Department of Energy DE-FG02-92ER25133 - Final Report
Title: The High Performance and Wide Area Analysis and Mining of Scientific & Engineering Data

PI: Dr. Robert Grossman
National Center for Data Mining/Laboratory for Advanced Computing
University of Illinois at Chicago

December, 2002

All of the research goals initially outlined in this proposal were met or exceeded. This final report summarizes our accomplishments and findings and includes recent publications occurring in the final period of this award.

Research Accomplishments

Overview. One of our research goals was to develop algorithms and services for remote data analysis and distributed data mining which scaled from the commodity internet to high performance networks. When we began the project there was no effective mechanisms to achieve high end to end performance for data intensive applications over wide area, high bandwidth networks. For this reason, we developed algorithms and services for Layers 2, 3, and 4 in the simple data web application stack below. We describe our research accomplishments for each of these layers in turn.

A Simple Application Stack for Data Webs	
Layer	Description
4	Data Web Applications
3	Data Web Services
2	Network Protocol Services
1	IP

Network Protocols for Data Intensive Applications.

It is currently an open research problem to develop protocols which enable high end to end performance for data intensive applications operating over wide area high performance networks. For example, a standard data mining application operating between Chicago and Amsterdam can process data at about 5 Mb/s due to the time required to acknowledge TCP packets.

Data Transfer between Chicago and Amsterdam			
Approach	standard TCP	striped TCP	reliable UDP
Implementation	Linux	PSockets	SABUL
Bandwidth Mb/s	4.36	400 (est.)	907.1

DOE Patent Clearance Granted
MPDvorscak 6·6·03
Mark P. Dvorscak Date
(630) 252-2393
E-mail: mark.dvorscak@ch.doe.gov
Office of Intellectual Property Law
DOE Chicago Operations Office

DISCLAIMER

1. We developed a software library called PSockets which stripes data across several network sockets so that applications can use effectively the bandwidth available on high performance networks [Grossman 2002a]. We also integrated PSockets with local parallel i/o libraries, such as MPIO, as well as local RAID devices. PSockets is an application level library which does not require tuning the TCP window, an alterative approach which can be relatively labor intensive. PSockets is a layer 2 service in the data web application stack.

2. We performed a variety of experimental studies using PSockets as the Layer 2 service for remote data analysis and distributed data mining applications deployed over OC-3 (155 Mb/s) and OC-12 (622 Mb/s) wide area networks [Grossman 2002b]. In particular, we found that a neural network could be used to predict accurately the near optimal number of parallel sockets required for a particular application on a particular network at a particular time.

3. We found that PSocket's performance began to level out after 10-30 sockets, even on OC-12 and higher bandwidth links. For this reason, we developed an alternative protocol called SABUL. The idea behind SABUL is simple. SABUL combines the UDP protocol in order to send data at a high rate with the TCP protocol in order to do this in a reliable fashion. UDP has no flow control, rate control, or reliable transmission mechanisms. SABUL implements these control functions in a separate TCP control channel. SABUL is also an application library like PSockets, and, hence does not require tuning the network. This approach is in contrast to the approach of other high performance protocols such as NETBLT [Clark] which combine the data and control channels. SABUL is a layer 2 service in the data web application stack.

Data Webs and Data Web Services. Data webs which are web based infrastructures for data [Grossman 2002c]. Unlike grid middleware which is built over authentication, authorization and access (AAA) control mechanisms for rationing and scheduling presumably scarce high performance computing resources [Foster] data webs are built using W3C standards and emerging standards for web services and packaging (SOAP and XML). Data webs in contrast to data grids are designed to encourage the open sharing of data resources without AAA controls, in the same way that the web today encourages the sharing of document resources without AAA controls [Semantic Web].

The table above from [Grossman 2002d], lists several technologies for exploring remote and distributed numerical data: data webs [Grossman 2002c], data grids [Datagrid], and semantic webs [Semantic Web] can all provide web based access to numerical data. Data webs provide direct access to distributed rows and columns of data. Data grids enable large scale resource sharing of computational and data resources. Semantic webs provide knowledge based access to data using ontologies, RDF and agent based architectures.

	View	Mine/Discover	Compute
Knowledge	Digital Libraries	Knowledge Mining	Semantic Webs
Rows/Columns	Web Databases	Data Webs	Data Grids
Files	Persistent Archives	Dist. Data Mining	Grids

1. We developed a protocol for moving data between data web servers and data web clients called the Data Web Transfer Protocol or DWTP. An earlier name for DWTP was the DataSpace Transfer Protocol or DSTP. DWTP includes functions for retrieving data, metadata, and keys, for selecting rows and columns of data, and for sampling [Grossman 2002e]. DWTP is a layer 3 service in the data web application stack.

2. We designed and implemented a DWTP Server [Grossman 2002e] for the remote analysis and distributed mining of data over the commodity internet. This DWTP Server uses TCP as the layer 2 service and is designed for a single server.

3. We designed and implemented a DWTP Server [Grossman 2000]for the remote analysis and distributed mining of data over wide area, high performance networks. This DWTP Server uses SABUL as the layer 2 service and can be deployed over clusters of workstations.

4. We designed and created a reference implementation of a second version of DWTP. Version 2.0 of DWTP supports vector valued keys and missing data, two enhancements which turned out to be critical for some of the applications described below [Grossman 2002f].

Data Web Applications.
We have developed data web applications for astronomical, earth sciences, bioinformatics, and social science data.

1. We developed a data web application showing how data webs can be used to create virtual observatories. For this application we showed how a DWTP client could overlay distributed data, with one data source a DWTP server containing 2MASS data and another data source a DWTP server containing DPOSS data [Grossman 2001].

2. We developed a data web application containing simulation data from the Community Climate Model 3 maintained by NCAR, as well as related earth science data, such as El Nino anomaly data. Using this application, scientists from other domains can easily overlay CCM3 data with data from their own applications [Grossman 2002e] and [Grossman 2002g]. As another example, this DWTP server can be used by someone with cholera data to look for correlations between El Nino and Cholera, a relationship reported by some investigators.

3. We developed several data web applications for sequence, proteonomic, and chemical compound data. In particular, we show how a chemical compound from one DWTP server could be docked with a protein from another DWTP server [Grossman 2002h].

Infrastructure. Data webs became truly useful as more and more data is added to them and as more and more software interoperates with them. For this reason, part of our effort has been focused on infrastructure issues, including developing standards, writing open source software, and operating testbeds.

1. We have put a variety of data sets on DWTP servers for public consumption, including some NCAR CCM 3 data, data from the Protein Data Bank (PDB), chemical compound data from the National Cancer Institute (NCI), and some data from the US Census. This data is maintained at the site www.dataspaceweb.net.

2. We help develop standards for data mining. In particular, we are active members of the Data Mining Group, which is a collaboration of over 12 vendors developing an XML standard called the Predictive Model Markup Language (PMML) for data mining [Grossman 2002c]. In particular, our group developed reference implementations of PMML Versions 0.8, 0.9 and 1.0. PMML is the interchange language used by data web applications to exchange data mining models. The home for PMML is the site www.dmg.org which we maintain.

3. As mentioned above, we developed a protocol called the Data Web Transfer Protocol enabling data web clients to communicate efficiently with data web servers. We are currently preparing an IETF draft for this emerging standard.

4. We have operated two data mining testbeds for high performance and distributed data mining: one called the Terabyte Challenge Testbed [Grossman 2002i]and one called the Terra Wide Data Mining Testbed [Grossman 2002j]. Over the past several years, we have set a number of records using these testbeds. For example, in November, 2002, a cluster of three DWTP servers using SABUL as the Layer 2 service set a record for moving data across the Atlantic at 2.7 Gb/s. This outperformed a number of other protocols tested at the same event (iGrid 2002), including GridFTP. For comparison, data mining applications using TCP move data at about 5 Mb/s.

5. We maintain open source reference implementations of two DWTP servers using Source Forge. One is targeted at data webs built over the commodity internet and one at data webs built over high performance networks. DWTP is our Layer 3 service.

6. We maintain an open source reference implementation of SABUL on Source Forge. SABUL is one of our Layer 2 services.

Technical Publications
2002

1. Robert Grossman, and Marco Mazzucco, DataSpace - A Web Infrastructure for the Exploratory Analysis and Mining of Data, IEEE Computing in Science and Engineering, July/August, 2002, pages, 44-51.

2. Robert Grossman, Mark Hornick, and Gregor Meyer, Data Mining Standards Initiatives, Communications of the ACM, Volume 45-8, 2002, pages 59-61.

3. Robert Grossman, Mark Hornick, and Gregor Meyer, Emerging Standards and Interfaces in Data Mining, Handbook of Data Mining, Nong Ye, editor, Kluwer Academic Publishers, to appear.

4. Marco Mazzucco, Asvin Ananthanarayan, Robert L. Grossman, Jorge Levera, and Gokulnath Bhagavantha Rao, Merging Multiple Data Streams on Common Keys over High Performance Networks, Proceedings of Supercomputing 02.

5. Donald Hamelberg, Pavan K. Kasturi, and Robert L. Grossman, Data Webs for Bioinformatics Data, submitted for publication.

2001

1. H. Sivakumar, R. L. Grossman, M. Mazzucco, Y. Pan, Q. Zhang, Simple Available Bandwidth Utilization Library for High-Speed Wide Area Networks, Journal of Supercomputing, to appear.

2. H. Sivakumar, M. Mazzucco, and R. L. Grossman, Effect of Transmission Control Protocol on Network Striping over High Speed Wide Area Networks, submitted for publication.

3. Asvin Ananthanarayan, Rajiv Balachandran, Yunhong Gu, Robert Grossman, Xinwei Hong, Jorge Levera, Marco Mazzucco, Data Webs for Earth Science Data, submitted for publication.

4. R. Hollebeek and R. L. Grossman, The National Scalable Cluster Project: Three Lessons about High Performance Data Mining and Data Intensive Computing, in Handbook of Massive Data Sets, J. Abello, P. M. Pardalos, and M. G. C. Resende, editors., Kluwer, 2002.

5. R. L. Grossman and Yike Guo, Parallel Methods for Scaling Data Mining Algorithms to Large Data Sets, Handbook on Data Mining and Knowledge Discovery, Jan M Zytkow, editor, Oxford University Press, 2002.

6. R. L. Grossman, E. Creel, M. Mazzucco and Roy Williams, A DataSpace Infrastructure for Astronomical Data, in Scientific Data Mining, Kluwer Academic Publishers, 2001.

REFERENCES

[Clark] D. Clark, M. Lambert and L. Zhang, "NETBLT: A high throughput transport protocol", Frontiers in Computer Communications Technology: Proceedings of the ACM-SIGCOMM '87, Association for Computing Machinery, Stowe, Vermont, 353—359, 1987

[Datagrid] Datagrid, "Globus Data Grid", Retrieved from http://www.globus.org/datagrid/, September 2, 2002.

[Foster] I. Foster and C. Kesselman, "The Grid: Blueprint for a New Computing Infrastructure", "he Grid: Blueprint for a New Computing Infrastructure, Morgan Kaufmann, San Francisco, California, 1999.

[Semantic Web] W3C Semantic Web, Retrieved from www.w3.org/2001/sw/, September 2, 2002.

[Grossman 2000] S. Bailey, E. Creel, R. Grossman, S. Gutti and H. Sivakumar, "A High Performance Implementation of the Data Space Transfer Protocol (DSTP)", Large-Scale Parallel Data Mining, editors M. J. Zaki and C.-T. Ho, Springer-Verlag, Berlin, 2000, pages 55-64.

[Grossman 2001] Robert Grossman, Emory Creel, Marco Mazzucco and Roy Williams, editors R. L. Grossman, C. Kamath, W. Philip Kegelmeye, V. Kumar and R. Namburu, A DataSpace Infrastructure for Astronomical Data, Kluwer Academic Publishers, 2001, series Data Mining for Scientific and Engineering Applications.

[Grossman 2002a] Marco Mazzucco, Asvin Ananthanarayan, Robert L. Grossman, Jorge Levera and Gokulnath Bhagavantha Rao, "Merging Multiple Data Streams on Common Keys over High Performance Networks", Proceedings of Supercomputing 2002, IEEE and ACM, 2002.

[Grossman 2002b] H. Sivakumar, Marco Mazzucco and Robert L Grossman, "Effect of Transmission Control Protocol on Network Striping over High Speed Wide Area Networks", Submitted for publication.

[Grossman 2002c] R. L. Grossman, Standards and Infrastructures for Data Mining, Communications of the ACM, {45}{8}, 2002, pp.45-48.

[Grossman 2002d] Robert L. Grossman, "Some Frequently Asked Questions about DataSpace", Submitted for publication.

[Grossman 2002e] Robert Grossman and Marco Mazzucco, "DataSpace - A Web Infrastructure for the Exploratory Analysis and Mining of Data", IEEE Computing in Science and Engineering, July/August, 2002, pgs. 44—51.

[Grossman 2002f] Robert Grossman and Marco Mazzucco, "The DataSpace Transfer Protocol (DSTP): A Data Transport Protocol for Data Webs, Submitted for publication.

[Grossman 2002g] Asvin Ananthanarayan, Rajiv Balachandran, Yunhong Gu, Robert Grossman, Xinwei Hong, Jorge Levera and Marco Mazzucco, Data Webs for Earth Science Data, Submitted for publication

[Grossman 2002h] Donald Hamelberg, Pavan K. Kasturi and Robert L. Grossman, "Data Webs for Bioinformatics Data", Submitted for publication.

[Grossman 2002i] R. Hollebeek and R. L. Grossman, The National Scalable Cluster Project: Three Lessons about High Performance Data Mining and Data Intensive Computing, Handbook of Massive Data Sets, editors J. Abello, P. M. Pardalos and M. G. C. Resende, Kluwer Academic Publishers, 2002.

[Grossman 2002j] Robert L. Grossman, Yunhong Gu, Dave Hanley, Xinwei Hong, Dave Lillethun, Jorge Levera, Joe Mambretti, Marco Mazzucco and Jeremy Weinberger, "Photonic Data Services: Integrating Path, Network and Data Services to Support Next Generation Data Mining Applications, Proceedings of the NSF Workshop on Next Generation Data Mining, 2002.