Data Logistics in Network Computing: The Logistical Session Layer

Martin Swany and Rich Wolski
Computer Science Department
University of Tennessee
1122 Volunteer Avenue Knoxville, TN, 37996
{*swany,rich*}*@cs.utk.edu*

Abstract

In this paper, we present a strategy for optimizing end-to-end TCP/IP performance over long-haul networks. Our approach defines a Logistical Session Layer (LSL) that uses intermediate process-level "depots" along the network route from source to sink to implement an end-to-end communcation session. Despite the additional processing overhead resulting from TCP/IP protocol stack Unix kernel boundary traversals at each depot, our experiments show that dramatic end-to-end bandwidth improvements are possible. We also describe the prototype implementation of LSL that does not require Unix kernel modification or root access privilege that we used to generate the results, and discuss its utility in the context of extant TCP/IP tuning methodologies.

1 Introduction

The need for flexible and high-performance access to distributed resources has driven the development of networking since its inception. With the maturing of "The Internet" this community continues to increase its demands for network performance to support a raft of emerging applications including distributed collaboratoria, full-motion video, and Computational Grid programs.

Traditional models of high-performance computing are evolving hand-in-hand with advanced networking [14]. While distributed computation control and network resource control [15] techniques are currently being developed, we have been studying the use of time-limited, dynamically allocated network buffers [29] as a way of provisioning the communication medium. We term this form of networking *Logistical Networking* [7] to emphasize the higher-level control of buffer resources it entails.

In this paper, we present a novel approach to optimizing end-to-end TCP/IP performance using Logistical Networking. Our methodology inserts application-level TCP/IP

"depots" along the route from source to destination and, despite having to doubly traverse a full TCP/IP protocol stack at each depot, improves bandwidth performance. In addition, we have implemented the the communication abstractions that are necessary to manage each communication without kernel modifications as a set of session-layer semantics over over the standard byte-stream semantics supported by TCP/IP sockets. As a result, we term the abstractions we have implemented the *Logistical Session Layer* (LSL).

LSL improves end-to-end network performance by breaking long-haul TCP/IP connections into shorter TCP segments between depots stationed along the route. Staging data at the session layer in a sequence of depots increases the overhead associated with end-to-end communication. In the LSL case, data emanating from the source must be processed twice (ingress and egress) at each depot thereby increasing the overall protocol processing overhead. In this paper, we show that this performance penalty is dramatically overshadowed by the performance improvement that comes from moving TCP end-points closer together. It is counter-intuitive that adding the processor overhead incurred by traversing the protocol stack on an additional machine could actually *improve* performance. Indeed, for some time the networking community has focused on TCP/IP overhead [9, 22] and examined ways to mitigate it [23, 33, 38]. To introduce additional protocol processing runs against the current optimization trends in high-performance wide-area networking and computing. However, *despite the additional processing overhead that comes from moving the data in and out of the kernel at each depot (including checksumming costs), moving TCP end-points closer together can improve end-to-end performance.*

We present this work in the context of recent networking trends that focus on state management in the network fabric itself. While the Internet Protocol suite (as typically implemented) mandates the communication state be managed at the end-points [35], new "stateful" facilities [8, 27] which relax this restriction have been proposed. In this vein, we believe that there are several reasons that intermediate TCP processing helps, rather than hurts, end-to-end bandwidth performance. First, since the round-trip time (RTT) between any two depots is shorter than the end-to-end round-trip-time, LSL allows the inherent TCP congestion-control mechanism to sense the maximally available throughput more quickly. That is, even though the sum of the RTTs between depots may be longer than the end-to-end RTT, because the maximum RTT between any two depots is shorter, the congestion-control mechanisms adapt more rapidly. Secondly, a retransmission that results from a lost packet need not originate at the source, but rather, can be generated from the last depot to forward the data. Finally, recent advances in the processing speed, memory bandwidth, and I/O performance of commonly available processors has lowered protocol processing and data movement costs relative to available network performance. We describe, more completely, the confluence of these effects in Section 3.

In Section 2, we describe the architecture of a prototype application-layer LSL imple-

mentation that we have developed. The advantage of providing a session-layer interface is that applications do not need to employ their own customized buffer management strategies in order to use Logistical Networking to enhance end-to-end network performance. As such, our work not only provides a general methodology for improving deliverable network performance, but it also constitutes an important early example of a Grid-enabling network abstraction. At the same time, since our implementation does not require kernel modification, it is portable and easy to deploy.

Finally, in Section 4 we detail the effect of using intermediate TCP depots and LSL on end-to-end bandwidth, independent of end-point buffer settings, both with and without the RFC 1323 [21] window-scaling. Our results show that, using LSL, an application can gain a substantial end-to-end increase in bandwidth over standard TCP/IP sockets, even if the socket connections have been "tuned" for performance.

2 Architecture

The Logistical Session Layer (LSL) is a "session" layer (layer 5) in terms of the OSI protocol model. The session layer lies above the Transport layer (TCP, in the Internet Protocol suite). Recall that a transport layer conversation consists of multiple hops of network layer conversations. In an analogous fashion, a session layer conversation can consist of multiple hops of transport layer conversations. [19]. A connection that is initiated through the LSL will pass through a number of LSL-aware routers, or "depots." These devices can actually be thought of as "transport layer switches" in that they multiplex session-layer conversations onto sets of transport layer conversations. While we believe that a kernel-level implementation or dedicated system versions of these switches will, ultimately, improve performance over the results we report in the next section, we have chosen a non-privileged, application-level implementation initially. By doing so, we are able to gain two important benefits. First, because the LSL depots use standard, user-level sockets, our initial implementation of LSL does not (and, indeed, cannot) violate the current TCP congestion-control mechanisms. From the perspective of the network, an LSL session appears to be a series of user-level applications communicating in a chain. All resource control mechanisms governing "normal" user-applications (such as flow-control, congestion-control, memory-size, etc.) remain functional and need not be disabled.

Secondly, because LSL depots can run under any user login id (i.e. do not require root access), security and stability concerns are greatly reduced. It is not possible for an LSL depot to be used as a vehicle for obtaining root access because it does not run as root and it does not execute any functions not compiled into its binary image.

Additionally, our first implementation of the LSL client API closely mimics the familiar Unix sockets interface. This design choice allows easy incorporation into legacy applications. Users of the socket interface are familiar with the "Internet" address

family, denoted with **AF_INET**. We designate a new family, which we label **AF_LSL**. So, for a given program to use LSL, a simple text substitution in the source code would enable use of the system. The connection would "fall back" to using a direct TCP connection if necessary to make the change less intrusive.

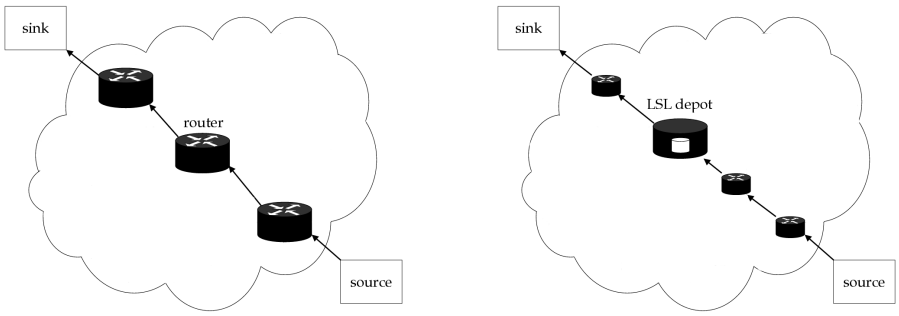

Without LSL With LSL

Figure 1: Network communication with and without LSL

Observe Figure 1 for an illustration of this architecture. This depicts the network communication between two hosts with and without LSL. A given session may pass through one or more LSL depots.

An application should be able to direct the LSL session to use a given depot if necessary to allow application-level tuning. In addition, we plan an end-to-end routing service based on the Network Weather Service [42, 43] that determines a "good" route for each client. In either case, utilizing the Sockets interface for this simply entails specifying a source routed path. When an LSL connection is initiated, a predicted path may be specified or local forwarding decisions may be relied upon. To specify a path explicitly, the sender will use the strict source route options with the LSL socket. In fact a combination of local and global forwarding strategies may be employed by specifying a loose source route in the same fashion.

To test out the effectiveness of LSL and begin to understand its potential performance benefits, we have implemented and deployed a rudimentary prototype having three components:

- a simple file server program called *lsrv*,
- a per-depot session-level daemon that establishes and releases TCP streams traversing each depot called *lsd*, and
- a client (responsible for choosing end-to-end routes) called *lget*.

Our intention is to use this framework to study both the performance characteristics of LSL and how LSL may be implemented for computational Grid settings. By thus modularizing the LSL system, we will be able to take advantage of the functionality provided by systems such a Globus [13], Legion [16], and the Network Weather Service [43].

2.1 The End-to-End Argument and the Session Layer

The architecture of the current Internet Protocol suite has been guided by what is known as the "end-to-end" argument [35]. This model (as commonly understood) dictates that state be kept at the end nodes and that the core of the network be stateless primarily for reasons of reliability. Recent trends in network service provision [8, 27], however, relax the requirement of statelessness in favor of better performance and service quality control. Indeed, the general question of end-to-end versus stateful networking is also being considered explicitly by many, including the original authors of the end-to-end argument [8, 32]. LSL is another example of how careful state management within the network fabric itself can improve delivered network performance while, at the same time, preserving the stability and reliability characteristics that the Internet Protocol suite provides. In addition, the architecture we have defined is compatible the current implementations of TCP/IP while offering a similar programming interface to that provided by the Unix socket abstraction.

In short, the general application of the principle is somewhat different than the networking community at large has come to understand. However, since we use the semantics of a session layer to implement our system, even the most dogmatic network engineer will be left without argument.

3 Observations

The key idea behind LSL is that, by allowing an application to temporarily and anonymously utilize buffers "in" the network, end-to-end performance will be enhanced. It is intuitive that there is a fundamental cost associated with buffering unacknowledged segments for retransmission. Moreover, it is clear that the problem is only exacerbated as network speeds increase.

By its very definition LSL causes the end-to-end connection to have a larger *aggregate window*. We define this to be the sum of the minimum of the congestion window (*cwnd*) and the advertised window (*rwnd*) over each of the TCP connections. In exposing the pipeline we have, in fact, increased the "capacity" of the network.

In addition, LSL optimizes the end-to-end bandwidth performance in two ways: by improving the response of the congestion-control mechanisms that are currently in use, and by exploiting locality for packet retransmission.

3.1 Congestion Control and Cascaded TCP

By cascading TCP streams, LSL affects TCP congestion control in two ways. First, it shortens the RTT that any constituent TCP stream uses to "clock" the rate at which congestion-window modifications are made. Secondly, it isolates RTT variance so that retransmission times across low variance links are based on more accurate estimates of variance.

By staging data above the transport layer at depots along the path from source to sink, LSL reduces the bandwidth-delay product associated with any single TCP connection. As such, it allows the adaptive congestion-control mechanisms [2]) to achieve maximal throughput and recover from loss more quickly. The flow-control and congestion-control mechanisms used by TCP require that an acknowledgement (ACK) be sent from the receiver. This stream of ACKs acts as a clock for strobing packets into the network [20]. The speed with which slow-start allows the TCP connection to approach the advertised flow-control window is determined with the RTT (measured as the sum of the transit time of a packet and its ACK.) The effects of RTT have been observed [24,26] but intuitively, since increase in congestion window requires a full RTT, the longer the RTT, the longer it takes TCP to reach full link capacity. By using multiple LSL depots with stream between each, TCP can discern the congestion-and flow-control-levels for each communication and achieve maximal link capacity more quickly.

The second effect on TCP comes from a reduced variance in RTT. By shortening the length of the links that TCP traverses, LSL potentially isolates variance in packet delay. The TCP protocol uses an estimate of the variance in RTT to determine when a retransmission is necessary in response to a lost packet. if the TCP stream traverses a large number of hops, high variance in queue delay at any point along the route affects the retransmission time out for the entire route. By breaking the end-to-end route up into segments between depots, LSL allows TCP to better estimate RTT variance on a link-by-link basis. The result is that retransmissions are more accurately triggered, and performance is improved.

3.2 Locality and Packet Retransmission

By buffering data at intermediate points along the route from source to sink, LSL reduces the overhead associated with retransmission. With the current TCP buffering model, a lost packet at any point between source and sink requires a retransmission from the communication source point. The retransmitted packet, then, must traverse the entire network using valuable bandwidth along the entire route. For example, consider an end-to-end communication between two hosts separated by 10 routers or gateways in which packets are being dropped at the last gateway before the sink. Every time the 10th router drops a packet, the resulting retransmission must traverse the other 9 routers, taking up scarce buffer and bandwidth resources. By buffering data at the session layer,

LSL ensures that any retransmits traverse only the links between depots. The result is a savings in the bandwidth that would otherwise be wasted from end-to-end moving retransmitted data.

4 Results

The place that we expect the LSL optimization effects to be most apparent is in long-running transfers over networks where the bandwidth-delay product is high. In this section, we examine several example transfer paths that terminate at the University of Tennessee, Knoxville (UTK). Since we had complete control over the UTK machines, we were able to investigate the effects of different kernel-level TCP settings.

In the first test we study transfers from Argonne National Laboratory (ANL) to the University of Tennessee (UTK). To do so, we deploy an LSL daemon at Oak Ridge National Laboratory (ORNL) to serve as a depot between UTK and ANL. UTK is directly connected to ORNL via an OC-3 (155 Mb/sec) link, and ORNL and ANL are both connected to the Energy Sciences Network (ESnet) at OC-12 (622 Mb/sec) [11]. Appendix A is the output of the *traceroute* command from UTK to both other sites, Appendix B and C are from ANL and ORNL, respectively.

The ANL, ORNL, and UTK machines were configured to use the RFC1323 [21] window-scaling optimizations and large kernel buffers. For this experiment, we set the kernel buffers (through the Unix *setsockopt()* command) to be eight megabytes at both ends, and verified that the correct window size was bring quoted using *getsockopt* and *tcpdump* at the UTK end.

The results in Figure 2 represent roughly 280 experiments in total. Along the x-axis we show a series of different transfer sizes. The y-axis of the figure indicates the observed, end-to-end throughout in megabits per second. Each data point represents the average throughput observed over 20 different transfers at transfer size corresponding to its x coordinate.[1] In addition, the x-axis is shown on a log scale.

Figure 2 shows that the LSL does indeed optimize end-to-end transfers of 256KBytes and larger. For transfers of 32Mbytes, use of the LSL depot at ORNL increases the average performance by well over a factor of 2.5.

While investigating these results, we observed that the route between UTK and ANL is asymmetric as the traceroute from ANL to UTK (in Appendix B indicates. As this is the case, part of this improvement can be attributed to the fact that by explicitly routing through ORNL, we are enforcing symmetric paths. This "user-controlled" routing is somewhat analogous to the IP "source route" option and we discuss similarities to this and other approaches in Section 5.

[1] Appendix G shows some summary statistics and average transfer rates for each transfer size we consider in this paper. For visual clarity, the figures we present depict the averages only.

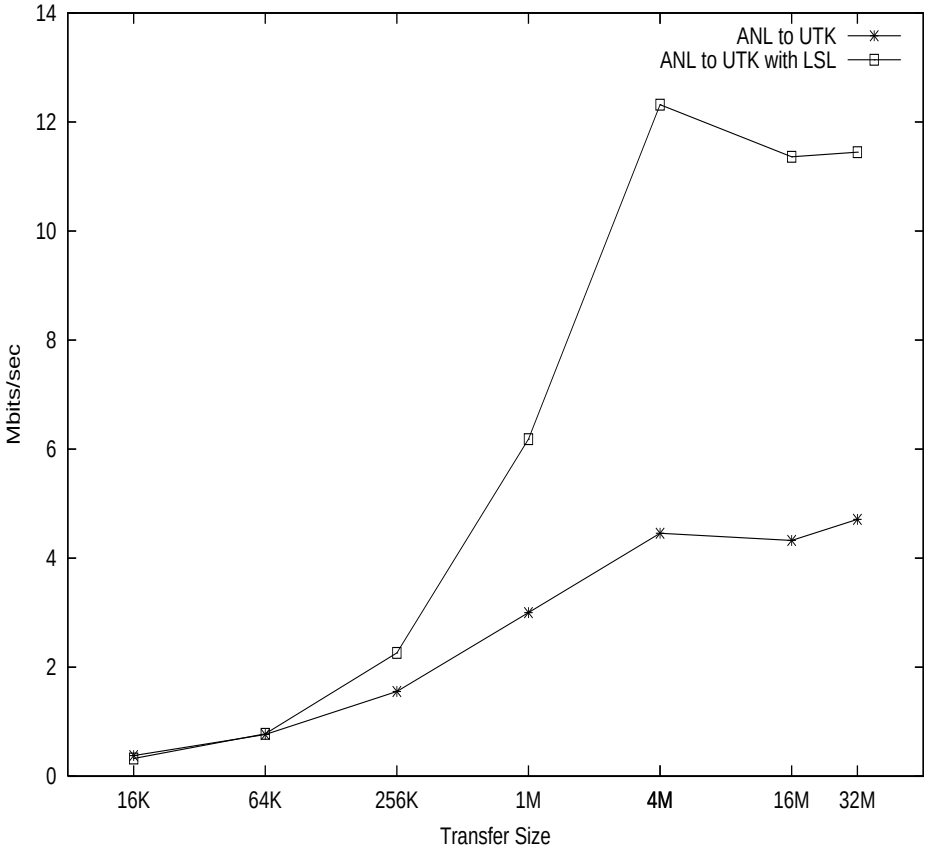

Figure 2: Data transfer from ANL to UTK

Since this might be viewed as a pathological case (although recent work [28] indicates otherwise) we sought a case in which the LSL route corresponded with the non-LSL route in both directions. We chose a path between the University of Tennessee (UTK) and the University of California at San Diego (UCSD) using a machine located at Texas A&M University, South Campus which is very near the Abilene [1] "Point of Presence" (POP) in Houston. Traffic from UTK to UCSD and vice versa traverses this POP. Appendix D shows the traceroute from UTK to UCSD and TAMUS. Appendix E and F show traceroutes from UCSD and TAMUS, respectively.

Figure 3 illustrates the comparison of LSL-enabled and non-LSL-enabled end-to-end bandwidth performance from UCSD to UTK, again using 8 MB buffers and windows. Here, LSL offers as much as a %50 improvement over direct TCP *despite* adding to the gross latency and protocol processing overhead along the path from source to sink.

For high-capacity long-haul networks (like Abilene), large window sizes are necessary so that the sending side does not block due to flow-control before an acknowledgement

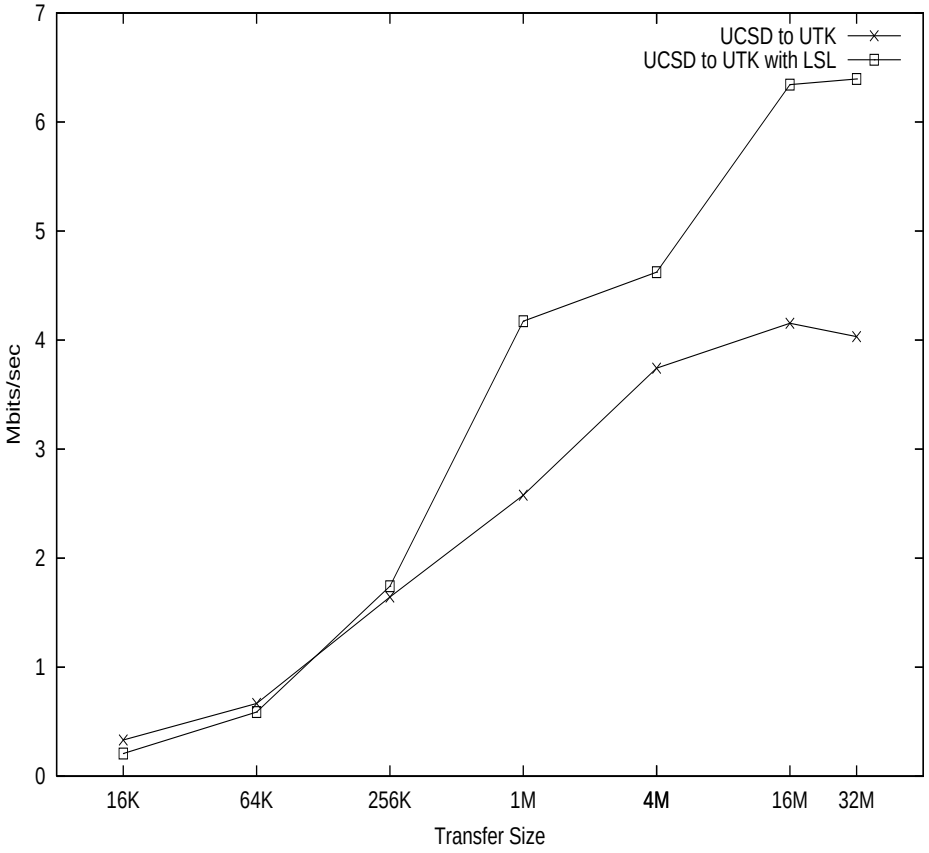

Figure 3: Data transfer from UCSD to UTK

has time to return from the receiver. That is, the buffering must allow the amount of data in flight to exceed the bandwidth-delay product. By choosing 8 MB windows and buffers, we ensure that the LSL effects we observe not not simply due "bad" buffering choices at the end points. We believe that an optimized TCP stream using large buffers at either end would see similar performance improvements.

Although we recognize that buffers of this size may not be optimally tuned, this does not effect our results. The danger in over-sizing buffers is in wasting resources on the host, not in causing poor TCP performance [38]. We judged this to be acceptable for this experiment, but we do await the products produced by groups like the Web100 [41] and the Internet2 End-to-End Performance initiative [18].

However, not all hosts support (through design or configuration) large window sizes. For this case, we wanted to investigate how using an LSL depot with large windows might enhance the performance of "untuned" TCP streams. Figure 4 shows the average transfer rates when the buffers at the sending and receiving ends are restricted to 64K

bytes.

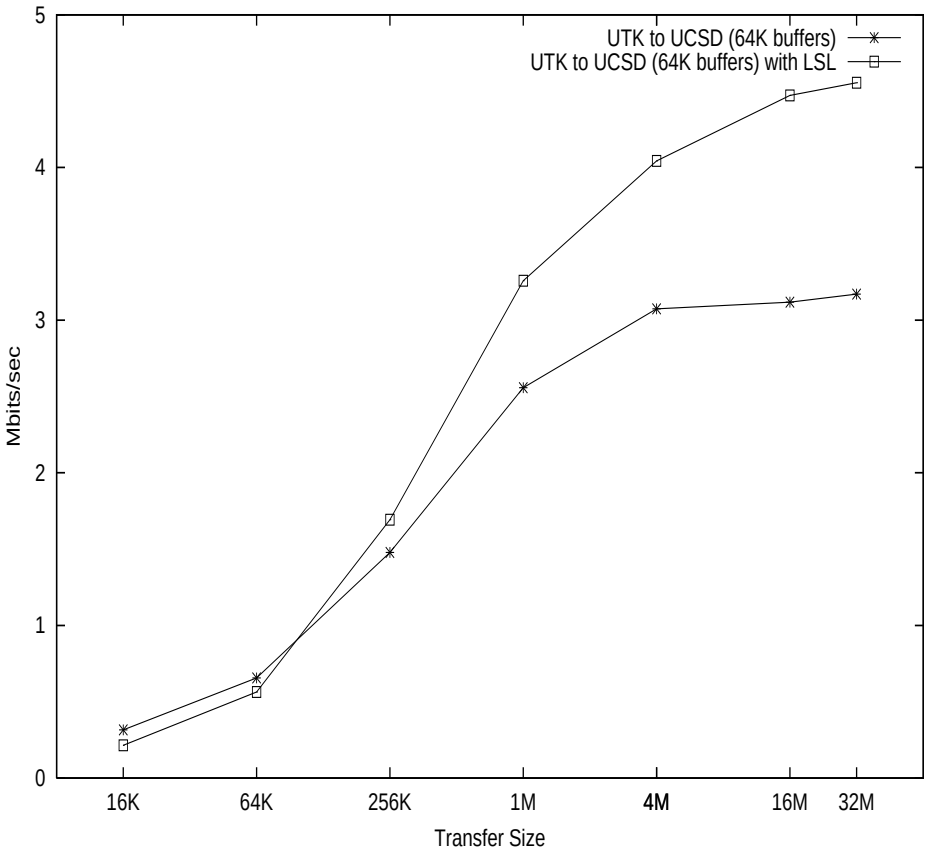

Figure 4: Data transfer from UCSD to UTK with 64KB buffers

As expected, the absolute performance is lower. However, the LSL-enabled stream was still able to outperform the non-enabled stream by %43 for the largest transfer size. As a non-intrusive optimization, however, we believe that such a performance improvement will be useful in many application settings.

Finally, Figure 5 shows the effect of using LSL on a TCP stream that does no *setsockopt()* buffer conditioning whatsoever.

Again, the performance is not as dramatic as in the tuned cases, but the LSL effect is still present.

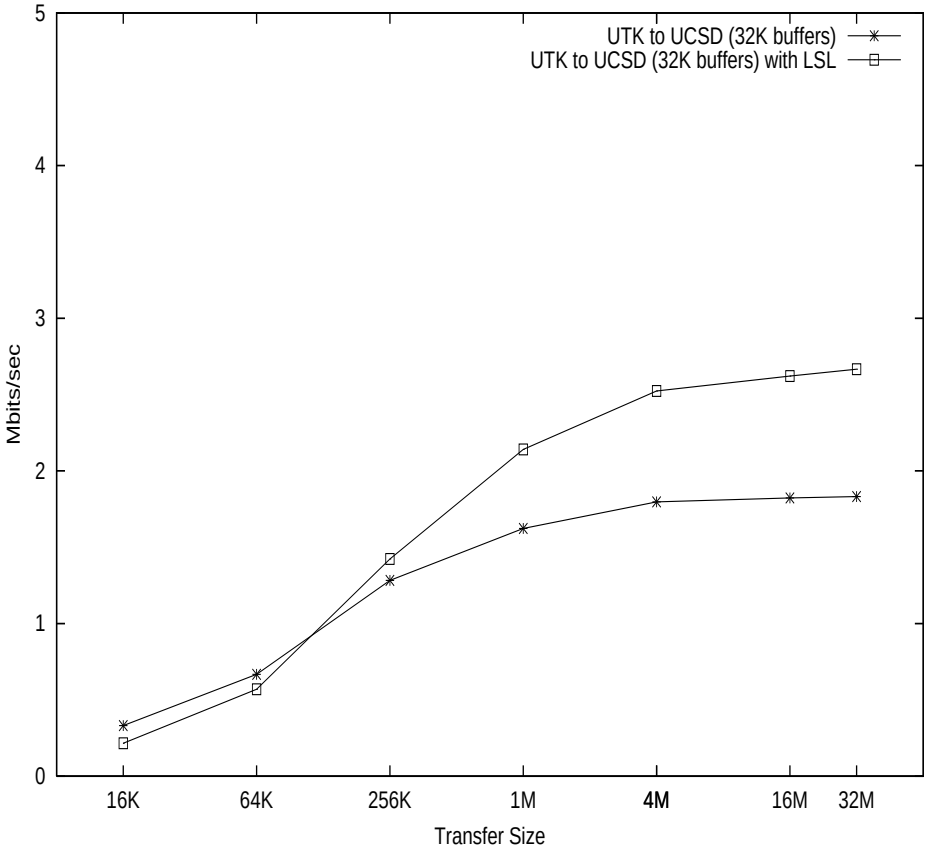

Figure 5: Data transfer from UCSD to UTK with default (32KB) buffers

5 Related Work

There are many areas in which facets of our results and similar mindsets can be seen in the community. Two broad categories are TCP effects (retransmission cost and locality) and control over the topology.

Techniques developed for wireless networks [5,6] seek to mitigate the cost of retransmission in lossy environments. However, they violate layering to do so. Systems to proxy TCP have been developed with the same goals in mind. One example is TPOT [34], which alters TCP to allow this mode of operation (and therefore has dubious possibility for actual deployment in the global Internet.) A similar approach targeting caching of web objects also proposes modifications to TCP [25].

There are many approaches to reducing the cost of retransmission with the network's assistance. One of the areas that has pushed this notion forward is the wireless commu-

nity. Since wireless links at the edges of the network tend to be much less reliable than other parts of the network, the almost inevitable packet loss hampers the performance of TCP. The "snoop" technique watches TCP traffic and can perform retransmits from and intermediate point in the network. [6]. This protocol comes from a very similar mindset, but could be considered inappropriate in that it violates layering. Further, flow-level snooping is expensive, so the scalability of this approach for high-bandwidth networks is questionable. Also targeting wireless is Indirect TCP [5], which is similar to our approach. Another similar approach is that proposed for use in multicast video transmission [39] in that it allows a data stream to take advantage of retransmit locality and finer-grained adaptability.

The PSockets [12] work addresses the difficulty of getting high-bandwidth from long, fat networks. This work is similar to LSL in that they preserve the syntax of the well-known socket interface while taking liberties underneath. They differ in mechanism and spirit, however. This approach to higher bandwidth uses multiple TCP streams and keeps aggregate throughput high by amortizing loss over several streams. The parallel sockets approach is to ignore TCP's inefficiencies in the wide area and aim for better average throughput. However, the streams are then inducing the congestion that the other streams are sensing. So, Psockets provides better average performance but worse link utilization. LSL allows the network to be "articulated" and allows TCP to respond to congestion where it exists without introducing additional load.

The need to control the topology by tunneling from one host to another has been identified by a body of engineering and by the sheer number of Virtual Private Networks that are in use. Even outside the desire to provide an encrypted tunnel, virtual topology systems are abundant [36, 40]. The performance benefits of user-optimization of paths has been discussed [3, 31, 37]. While LSL has the functionality of these systems, our approach is different in that we consider this to be part of a "session" layer of services. Our empirical evidence does establish the viability of this line of research.

Finally, within the networking community, the notion has long existed that the end user will occasionally need to have explicit control over the route that traffic follows. Loose and strict source routing were defined in RFC 791 [30], which defines the Internet Protocol. Store and forward connectivity has been used for quite some time in the networking community and there are many situations in which data transfer need not be synchronous or connection-oriented. SMTP [10], USENET [17] and its successor NNTP [4], all use hop-oriented, connectionless paradigms to send data.

6 Conclusion

With the maturing of network infrastructure, both in terms of ubiquity and quality, comes the possibility of increasing the state held in the network, and the time duration over which it is maintained. Logistical Networking [7] attempts to define the parameters

under which the added cost associated with maintaining state is overshadowed by an increase in delivered network performance. The Logistical Session Layer (LSL) is an implementation of Logistical Networking concepts to improve end-to-end communication performance between applications that currently use the TCP socket interface. Our early results show that LSL can result in dramatic throughput performance improvements despite greater protocol processing overhead. In addition, our initial prototype is serving as an architectural framework within which he hope to generalize these results.

Important research questions exist such as the permanence of the buffering, the nature of the flow control mechanisms and alternate architectures for a mechanism such as this. We believe that we have demonstrated the efficacy of this system and that it is novel, viable approach.

References

[1] Abilene. http://www.ucaid.edu/abilene/.

[2] M. Allman, V. Paxson, and W. Stevens. TCP congestion control. *RFC 2581*, page 13, 1999.

[3] D. Andersen, H. Balakrishnan, M. Kaashoek, and R. Morris. The case for resilient overlay networks. In *8th Annual Workshop on Hot Topics in Operating Systems (HotOS-VIII)*, May 2001.

[4] P. L. B. Kantor. A proposed standard for the stream-based transmission of news. RFC 977, February 1986.

[5] A. Bakre and B. R. Badrinath. I-TCP: Indirect TCP for mobile hosts, 1995.

[6] H. Balakrishnan, S. Seshan, and R. H. Katz. Improving reliable transport and handoff performance in cellular wireless networks. *ACM Wireless Networks*, 1(4), 1995.

[7] M. Beck, T. Moore, J. Plank, and M. Swany. Logistical networking: Sharing more than the wires. In *Proc. of 2nd Annual Workshop on Active Middleware Services*, August 2000.

[8] R. Braden, D. Clark, and S. Shenker. Integrated services in the internet architecture: an overview. RFC 1633, June 1994.

[9] D. D. Clark, J. Romkey, and H. Salwen. An analysis of TCP processing overhead. In *Proc. 13th Conference on Local Computer Networks*, pages 284–291, Minneapolis, Minn., 10-12 1988. IEEE Computer Society.

[10] D. Crocker. Standard for the format of arpa internet text messages. RFC 822, August 1992.

[11] Esnet. http://www.es.net.

[12] X. fixme. Psockets: The case for application-level network striping for data intensive applications using high speed wide area networks. SC2000, Nov 2000.

[13] I. Foster and C. Kesselman. Globus: A metacomputing infrastructure toolkit. *International Journal of Supercomputer Applications*, 1997.

[14] I. Foster and C. Kesselman. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann Publishers, Inc., 1998.

[15] I. Foster, C. Kesselman, C. Lee, B. Lindell, K. N. dt, and A. Roy. A distributed resource management architecture that supports advance d reservation and co-allocation. In *International Workshop on Quality of Service*, 1999.

[16] A. S. Grimshaw, W. A. Wulf, J. C. French, A. C. Weaver, and P. F. Reynolds. Legion: The next logical step towrd a nationwide virtual computer. Technical Report CS-94-21, University of Virginia, 1994.
[17] M. Horton. Standard for interchange of usenet messages. RFC 850, June 1983.
[18] Internet2 end-to-end performance initiative. http://www.internet2.edu/e2eperf/.
[19] I. T. O. S. I. International Organization for Standardization. ISO/IEC 8827, 1990.
[20] V. Jacobson. Congestion avoidance and control. *ACM Computer Communication Review; Proceedings of the Sigcomm '88 Symposium in Stanford, CA, August, 1988*, 18, 4:314–329, 1988.
[21] V. Jacobson, R. Braden, and D. Borman. Tcp extensions for high performance. Network Working Group, Internet Engineering Task Force. Request For Comments: 1323, May 1992.
[22] J. Kay and J. Pasquale. The importance of non-data touching processing overheads in TCP/IP. In *SIGCOMM*, pages 259–268, 1993.
[23] J. Kay and J. Pasquale. Profiling and reducing processing overheads in tcp/ip. *IEEE/ACM Transactions on Networking*, 4(6):817–828, December 1996.
[24] T. Lakshman and U. Madhow. The performance of tcp/ip for networks with high bandwidth-delay products and random loss, 1997.
[25] U. Legedza and J. Guttag. Using network-level support to improve cache routing. *Computer Networks and ISDN Systems*, 30(22–23):2193–2201, 1998.
[26] M. Mathis, J. Semke, J. Mahdavi, and T. Ott. The macroscopic behavior of the tcp congestion avoidance algorithm, 1997.
[27] K. Nichols, S. Blake, F. Baker, and D. Black. Definition of the differentiated services field, 1998.
[28] V. Paxson. End-to-end internet packet dynamics. *IEEE/ACM Transactions on Networking*, 7(3):277–292, 1999.
[29] J. Plank, , A. Bassi, M. Beck, T. Moore, M. Swany, and R. Wolski. The internet backplane protocol: Storage in the network. IEEE Internet Computing (to appear), 2001.
[30] J. Postel. Internet protocol. RFC 791, USC/Information Sciences Institute, September 1981.
[31] N. Rao. Netlets: End-to-end qos mechanisms for distributed computing over internet using two-paths. Int. Conf. on Internet Computing, 2001.
[32] D. Reed, J. Saltzer, and D. Clark. Commentaries on the active networking and end-to-end arguments, 1998.
[33] S. H. Rodrigues, T. E. Anderson, and D. E. Culler. High-performance local-area communication with fast sockets. In *USENIX 1997 Annual Technical Conference*, pages 257–274. USENIX, 1997.
[34] P. Rodriguez, S. Sibal, and O. Spatscheck. Tpot: Translucent proxying of tcp, 2000.
[35] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4):277–288, 1984.
[36] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan. Detour: a case for informed internet routing and transport. Technical Report TR-98-10-05, University of Washington, 1998.
[37] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson. The end-to-end effects of internet path selection. In *SIGCOMM*, pages 289–299, 1999.
[38] J. Semke, J. Mahdavi, and M. Mathis. Automatic TCP buffer tuning. In *SIGCOMM*, pages 315–323, 1998.
[39] S. Sen, D. Towsley, Z.-L. Zhang, and J. Dey. Optimal multicast smoothing of streaming video over an internetwork. Technical Report UM-CS-1998-077, 1998.

[40] J. Touch. The xbone. Workshop on Research Directions for the Next Generation Internet, May 1997.
[41] Web100. http://www.web100.org.
[42] R. Wolski. Dynamically forecasting network performance using the network weather service. *Cluster Computing*, 1:119–132, January 1998. also available from http://www.cs.utk.edu/ rich/publications/nws-tr.ps.gz.
[43] R. Wolski, N. Spring, and J. Hayes. The network weather service: A distributed resource performance forecasting service for metacomputing. *Future Generation Computer Systems*, 15(5-6):757–768, October 1999. available from `http://www.cs.utk.edu/~rich/publications/nws-arch.ps.gz`.

A Traceroute from UTK to ANL and UTK to ORNL

```
traceroute to pitcairn.mcs.anl.gov (140.221.9.180), 30 hops max, 38 byte packets
 1  R6HM01V150.NS.UTK.EDU (160.36.56.1)  0.220 ms  0.181 ms  0.168 ms
 2  128.169.192.241 (128.169.192.241)  1.631 ms  0.711 ms  0.770 ms
 3  R7DH03G11-0-0.NS.UTK.EDU (192.168.101.3)  2.318 ms  2.021 ms  1.652 ms
 4  UTK-GATECH.NS.UTK.EDU (128.169.50.246)  68.930 ms  65.490 ms  65.199 ms
 5  esnet-sox-rtr.sox.net (199.77.192.6)  48.288 ms  48.575 ms  48.643 ms
 6  orn-gsu.es.net (134.55.208.65)  48.211 ms  48.312 ms  48.502 ms
 7  nyc-s-orn.es.net (134.55.205.109)  81.104 ms  81.595 ms  80.854 ms
 8  chi-s-nyc.es.net (134.55.205.106)  101.551 ms  100.881 ms  101.062 ms
 9  anl-chi-ds3.es.net (134.55.208.150)  102.894 ms  103.809 ms  102.939 ms
10  anl-esanl2.es.net (198.124.254.166)  105.927 ms  105.344 ms  105.025 ms
11  stardust-msfc-20.mcs.anl.gov (140.221.20.124)  106.386 ms  108.505 ms  106.758 ms
12  pitcairn.mcs.anl.gov (140.221.9.180)  106.309 ms *  105.946 ms


traceroute to falcon0j.ccs.ornl.gov (160.91.199.21), 30 hops max, 38 byte packets
 1  R6HM01V150.NS.UTK.EDU (160.36.56.1)  0.242 ms  0.184 ms  0.171 ms
 2  128.169.192.241 (128.169.192.241)  0.942 ms  0.848 ms  0.664 ms
 3  R7DH03G11-0-0.NS.UTK.EDU (192.168.101.3)  1.743 ms  1.805 ms  1.657 ms
 4  mmesgwya32.ctd.ornl.gov (192.31.96.17)  3.222 ms  3.219 ms  3.689 ms
 5  192.31.96.225 (192.31.96.225)  2.203 ms  2.493 ms  2.061 ms
 6  ornlgwy-ext.ens.ornl.gov (198.124.42.162)  2.358 ms  2.475 ms  2.736 ms
 7  ccsrtr.ccs.ornl.gov (160.91.0.66)  2.198 ms  2.421 ms  2.094 ms
 8  falcon0j.ccs.ornl.gov (160.91.199.21)  2.194 ms  2.574 ms  2.289 ms
```

B Traceroute from ANL to UTK and ORNL

```
traceroute to modulus.cs.utk.edu (160.36.59.32): 1-30 hops, 38 byte packets
 1  stardust-msfc-11.mcs.anl.gov (140.221.11.251)  6.49 ms  11.1 ms  4.73 ms
 2  kiwi.anchor.anl.gov (140.221.20.97)  0.614 ms  0.721 ms  0.639 ms
 3  abilene-anl.anchor.anl.gov (192.5.170.170)  5.97 ms  5.35 ms  5.29 ms
```

```
 4  atla-ipls.abilene.ucaid.edu (198.32.8.42)  15.7 ms  15.0 ms  15.2 ms
 5  sox-rtr.abilene.sox.net (199.77.193.9)  41.1 ms 40.7 ms  40.2 ms
 6  R7DH03A1-0-2.NS.UTK.EDU (128.169.50.245)  107 ms  107 ms  107 ms
 7  *  *  *
 8  128.169.192.242 (128.169.192.242)  107 ms  106 ms  106 ms
 9  MODULUS.CS.UTK.EDU (160.36.59.32)  106 ms  106 ms  106 ms

traceroute to falcon0j.ccs.ornl.gov (160.91.199.21): 1-30 hops, 38 byte packets
 1  stardust-msfc-11.mcs.anl.gov (140.221.11.251)  0.728 ms  0.792 ms  0.515 ms
 2  kiwi.anchor.anl.gov (140.221.20.97)  0.377 ms  0.315 ms  0.279 ms
 3  esanl2-anl.es.net (198.124.254.165)  0.783 ms  0.989 ms  0.504 ms
 4  chi-anl-ds3.es.net (134.55.208.149)  3.33 ms  2.34 ms  2.35 ms
 5  nyc-s-chi.es.net (134.55.205.105)  23.1 ms  23.0 ms  22.7 ms
 6  orn-s-nyc.es.net (134.55.205.110)  56.2 ms  55.8 ms  54.9 ms
 7  ornl-orn.es.net (134.55.208.62)  76.0 ms  75.1 ms  75.0 ms
 8  192.31.96.225 (192.31.96.225)  74.2 ms  74.8 ms  75.7 ms
 9  ornlgwy-ext.ens.ornl.gov (198.124.42.162)  75.0 ms  75.5 ms  74.9 ms
10  ccsrtr.ccs.ornl.gov (160.91.0.66)  75.3 ms  74.6 ms  75.3 ms
11  falcon0j.ccs.ornl.gov (160.91.199.21)  75.5 ms  75.1 ms  74.7 ms
```

C Traceroute from ORNL to UTK and ANL

```
traceroute to modulus.cs.utk.edu (160.36.59.32): 1-30 hops, 38 byte packets
 1  ccsrtr-003.ccs.ornl.gov (160.91.199.1)  0.341 ms  0.270 ms  0.418 ms
 2  160.91.0.65 (160.91.0.65)  0.490 ms  0.299 ms  0.357 ms
 3  orgwy2.ens.ornl.gov (198.124.42.185)  0.330 ms  0.385 ms  0.327 ms
 4  mmesgwy-ext-fe.cind.ornl.gov (192.31.96.235)  2.5 ms  1.54 ms  1.95 ms
 5  utk-rtr.ctd.ornl.gov (192.31.96.18)  3.54 ms  2.52 ms  3.8 ms
 6  * * *
 7  128.169.192.242 (128.169.192.242)  3.27 ms  2.40 ms  2.19 ms
 8  MODULUS.CS.UTK.EDU (160.36.59.32)  2.46 ms  2.30 ms  2.47 ms

traceroute to pitcairn.mcs.anl.gov (140.221.9.180): 1-30 hops, 38 byte packets
 1  ccsrtr-003.ccs.ornl.gov (160.91.199.1)  2.55 ms  0.575 ms  0.342 ms
 2  160.91.0.65 (160.91.0.65)  0.338 ms  4.44 ms  0.796 ms
 3  orgwy2.ens.ornl.gov (198.124.42.185)  1.69 ms  0.290 ms  0.410 ms
 4  ornl-rt3-ge.cind.ornl.gov (192.31.96.230)  0.674 ms  0.503 ms  3.35 ms
 5  orn-ornl.es.net (134.55.208.61)  18.4 ms  19.9 ms  19.6 ms
 6  nyc-s-orn.es.net (134.55.205.109)  52.6 ms  52.5 ms  52.8 ms
 7  chi-s-nyc.es.net (134.55.205.106)  73.1 ms  72.9 ms  72.4 ms
 8  anl-chi-ds3.es.net (134.55.208.150)  73.5 ms  73.9 ms  81.2 ms
 9  anl-esanl2.es.net (198.124.254.166)  75.8 ms  74.1 ms  74.1 ms
10  stardust-msfc-20.mcs.anl.gov (140.221.20.124)  74.5 ms  76.1 ms  76.4 ms
11  pitcairn.mcs.anl.gov (140.221.9.180)  74.6 ms *  75.8 ms
```

D Traceroutes from UTK to UCSD and TAMUS

```
traceroute to freak.ucsd.edu (132.239.95.142), 30 hops max, 38 byte packets
 1  R5HM01V277.NS.UTK.EDU (128.169.92.1)  95.333 ms  13.187 ms  3.477 ms
 2  R7DH03G11-0-0.NS.UTK.EDU (192.168.101.3)  6.608 ms  1.406 ms  2.253 ms
 3  UTK-GATECH.NS.UTK.EDU (128.169.50.246)  65.648 ms  66.254 ms  66.578 ms
 4  199.77.193.10 (199.77.193.10)  65.590 ms  65.773 ms  67.116 ms
 5  hstn-atla.abilene.ucaid.edu (198.32.8.33)  86.127 ms  84.799 ms  85.412 ms
 6  losa-hstn.abilene.ucaid.edu (198.32.8.21)  116.830 ms  117.402 ms  117.387 ms
 7  USC--abilene.ATM.calren2.net (198.32.248.85)  117.384 ms  117.809 ms  117.693 ms
 8  UCSD--USC.POS.calren2.net (198.32.248.34)  120.569 ms  120.970 ms  120.841 ms
 9  sdsc2--UCSD.ATM.calren2.net (198.32.248.65)  121.533 ms  122.205 ms  121.619 ms
10  cse-rs.ucsd.edu (132.239.254.45)  122.926 ms  122.876 ms  122.653 ms
11  freak.ucsd.edu (132.239.95.142)  122.757 ms  122.534 ms  122.110 ms

traceroute to i2-dsi.ibt.tamus.edu (165.95.234.23), 30 hops max, 38 byte packets
 1  R6HM01V150.NS.UTK.EDU (160.36.56.1)  0.230 ms  0.189 ms  0.172 ms
 2  128.169.192.241 (128.169.192.241)  0.890 ms  1.002 ms  0.939 ms
 3  192.168.101.3 (192.168.101.3)  2.096 ms  1.592 ms  1.197 ms
 4  UTK-GATECH.NS.UTK.EDU (128.169.50.246)  66.649 ms  66.018 ms  65.296 ms
 5  atla.abilene.sox.net (199.77.193.10)  65.601 ms  65.199 ms  65.430 ms
 6  hstn-atla.abilene.ucaid.edu (198.32.8.33)  84.666 ms  84.745 ms  84.695 ms
 7  LINK2ABILENE.GIGAPOP.GEN.TX.US (198.32.236.13)  85.458 ms  97.486 ms  87.325 ms
 8  LINK2IBT.GIGAPOP.GEN.TX.US (198.32.236.34)  86.562 ms  86.351 ms  86.420 ms
 9  ibtx2-atm10-401.ibt.tamus.edu (165.95.232.6)  87.595 ms  87.494 ms  87.433 ms
10  i2-dsi.ibt.tamus.edu (165.95.234.23)  87.476 ms  87.740 ms  87.811 ms
```

E Traceroutes from UCSD to UTK and TAMUS

```
traceroute to modulus.cs.utk.edu (128.169.94.246), 30 hops max, 38 byte packets
 1  cse-danger-gateway.ucsd.edu (132.239.95.1)  0.723 ms  0.702 ms  0.622 ms
 2  bigmama.ucsd.edu (132.239.254.5)  1.860 ms  1.321 ms  1.224 ms
 3  UCSD--sdsc2.ATM.calren2.net (198.32.248.66)  1.270 ms  1.153 ms  1.542 ms
 4  USC--UCSD.POS.calren2.net (198.32.248.33)  4.997 ms  5.054 ms  5.961 ms
 5  abilene--USC.ATM.calren2.net (198.32.248.86)  5.412 ms  5.223 ms  4.955 ms
 6  hstn-losa.abilene.ucaid.edu (198.32.8.22)  37.056 ms  37.352 ms  37.429 ms
 7  atla-hstn.abilene.ucaid.edu (198.32.8.34)  57.223 ms  56.851 ms  56.922 ms
 8  199.77.193.9 (199.77.193.9)  56.819 ms  56.770 ms  57.155 ms
 9  R7DH03A1-0-2.NS.UTK.EDU (128.169.50.245)  124.427 ms  124.783 ms  124.299 ms
10  192.168.101.40 (192.168.101.40)  169.204 ms  125.048 ms  124.302 ms
11  MODULUS.CS.UTK.EDU (128.169.94.246)  122.207 ms  121.958 ms  122.852 ms

traceroute to i2-dsi.ibt.tamus.edu (165.95.234.23), 30 hops max, 38 byte packets
 1  cse-danger-gateway.ucsd.edu (132.239.95.1)  0.765 ms  0.661 ms  0.657 ms
 2  nodeb-rs-backbone.ucsd.edu (132.239.254.26)  1.099 ms  1.288 ms  2.827 ms
```

```
 3  nodeB-6500-5500-GE.ucsd.edu (132.239.255.150)  0.842 ms  0.643 ms  0.706 ms
 4  ucsd-gw-nodeB.ucsd.edu (132.239.255.142)  0.872 ms  0.737 ms  0.714 ms
 5  198.32.248.185 (198.32.248.185)  1.508 ms  0.826 ms  0.673 ms
 6  USC--UCSD.POS.calren2.net (198.32.248.33)  4.374 ms  4.206 ms  5.037 ms
 7  Abilene--USC.ATM.calren2.net (198.32.248.86)  4.978 ms  5.411 ms  5.360 ms
 8  hstn-losa.abilene.ucaid.edu (198.32.8.22)  36.870 ms  36.619 ms  36.987 ms
 9  LINK2ABILENE.GIGAPOP.GEN.TX.US (198.32.236.13)  38.636 ms  37.631 ms  37.596 ms
10  LINK2IBT.GIGAPOP.GEN.TX.US (198.32.236.34)  38.995 ms  38.918 ms  38.093 ms
11  ibtx2-atm10-401.ibt.tamus.edu (165.95.232.6)  40.486 ms  39.158 ms  39.347 ms
12  i2-dsi.ibt.tamus.edu (165.95.234.23)  41.655 ms  40.401 ms  39.469 ms
```

F Traceroutes from TAMUS to UTK and UCSD

```
traceroute to modulus.cs.utk.edu (160.36.59.32), 30 hops max, 40 byte packets
 1  ibtx2-atm10-1 (165.95.234.254)  1.015 ms  0.802 ms  0.757 ms
 2  ibtx1-atm10-401 (165.95.232.5)  1.451 ms  1.260 ms  1.239 ms
 3  198.32.236.33 (198.32.236.33)  1.915 ms  1.997 ms  1.795 ms
 4  ABILENE.GIGAPOP.GEN.TX.US (198.32.236.14)  2.224 ms  2.279 ms  2.129 ms
 5  atla-hstn.abilene.ucaid.edu (198.32.8.34)  22.184 ms  21.755 ms  21.754 ms
 6  sox-rtr.abilene.sox.net (199.77.193.9)  21.936 ms  21.987 ms  21.861 ms
 7  R7DH03A1-0-2.NS.UTK.EDU (128.169.50.245)  87.940 ms  89.067 ms  87.372 ms
 8  192.168.101.40 (192.168.101.40)  87.845 ms  87.509 ms  87.301 ms
 9  128.169.192.242 (128.169.192.242)  86.897 ms  87.058 ms  87.171 ms
10  MODULUS.CS.UTK.EDU (160.36.59.32)  87.495 ms  87.104 ms  87.109 ms

traceroute to freak.ucsd.edu (132.239.95.186), 30 hops max, 40 byte packets
 1  ibtx2-atm10-1 (165.95.234.254)  0.991 ms  0.793 ms  0.759 ms
 2  ibtx1-atm10-401 (165.95.232.5)  1.413 ms  1.299 ms  1.273 ms
 3  198.32.236.33 (198.32.236.33)  2.069 ms  1.811 ms  1.852 ms
 4  ABILENE.GIGAPOP.GEN.TX.US (198.32.236.14)  2.213 ms  2.260 ms  2.425 ms
 5  198.32.8.21 (198.32.8.21)  34.722 ms  34.373 ms  34.278 ms
 6  USC--abilene.ATM.calren2.net (198.32.248.85)  34.651 ms  34.705 ms  34.526 ms
 7  UCSD--USC.POS.calren2.net (198.32.248.34)  38.364 ms  38.095 ms  38.284 ms
 8  198.32.248.186 (198.32.248.186)  38.489 ms  38.819 ms  38.191 ms
 9  nodeB-ucsd-gw.ucsd.edu (132.239.255.141)  38.399 ms  38.286 ms  38.631 ms
10  nodeB-5500-6500-GE.ucsd.edu (132.239.255.149)  38.969 ms  41.245 ms  38.884 ms
11  cse-rs.ucsd.edu (132.239.254.45)  39.254 ms  39.132 ms  40.319 ms
12  freak.ucsd.edu (132.239.95.186)  39.194 ms  39.299 ms  39.676 ms
```

G Statistics

	anl-utk-8M		
	max	min	average
Xfer Size			
16K	0.40300	0.331250	0.377521
64K	0.782110	0.724233	0.763622
256K	2.041860	0.823960	1.551935
1M	3.820735	1.582323	2.998547
4M	6.075528	2.964298	4.455999
16M	7.172249	3.019620	4.323620
32M	5.590730	3.874349	4.711681

	anl-lsl-utk-8M		
	max	min	average
Xfer Size			
16K	0.329159	0.282999	0.324922
64K	0.821404	0.224252	0.778078
256K	2.358894	2.151340	2.260900
1M	6.806367	1.774656	6.182508
4M	15.474262	2.971618	12.317705
16M	22.538489	2.581877	11.362988
32M	24.572602	2.627871	11.447852

	ucsd-utk-8M		
	max	min	average
Xfer Size			
16K	0.335828	0.327585	0.332399
64K	0.671300	0.612586	0.666729
256K	1.770170	0.976332	1.643167
1M	3.208707	1.604943	2.577704
4M	4.738750	2.086455	3.742402
16M	6.517446	2.972433	4.154653
32M	5.236016	2.906003	4.031742

ucsd-lsl-utk-8M

Xfer Size	max	min	average
16K	0.223686	0.033156	0.207426
64K	0.602354	0.522283	0.588151
256K	1.801070	0.755655	1.742145
1M	5.168106	0.976739	4.173235
4M	9.254006	2.622432	4.622810
16M	14.895728	3.986542	6.343063
32M	13.465528	4.960214	6.395069

ucsd-utk-64K

Xfer Size	max	min	average
16K	0.336283	0.240713	0.315922
64K	0.670764	0.598264	0.655689
256K	1.608678	0.688810	1.478414
1M	2.913219	1.108043	2.557631
4M	3.615189	1.866385	3.074215
16M	3.529395	2.790902	3.117535
32M	3.546387	2.623760	3.170910

ucsd-lsl-utk-64K

Xfer Size	max	min	average
16K	0.217576	0.206132	0.214729
64K	0.599110	0.287751	0.563415
256K	1.720705	1.681558	1.692182
1M	3.533808	1.483824	3.257987
4M	4.855747	1.464077	4.042588
16M	5.212078	3.793412	4.471644
32M	4.994041	3.914145	4.554841