Protocol Wars

an exploration of PlanetLab, Internet2 and the UDT protocol

Jule Slootbeek
April 27th, 2005
Sponsor: Prof. Jerry Breecher

Introduction

- Idea & Motivation
- Protocol Outlines
- Internet2
- Abilene Network
- PlanetLab
- Methodology
- Results
- Conclusion
- Future

Idea & Motivation

- To compare TCP to new Developmental Protocols, i.e. UDT
- To learn about network protocols.
- To learn more about the planet lab network and its backbone.
- Determine whether is it plausible to continue development on UDT.
- Determine in what scenarios UDT is favored over TCP and vice versa.

Transmission Control Protocol

- ACK for every packet sent
- Good over short distances
- Good over short distance/low-bandwidth connections
  - Dial-up
  - {A}DSL
  - Cable
  - LAN

UDP-based Data Transport Protocol

- Time based ACK
- NACK when packet not received
  - Queued, resent at a later time
- Good over high-bandwidth/long distance connections
  - Optical networks
  - Global networks

- Consortium of 200+ universities, and 60+ companies to develop and deploy advanced network applications and technology.
- Abilene Backbone
  - 10 gigabits per second cross-country
  - 100 megabits per second between every Abilene connected desktop
- Supporters of National LambdaRail

Abilene Network Map

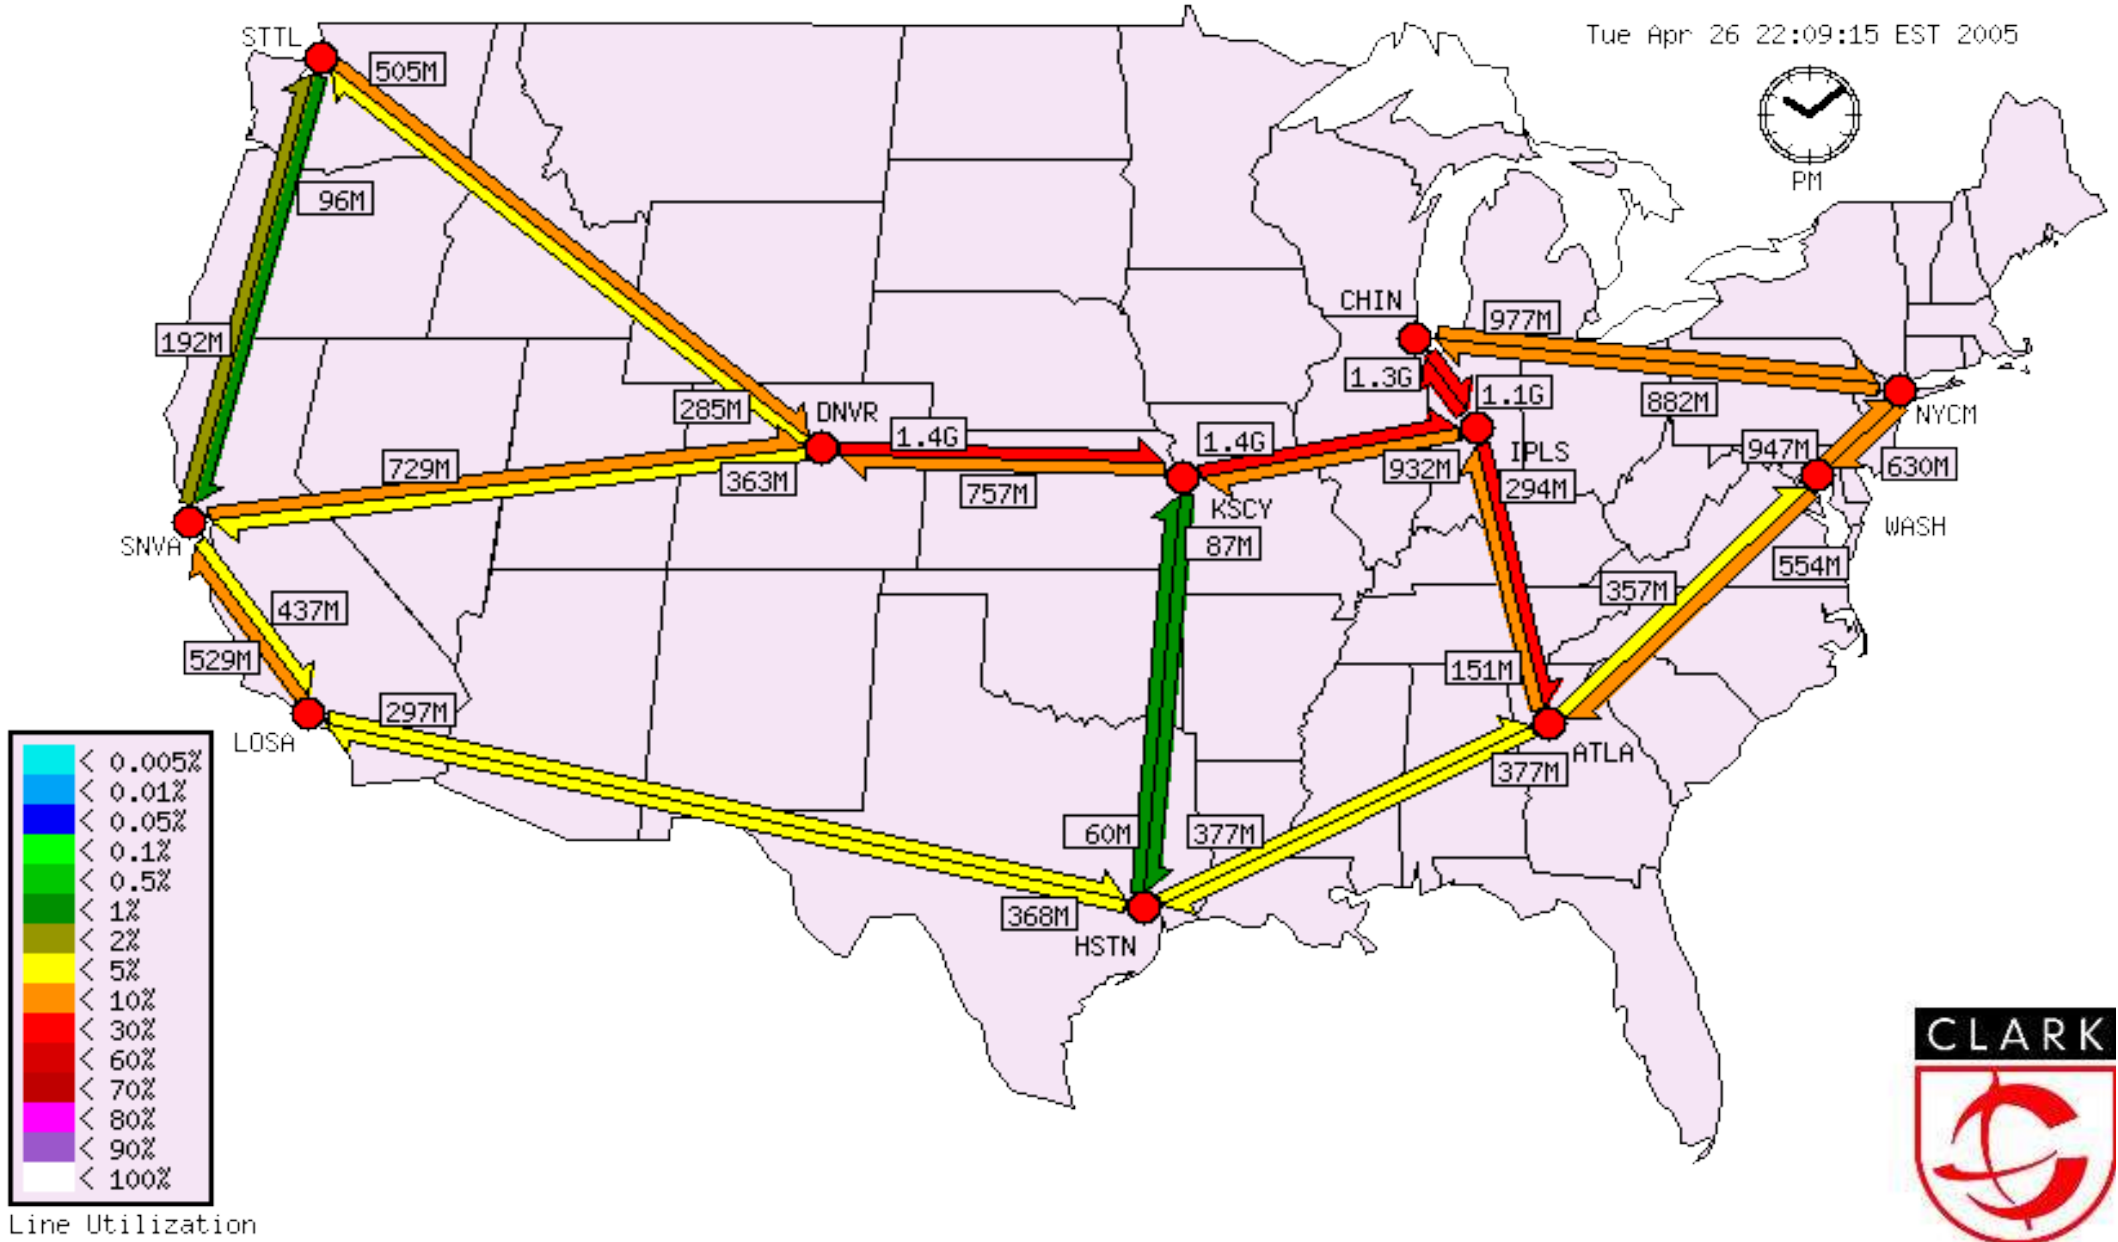

Abilene Network

- US high-performance backbone created by I2
- Created in 1999 with 2.5 gigabits/sec
- Expanded in 2004 with 10 gigabits/sec
- Aim to achieve 100 gigabits/sec between every node by the end of 2006
- Controversy around file swapping using i2hub
- LambdaRail
  - Transcontinental Ethernet Computer Network
  - Regional Optical Networks providing OC-192
    - OC-192: fiber optic line with SONET rate of 9953.28 megabit/sec

Synchronous Optical Networking

- Standard for digital communication over optical fiber in the US and Canada
- Synchronous Digital Hierarchy
  - Rest of the World
  - Main focus, SONET as a variation
- OC-1 -> OC-3072
  - 51 640 kb/sec -> 159 292 480 kb/sec

PLANETLAB

An open platform for developing, deploying, and accessing planetary-scale services

- Sites (262):
  - Physical Location of PlanetLab nodes
- Nodes (558):
  - Server that runs PlanetLab components
- Slices
  - Set of allocated resources
  - UNIX shell access to resources

CLARK

Quick PlanetLab History

- 2002
  - Started by Intel Research – Berkeley, and a whole bunch of major universities
  - 100 nodes
- 2003
  - HP joins the project.
  - NSF grants $4.5M for the project.
  - Reaches 200 node mark in December
- 2004
  - Version 2.0 is deployed
  - CANADRIE and RNP, added to the network
  - 400 node mark passed in July
  - CERNET joins PlanetLab
  - 500 node mark passed in December

Methodology

- Send 1G of data over four (4) scenarios
- Measure number of bytes received every second
- TCP
  - Use iperf
- UDT
  - Use UDT flooder

Bandwidth Comparison

- TCP: Used iperf to send 1G of data
- UDT: Used UDT-flooder to send 1G of data
- Ran every test five (5) times
- Four (4) scenarios
  - {TCP, UDT} over Internet2 backbone
    - `From: planetlab1.wash.internet2.planet-lab.org`
    - `To: planetlab1.atla.internet2.planet-lab.org`
  - {TCP, UPT} over a non-Internet2 backbone
    - `From: planetlab1.millennium.berkeley.edu`
    - `To: planetlab1.rn.informatics.scitech.susx.ac.uk`

Predicted Behavior

- UDT does not wait for ACK after every packet
- Lot of packages on the pipe
- Secure PlanetLab connection
  - Few hops
  - Fast connections
- TCP will wait for ACK
  - Slow down the process

Bandwidth Comparison (results)

- TCP over non-Internet2 Backbone
  - Average time ~1 hour 42 minutes (6140 seconds)
- UDT over non-Internet2 Backbone
  - Average time ~15 minutes (904 seconds)
- TCP over Internet2 Backbone
  - Average time ~9 minutes (580 seconds)
- UDT over Internet2 Backbone
  - Average time ~3 minutes (218 seconds)

planetlab1.wash.internet2.planet-lab.org -> planetlab.wash.internet2.planet-lab.org

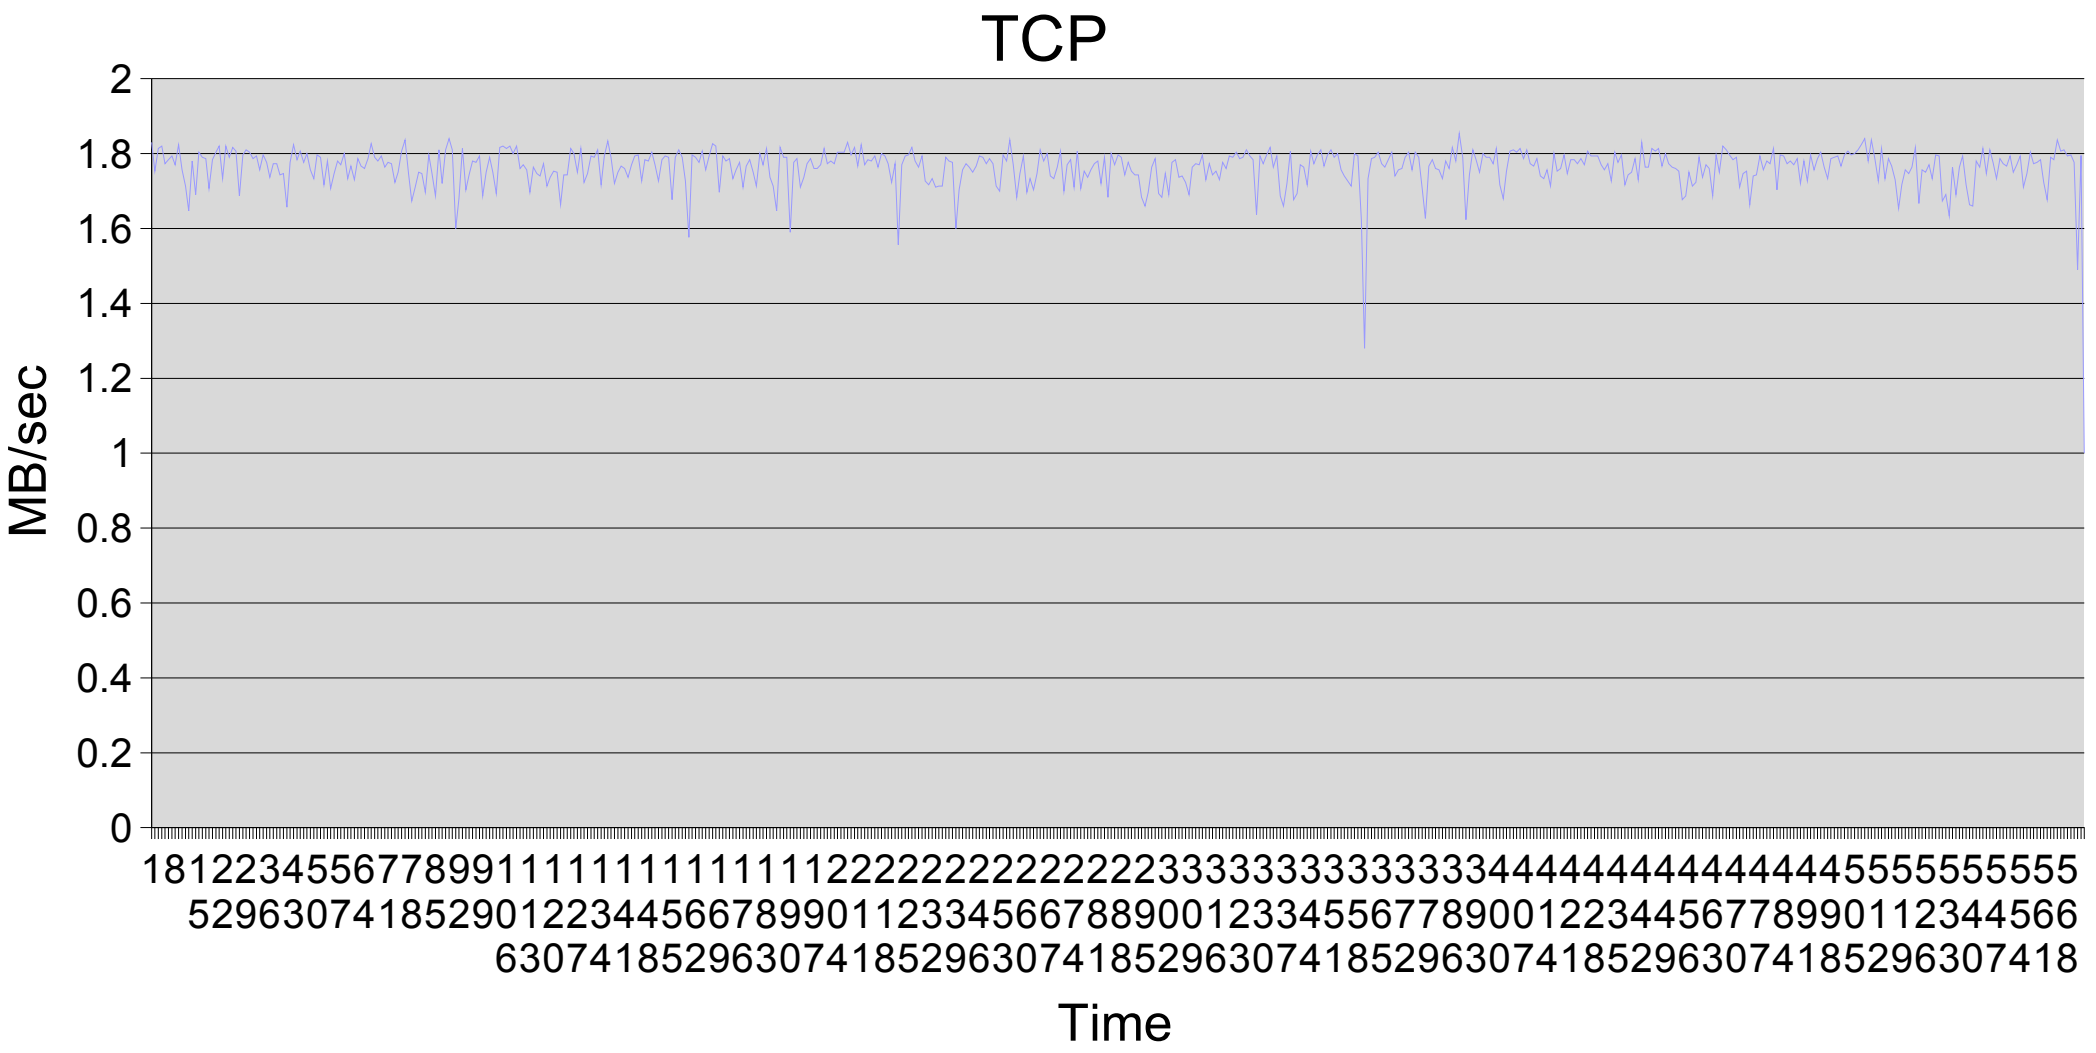

planetlab1.wash.internet1.planet-lab.org -> planetlab1.atla.internet1.planet-lab.org

UDT

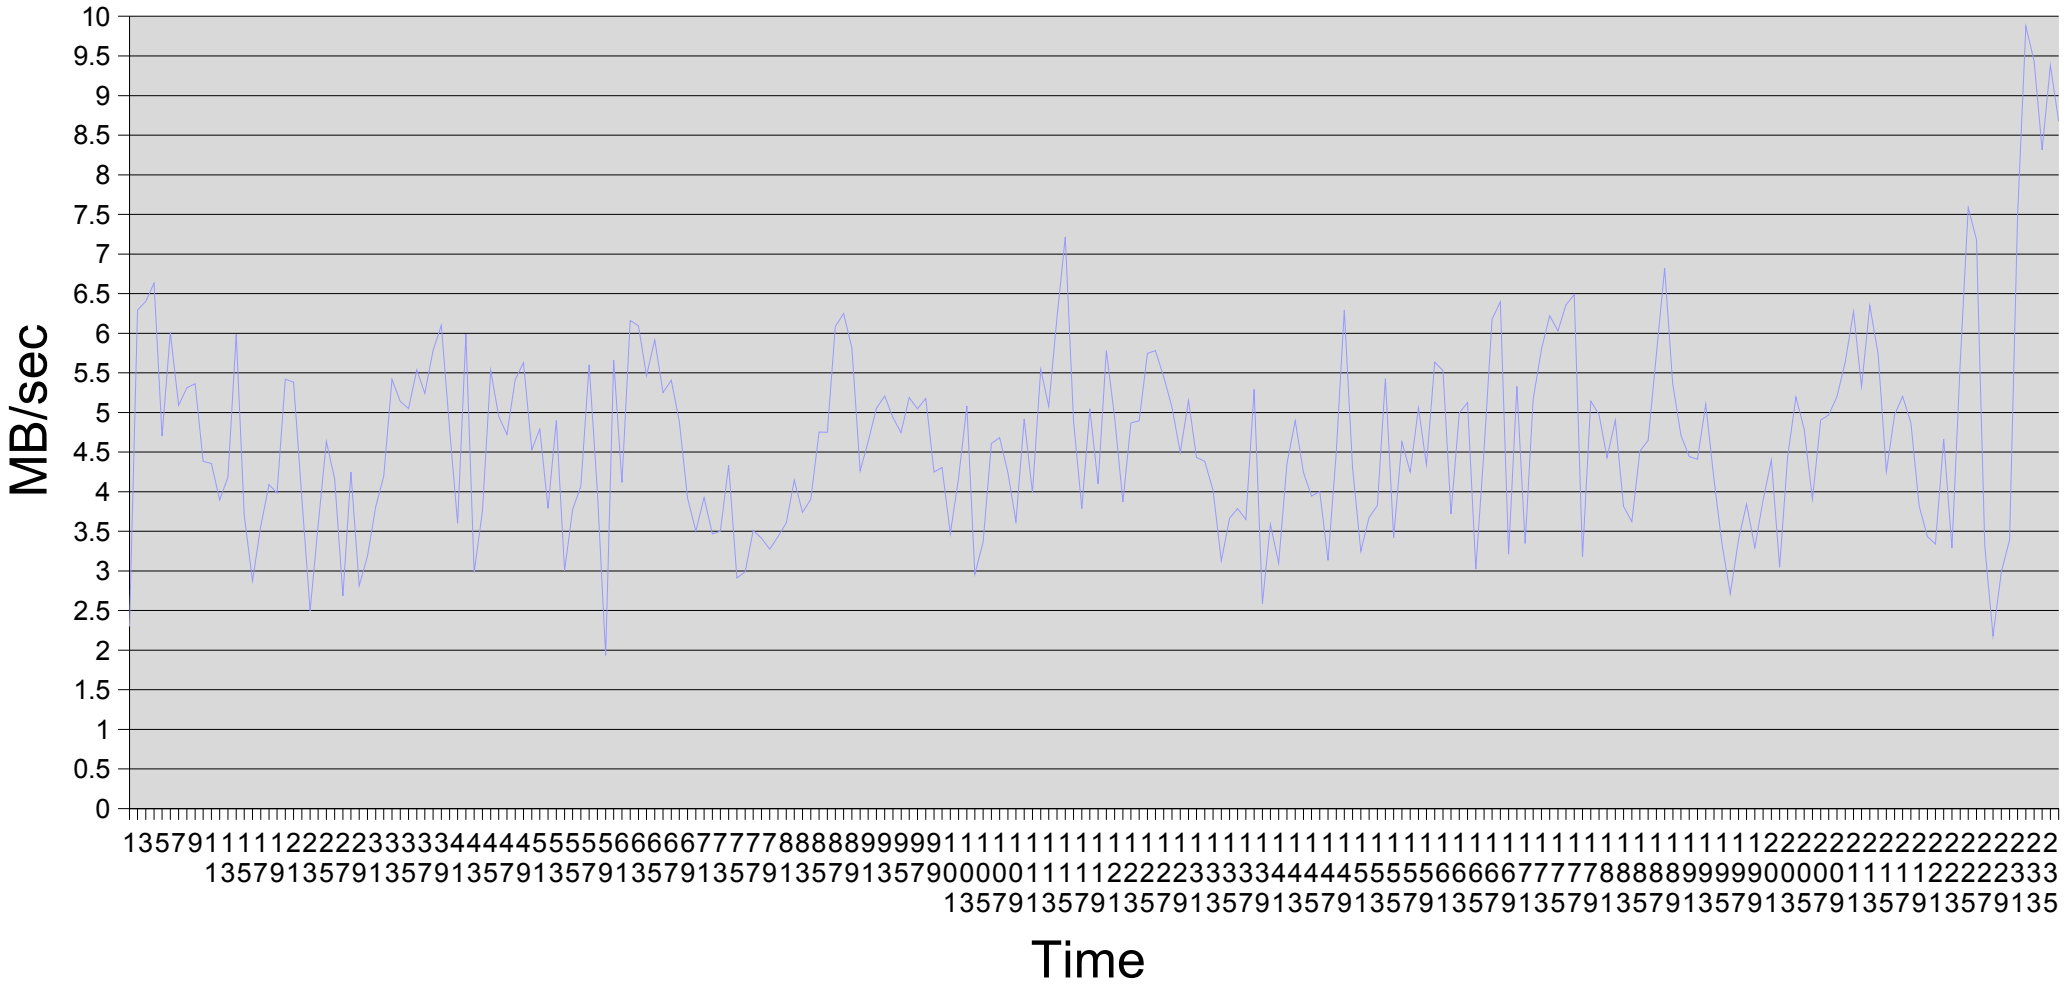

Conclusions

- TCP reaches maximum speeds on a 10/100 line. TCP and UDT perform evenly well in that scenario
- UDT reaches higher speeds over faster lines, since there are more packages out on the pipe
- UDT not good for household computing
- UDT optimal for large transfers over high speed networks

Future

- More detailed tests
- User space testing suite
  - Collision, outstanding packets, etc.
- Add UDT functionality to commonly used software applications
  - File Transport Protocol (FTP)
  - Peer2peer networks
  - Streaming Media
- Expand UDT usage on high performance computing applications

References

- Official UDT website, http://udt.sf.net
- Official PlanetLab website, http://www.planet-lab.org
- Official Internet2 website, http://www.internet2.org
- Wikipedia, http://en.wikipedia.org
- Andrew S. Tanenbaum. *Computer Networks,* Prentice Hall, Upper Saddle River, NJ, 1996.
- Thomas F.Herbert. *The Linux TCP/IP Stack: Networking for Embedded Systems,* Charles River Media, Hingham, MA, 2004
- Y. Gu and R. Grossman. *UDT: A Transport Protocol for Data Intensive Applications.* Internet Draft, work in progress.
- Y. Gu, Xinqei Hong and Robert L. Grossman. *Experiences in Design and Implementation of a High Performance Transport Protocol*, SC 2004, Nov 6-12, Pittsburgh, PA