Optimizing Data Transfer Speeds of the TeraGrid Computers

Jarett Collins
Dillard University and Fisk University

Research Alliance in Math and Science
Computer Science and Mathematics Division, Oak Ridge National Laboratory
Mentors: Dr. John W. Cobb and Gregory G. Pike

http://www.csm.ornl.gov/Internships/rams_06/abstracts/j_collins.pdf

Abstract

The advent of high speed network backbones such as the TeraGrid has brought the promise of ubiquitous computing closer to a reality. Users are less constrained by the overhead of transferring data between computational resources and are free to move between different platforms to select the one most appropriate for a particular stage of computation. Nevertheless, we are still not utilizing the full potential of our network bandwidth. Detailed exploration of network transport characteristics for LUSTRE and other common TeraGrid data transport protocols is shown in this poster. The general network throughput characteristics for these protocols and possible areas of improvement for general data transport are given.

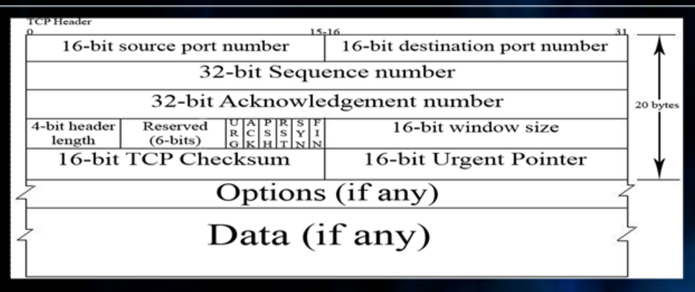

Fig 1. TCP Header

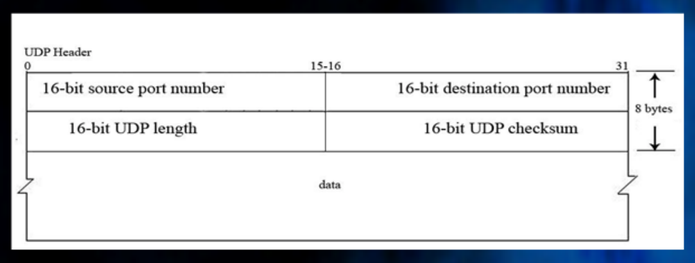

Fig 2. UDP Header

Fig 3. Cray XT3

Fig 4. TeraGrid Networks

Goals/Tasks

- analyze data transfer speed of TeraGrid computers
- analyze TCP dumps
- create a C/C++ program that
- parses essential data
- calculates bandwidth
- create graph dumps from different data transfer tools
- compare results of data transfer tools to determine which has superlative results

Processes

Research application level protocols

- LUSTRE
- Uberftp
- Sftp

C++ program/parser

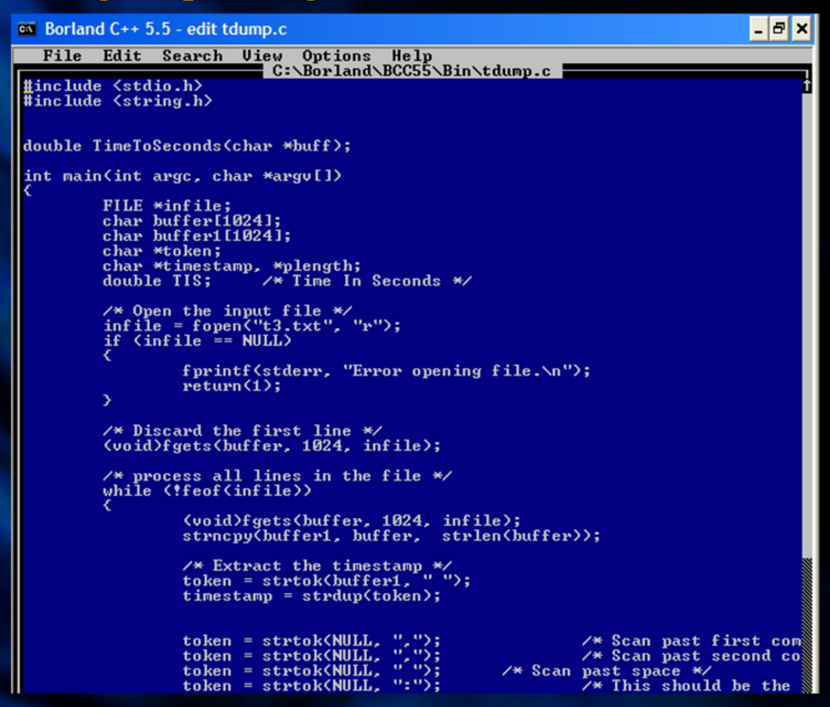

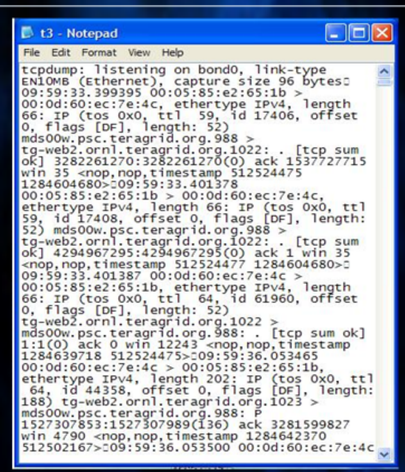

Before execution

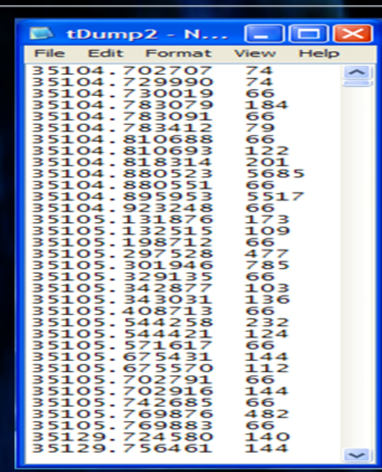

After execution

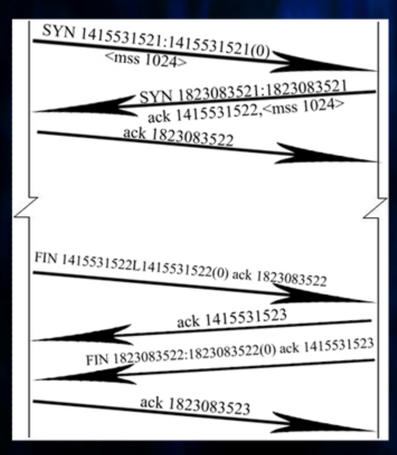

Fig 5. Example timeline of connection establishment and termination

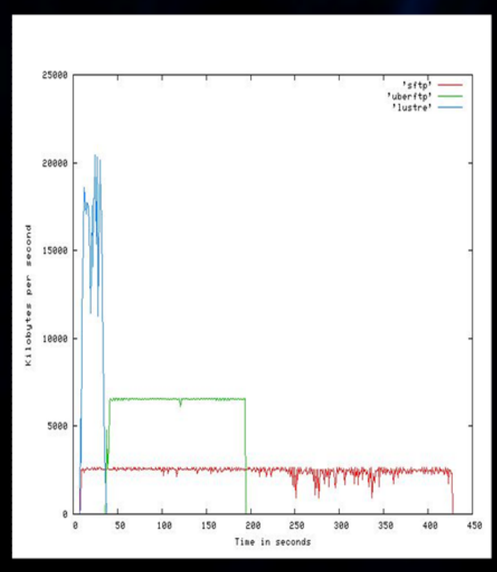

Results

Analysis of graph determines that LUSTRE is a better data transfer tool for the TeraGrid.

Prospective Research

- further research of other data transfer protocols
- implement other data transfer tools such as SABUL, Tsunami, UDT

The Research Alliance in Math and Science program is sponsored by the Mathematical, Information, and Computational Sciences Division, Office of Advanced Scientific Computing Research, U.S. Department of Energy. The work was performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725. This work has been authored by a contractor of the U.S. Government, accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

A special thanks goes out to faculty advisor Dr. Stephen Egarievwe for helping the author get this internship. Many thanks goes out to my mentors Dr. John W. Cobb and Gregory G. Pike for their efforts and for taking time out of their very busy schedule to help get data for this project. Finally, special thanks goes to Debbie McCoy, who made this research a possible and pleasant experience.

Office of Science U.S. DEPARTMENT OF ENERGY

TeraGrid EMPOWERING DISCOVERY

UT-BATTELLE

OAK RIDGE NATIONAL LABORATORY
U.S. DEPARTMENT OF ENERGY