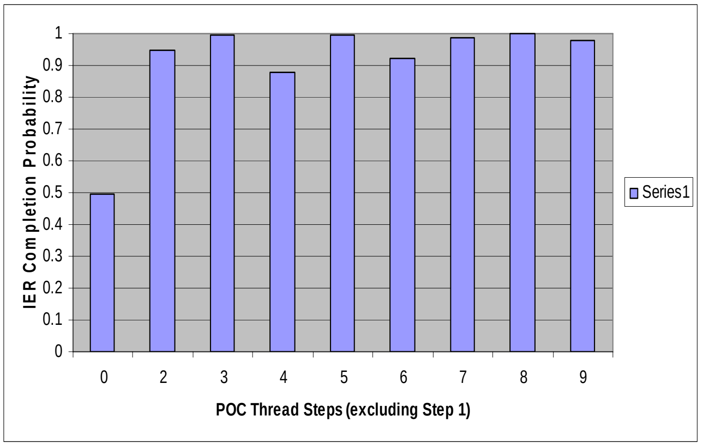

Figure 5: Avg IER Completion Rate per Thread Step for PEP-Config5 (Unicast only)

In summary, we found with our scenario that the split-connection PEP cases, followed by the SCPS-TP end-to-end cases, exhibited higher completion performance than the TCP cases investigated. However, this depended on the tuning of SCPS-TP, either end-to-end or between the PEPs. It should be noted that to effectively run TCP or SCPS-TP end-to-end in this hybrid network, their parameters need to be changed at each local host to a non-standard set of values (which is primarily needed to communicate across the hybrid network, not locally). When using the split-PEP, this step is not necessary, since each host can use the default settings, and only the PEPs need to be managed and tuned appropriately.

AREAS FOR FURTHER ANALYSES

There are a number of areas of further analysis, some of which are mentioned here.

Obviously, this evaluation considered one scenario, with its given node laydown, mobility, and traffic. It would be useful to evaluate these alternative transport mechanisms in at least a couple of dissimilar scenarios, to see if similar results are obtained.

Another area for analysis is the study of alternative transport mechanisms when Quality of Service mechanisms are enabled. The problems observed with congested queues could be improved using QoS and associated mechanisms, along with active queue management.

A third area to investigate is the use of alternative transport mechanisms to support reliable unicast data communications in wireless tactical networks, which operate in a broadcast environment with higher BERs and lower bandwidth. This is a research area, and with the use of PEPs, the connection over a tactical radio network could be better tuned for that environment.

A side issue has to do with the thread analyzed in the POC Scenario and, in general, data management. The thread consisted of 10 steps and traversed the heterogeneous path (between Washington, DC and Boise) three times. It seems that there should be better strategies in having the information needed by those in Boise located closer, and these should be given some thought.

Finally, in large heterogeneous networks, the applicability of Delay/Disruption Tolerant Network (DTN) concepts [12] should be further analyzed. DTN includes a store and forward concept to support networks that can become disconnected (e.g., due to a consistently poor performing, very congested network in a packet's path).

REFERENCES

[1] Comparetto, G., Schult, N., Mirhakkak, M., Chen, L., Wade, R., Duffalo, S., *"An End-to-End Modeling and Simulation Testbed (EMAST) to Support Detailed Quantitative Evaluations of GIG Transport Services"*, 10th International Command and Control Research and Technology Symposium (ICCRTS), McLean, VA, 13-16 June 2005.

[2] Durst, R., G. Miller, and E. Travis, “*TCP extensions for space communications*”, Wireless Networks 3, 1997, 389-403.

[3] Allman, M., and A. Falk, *“On Effective Evaluation of TCP”*, SIGCOMM ACM Computer Communications Review, Vol. 29, No. 5, October 1999, 59-70.

[4] Border, J., et al, “*Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations”*, Internet Engineering Task Force (IETF) RFC 3135, June 2001.

[5] Floyd, S., and T. Henderson, “*The NewReno Modification to TCP's Fast Recovery Algorithm”*, IETF RFC 2582, April 1999.

[6] Mathis, M., et al, “*TCP Selective Acknowledgement Options”*, IETF RFC 2018, October 1996.

[7] Fall, K., and S. Floyd, “*Simulation Based Comparisons of Tahoe, Reno, and SACK TCP”*, ACM Computer Communications Review, 26 (3), 5-12, 1996.

[8] Allman, M., et al, *“TCP Congestion Control”*, IETF RFC2581, April 1999.

[9] Paxson, V., et al, “*Known TCP Implementation Problems”*, IETF RFC 2525, March 1999.

[10] Tian, Y., and N. Ansari, “*TCP in Wireless Environments: Problems and Solutions”*, IEEE Radio Communications, March 2005, S27-S32.

[11] Comparetto, G., and M. Mirhakkak, “*End-to-End M&S Testbed (EMAST) Proof-of-Concept (POC) Analysis”*, November 2004.

[12] Burleigh, S., et al, “Delay-Tolerant Networking: An Approach to Interplanetary Internet”, IEEE Communications Magazine, June 2003, 128-136.

[13] Gurtov, A., and S Floyd, *“Modeling Wireless Links for Transport Protocols”*, ACM SIGCOMM Computer Communications Review, Vol 34, 2, April 2004, 85-96.

[14] Brakmo, L. and L. Peterson. *“TCP Vegas: End to End Congestion Avoidance on a Global Internet”* IEEE Journal on Selected Areas in Communication, Vol 13, 8, October 1995, 1465-1480.

[15] Casetti,, C., M. Gerla, S. Mascolo, M.Y. Sansadidi, and R. Wang, *"TCP Westwood: End-to-End Congestion Control for Wired/Wireless Networks"* Wireless Networks Journal, 8, 2002, 467-479.