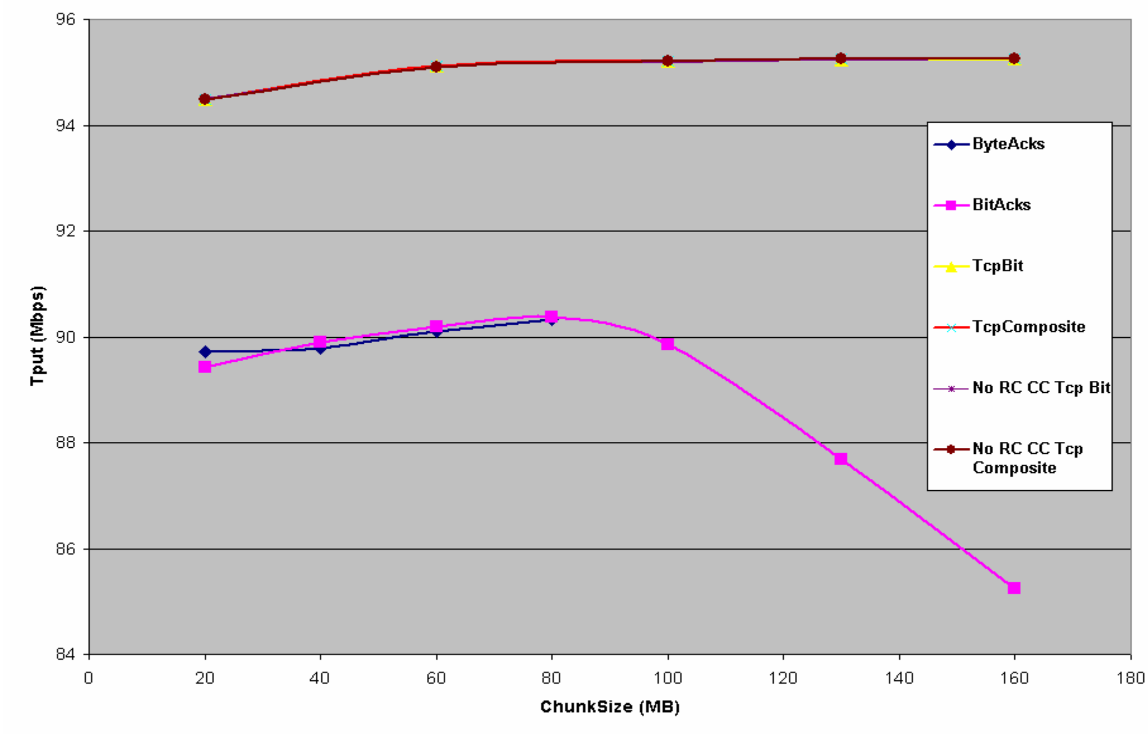

Figure 8 Throughput: ANL to NCSA.

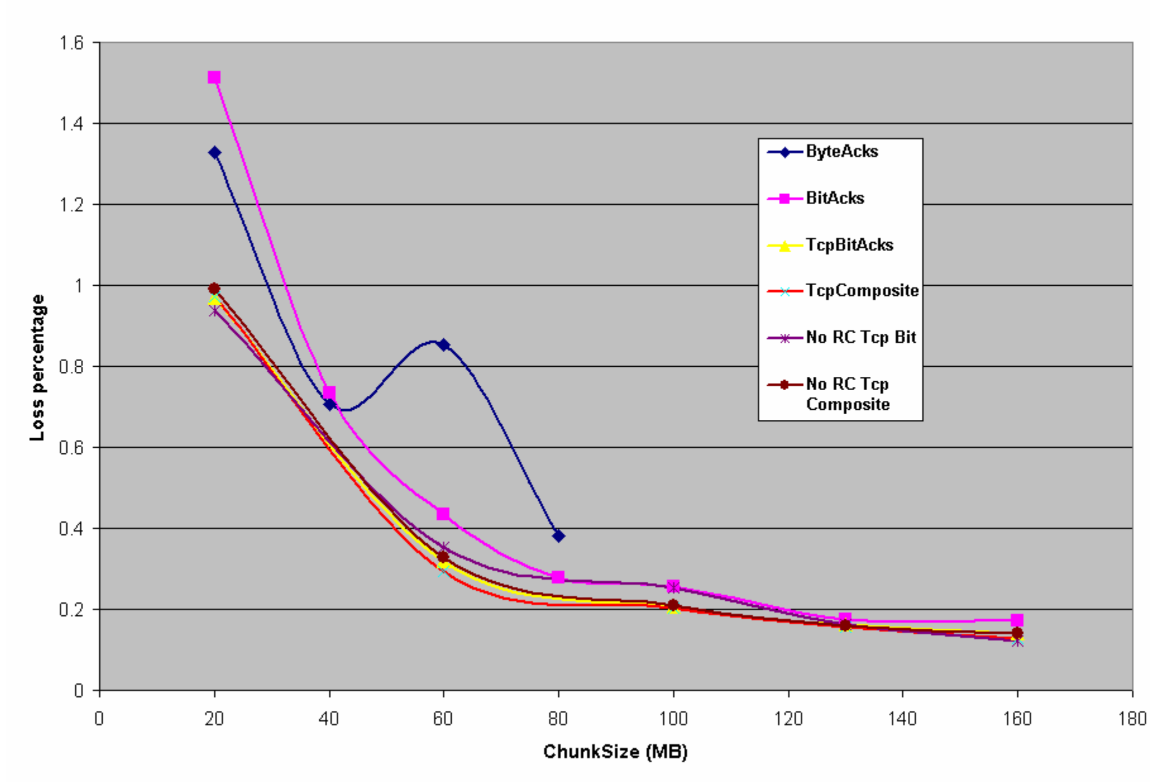

Figure 9 Loss: ANL to NCSA.

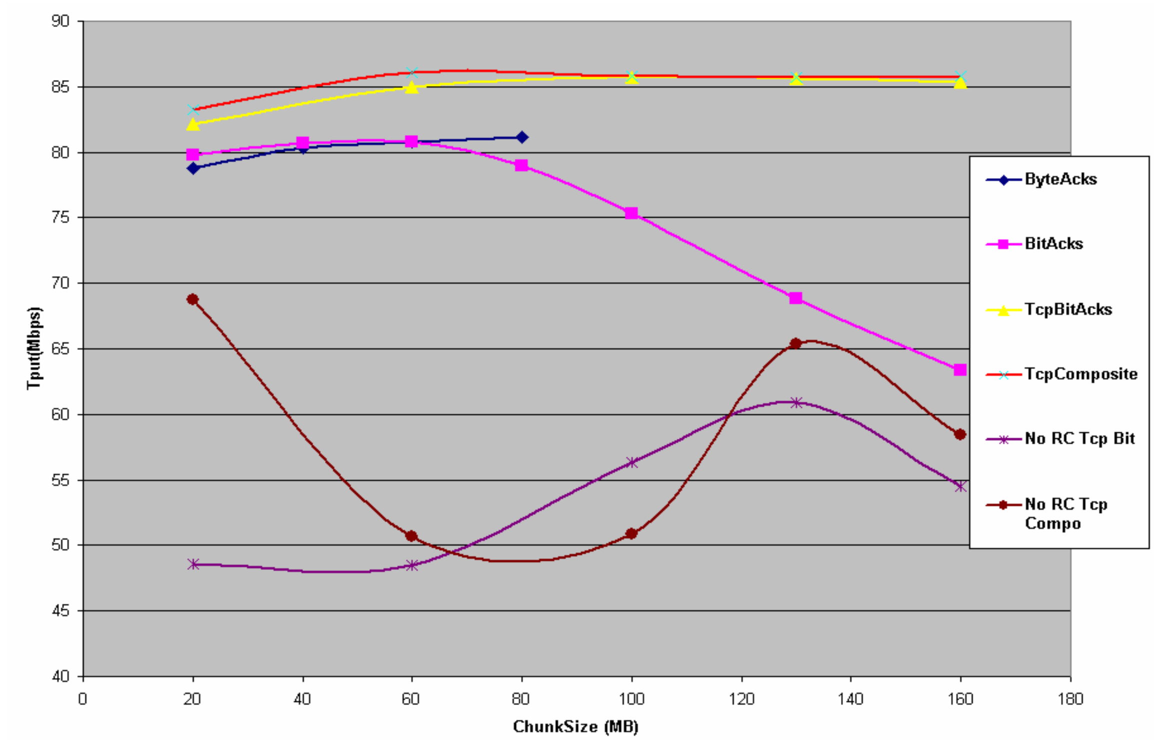

Figure 10 Throughput: NCSA to ANL.

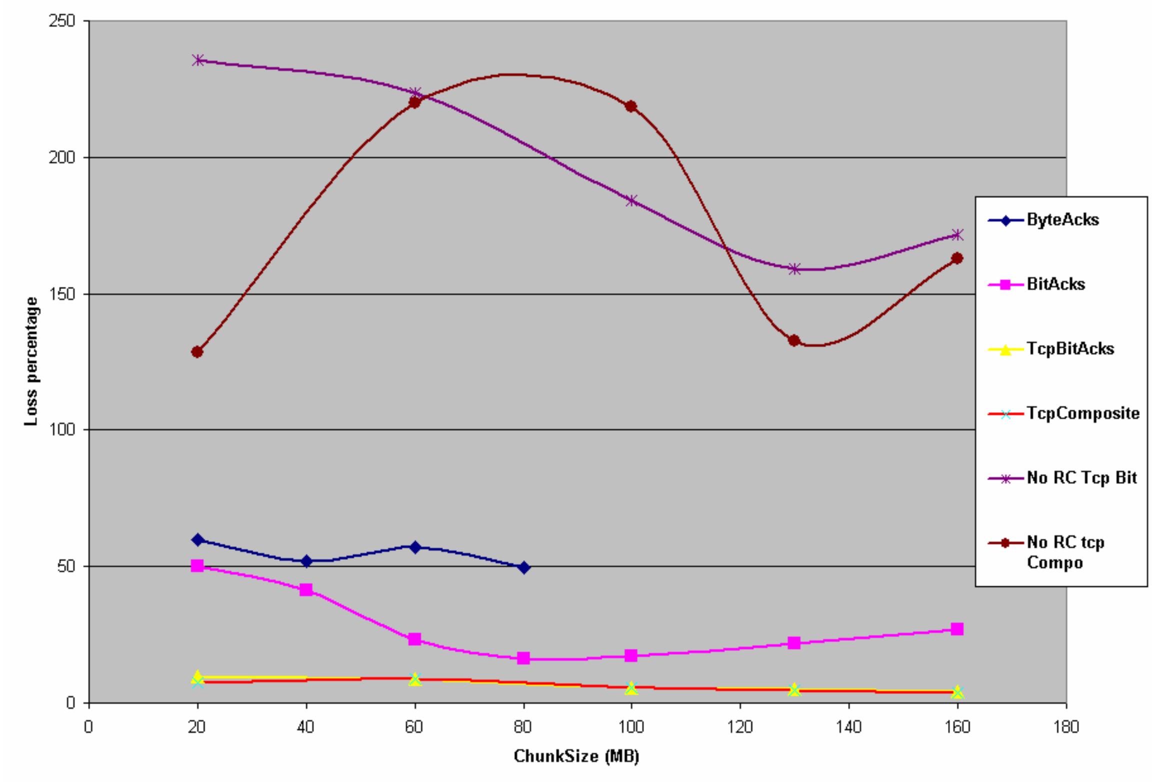

Figure 11 Loss: NCSA to ANL.

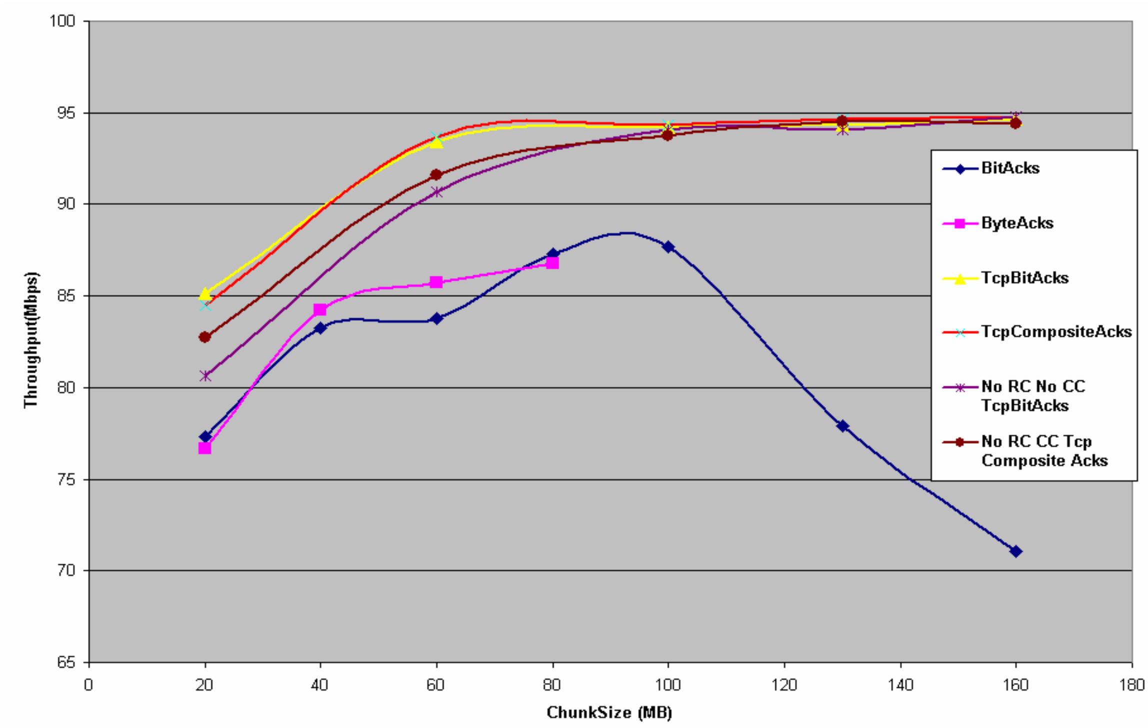

Figure 12 Throughput: NCSA to CACR.

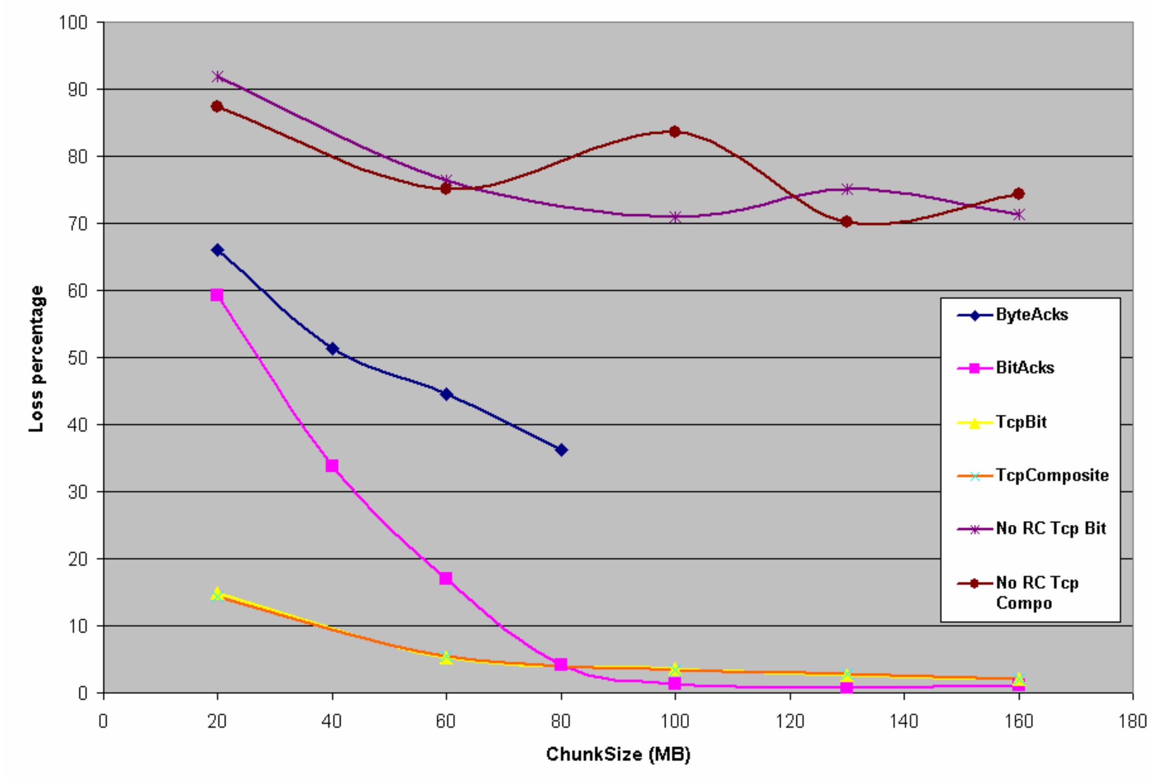

Figure 13 Loss: NCSA to CACR.

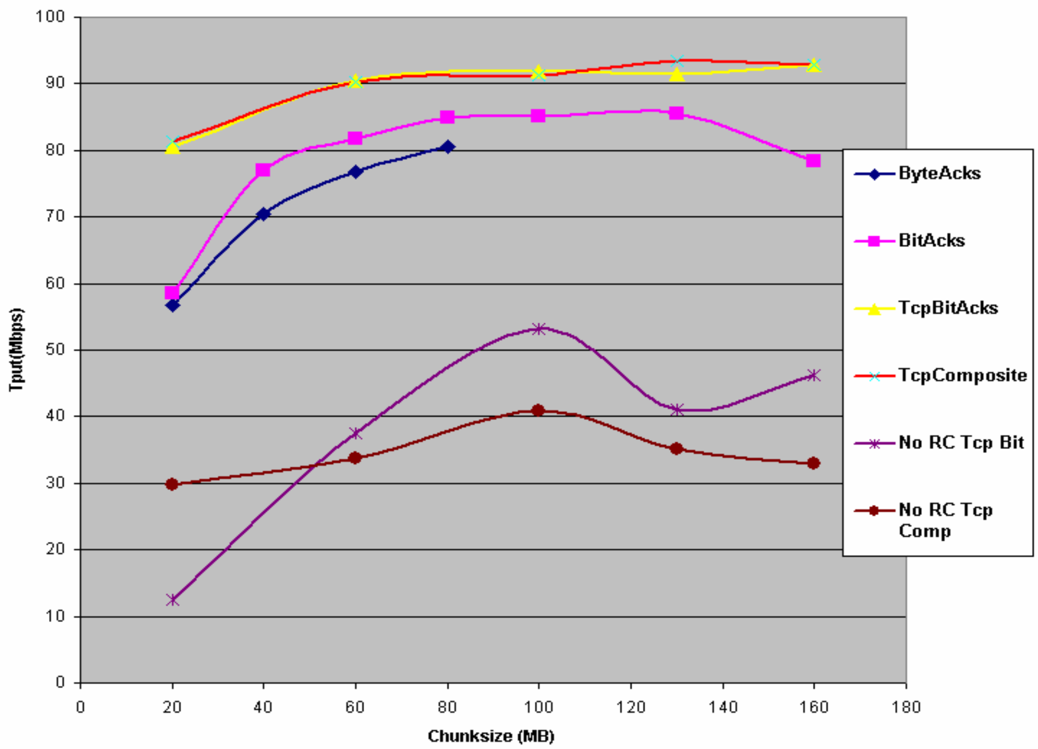

Figure 14 Throughput: CACR to NCSA.

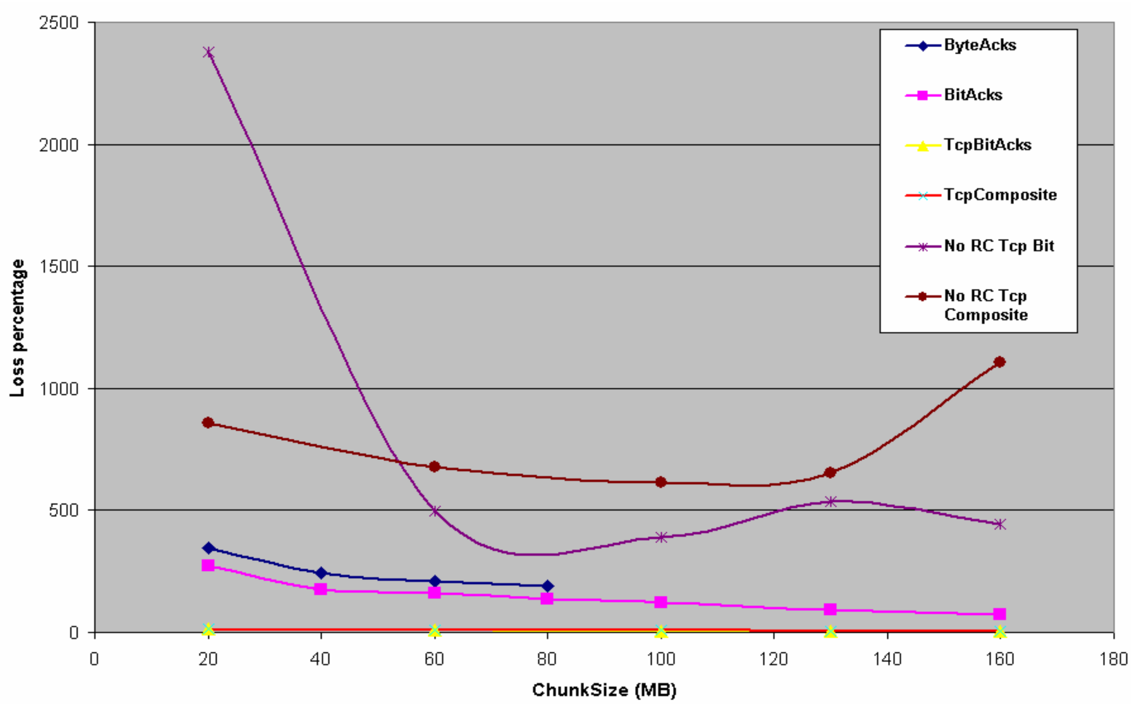

Figure 15 Loss: CACR to NCSA.

CHAPTER V

SUMMARY

5.1 Conclusions

This research work has shown that FOBS is a highly efficient, scalable, portable high-performance data transfer protocol. It is also portable and has been tested on all major platforms. The various mechanisms to ensure the reliability of transfer have been tested and proved to perform most efficiently. The rate and congestion control mechanisms have proved adequate for sharing the network resources with other applications and protocols.

Thus, as a result of the work described in this thesis, FOBS is deployable on the computational Grid, where a high-performance data transfer mechanism that utilizes a substantial percentage of the installed bandwidth and is able to co-exist with other protocols.

5.2 Recommendation for further studies

Future research is focusing on congestion control mechanisms. We are also collecting and analyzing the packet loss patterns to give us a system level understanding of the transfer dynamics. Various congestion control models, including an end-to-end system aware congestion model, are being developed that can be used as the network conditions dictate.

Another aspect to be considered is firewalls on the Grid. Most protocols are peer-to-peer, and thus firewalls are a major hassle that need to be overcome for deploying them successfully. This can be overcome by opening up specific reserved ports by the network administrator.

Appendix A

PERFORMANCE IMPLICATIONS OF USING UDP FOR DATA TRANSFER

A.1 Performance implications of choosing UDP for the data transport protocol

UDP provides a connectionless service for application-level procedures. Thus, UDP in itself is an unreliable service. However, the overhead of the protocol is low. A UDP packet consists of a header and payload. The header consists of a source port, destination port, field for length of the Datagram, and a field for checksum. The total overhead of the UDP packet is 8 octets (64 bits). FOBS itself adds a simple header of 4 octet length. Thus the total overhead of the FOBS data packet is 12 octets (56 bits) as compared to the TCP header size of 20 octets (160 bits). Thus, using UDP contributes to the better performance of FOBS as compared to data sent over TCP.

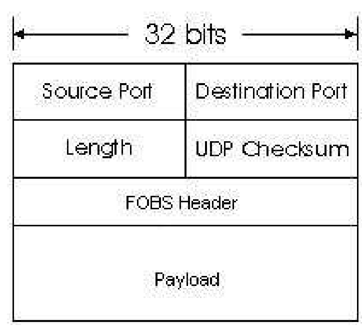

Figure 16 FOBS Data Packet

The payload size should be considered taking into account the underlying data link layer and the maximum transmission unit (MTU) on the network link between the sender and receiver. An MTU is the largest size of a packet that can be sent over a network without fragmenting it. There are many algorithms, as defined in [RFC1191], to determine the MTU. A packet, whose size is greater than the MTU size, will be fragmented into smaller packets, each not exceeding the MTU size. If any fragment thus created is lost, the rest of the fragments are discarded. Thus, the data packet should be as no larger than the MTU size. Thus, if the MTU were 1500 bytes (typical on a network link which includes Ethernet for some or all parts), the payload size of the FOBS packet would be 1488 bytes. For the same link and MTU, the size of the payload on a TCP segment would only be 1480 bytes.

BIBLIOGRAPHY

1 Abilene Network. *http://abilene.internet2.edu/*

2 Allcock and B. Bester and Bresnahan, J. and Chervenak, A. and Foster, I. and Kesselman, C. and Meder, S. and Nefedova, V. and Quesnet, D. and S. Tuecke. "Secure, Efficient Data Transport and Replica Management for High-Performance Data-Intensive Computing." Preprint ANL/MCS-P871-0201, Argonne National Laboratory. Feb 2001.

3 Andradas Veres and Miklos Boda. "The Chaotic Nature of TCP Congestion Control." In Proceedings of SC2000, November 2000.

4 Automatic TCP Window Tuning and Applications. *http://dast.nlanr.net/Projects*

5 BaBarCopy (BBCP) *http://www.slac.stanford.edu/~abh/bbcp/*

6 BabBarFTP (BBFTP) -- Large files transfer protocol. URL *http://doc.in2p3.fr/bbftp/*

7 Deepak Bansal and Hari Balakrishnan, "Binomial Congestion Control Algorithms,", Proceedings of IEEE INFOCOM 2001, April 2001.

8 Dickens, P. and Gropp, B. and P.Woodward. "High Performance Wide Area Data Transfers Over High Performance Networks." The 2002 International Workshop on Performance Modeling, Evaluation, and Optimization of Parallel and Distributed Systems. April 2002.

9 Dickens, P. and B.Gropp. "An Evaluation of Object-Based Data Transfers Over High Performance Networks." 11th High Performance Distributed Computing Conference. July 2002.

10 Enabling High Performance Data Transfers :Notes for Users and System administrators. *http://www.psc.edu/networking/perf_tune.html#intro.*

11 Feng, W. and P. Tinnakornsrisuphap. "The Failure of TCP in High-Performance Computational Grids." Super Computing 2000.

12 Feng, W., M. Fisk, M. Gardner and E. Weigle. "Dynamic Right Sizing: An automated Lightweight and scalable Technique for Enhancing Grid Performance." Proceedings of the 7th IFIP/IEEE International Workshop, PfHSN 2002, Berlin, April 2002

13 Globus Project. *http://www.globus.org*

14 Internet2. *http://www.internet2.edu*

15 Irwin, B. and M. Mathis. "Web100: Facilitating High-Performance Network Use." White Paper for the Internet2 End-to-End Performance Initiative. *http://www.internet2.edu/e2epi/*

16 Jacobson, V. and Braden, R. and D. Borman. "RFC 1323: TCP Extensions for high performance." *http://www.ietf.org/rfc/ rfc1323.txt?number=1323*

17 Leigh, J. and Yu, O. and Schonfeld, D. and Ansari, R. "Adaptive Networking for Tele-Immersion." Proceedings of the Immersive Projection Technology/Eurographics Virtual Environments Workshop (IPT/EGVE). May 2001.

18 List of sack implementations. *http://www.psc.edu/networking/all_sack.html.*

19 Mathis, M. and Mahdavi, J. and Floyd, S. and A. Romanow. "RFC 2018: TCP Selective Acknowledgment Options." *http://www.ietf.org/rfc/ rfc2018.txt?number=2018*

20 Ostermann, S. and Allman, M. and H. Kruse. "An Application-Level solution to TCP's Satellite Inefficiencies." Workshop on Satellite-based Information Services (WOSBIS). November, 1996.

21 R. Hobby. "Internet2 End-to-End Performance Initiative (or Fat Pipes Are Not Enough)." *http//www.internet2.org*

22 R. Venkat and P. Dickens. "WAMPI: An Alternative Protocol for High-Performance Wide Area MPI." M.S. Thesis, IIT, July 2001.

23 Sergey Gorinsky and Harrick Vin, "Additive Increase Appears Inferior." Tech Rep. TR2000-18, UT_Austin, May 2000.

24 Sivakumar, H. and Bailey, S. and R. Grossman. "PSockets: The Case for Application-level Network Striping for Data Intensive Applications using High Speed Wide Area Networks." Proceedings of Super Computing 2000 (SC2000).

25 Sivakumar, H. and Mazzucco, M. and Zhang, Q. and R. Grossman. "Simple Available Bandwidth Utilization Library for High Speed Wide Area Networks." Journal of Supercomputing, 2003.

26 The Very High Performance Backbone Service (VBNS). *http://www.vbns.net*

27 Tsunami: Experimental high speed network file transfer protocol. *http://www.anml.iu.edu/.*

28 Teragrid Project. *http://www.teragrid.org/*.

29 Y.Richard Yang, Min Sik Kim, and Simon S. Lam. "Transient Behaviors of TCP friendly Congestion Control Protocols." IEEE INFOCOM, April 2001.