46PaQ



Dimitris Miras, Saleem Bhatti, Peter Kirstein
Networks Research Group
Computer Science
UCL

Today's talk

- Overview
- Current Status and Results
- Future Work
- Questions

Overview

46PaQ overview [1]

- IPv**4** & IPv**6** **P**erformance **a**nd **Q**oS
  - http://www.cs.ucl.ac.uk/staff/S.Bhatti/46paq/
  - 01 Dec 2004 - 30 Nov 2006
- Sites:
  - Computer Science, UCL
  - Computer Science, St. Andrews (1 Jan 2006)
  - Computer Science, QML (1 Sep 2005)

46PaQ overview [2]

- **Network and transport protocols at multi-Gb/s:**
  - UCL leads (St. Andrews from 1 Jan 2006)
  - New protocols (e.g. DCCP, TCP ‘extensions’):
    - Collaboration with **ESLEA**
  - Engineering of the protocols and end-systems
  - Network feedback to applications for adaptation
- **Real-time traffic monitoring at multi-Gb/s:**
  - QML leads
  - Multi-protocol analysis:
    - Collaboration with **MASTS**
  - Traffic modelling

Current Status and Results

46PaQ **current** experiments

- Connectivity to UKLIGHT:
  - UCL: 2 × 1Gb/s to Chicago, remote loopback
- 1Gb/s loopback experiments on UKLIGHT from Q1/2005 (London-Chicago)
- High-speed tests within the lab:
  - **Bi-directional data** flows (not just TCP/FTP)
  - **Multiple data flows** (not just single flows)
  - 10Gb/s lab tests at UCL Q1/2006
- Examine end-to-end data path:
  - Interaction between application, hosts & network
  - Performance tuning for high-speed operation

Equipment setup

- End hosts (one as source, one as sink):
  - Gentoo Linux (AMD64 compiled), kernel 2.6.11
  - 2x Opteron 246, 2MB L2 Cache, 2.4GHz
  - Tigon chipset on motherboard
  - 2GB PC3200 DDR RAM on dual banks
  - Intel PCI-X 133MHz 64Bit dual 1Gb/s NIC
- Local switch:
  - Cisco Catalyst 3750
  - 2x VLAN, one for each of the 1Gb/s links to Chicago
- Tests run for two configurations:
  - Local-switch loopback 1Gb/s (RTT fraction of ms)
  - Chicago loopback 1Gb/s (RTT ~177ms)

UDP Baseline Tests

- Parameters:
  - UDP datagram sizes [bytes]:
    40, 500, 1000, 1470 (also 4096 and 16384)
  - Receive buffer sizes:
    64, 128, 256, 512, and 1024 KB
  - UDP flow duration: 2 mins
  - All experiments conducted using *iperf*

UDP Throughput

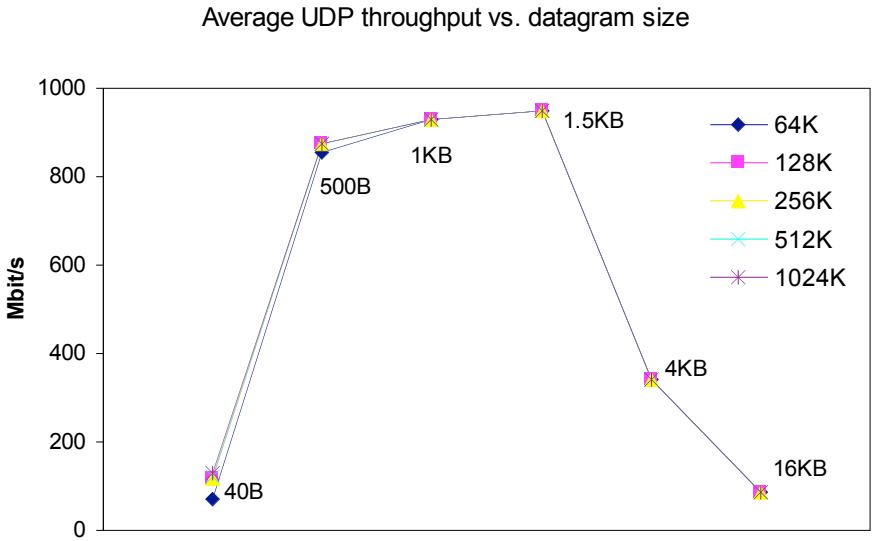

Local-switch loopback

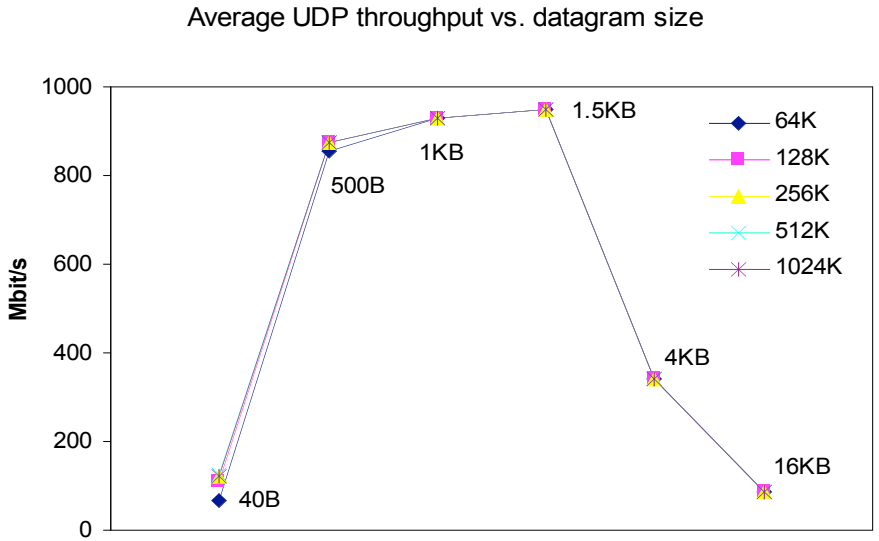

Chicago loopback

TCP Baseline Tests

- The following TCP settings were used:
  - net.ipv4.tcp_window_scaling = 1
  - net.ipv4.tcp_sack = 1
  - net.ipv4.tcp_rmem = 4096 87380 67108864
  - net.ipv4.tcp_wmem = 4096 87380 67108864
  - net.ipv4.tcp_mem = 196608 262144 67108864
  - net.core.rmem_default = 126976
  - net.core.wmem_default = 126976
  - net.core.rmem_max = 33554432
  - net.core.wmem_max = 33554432
- All experiments conducted using *iperf*

TCP Throughput, local-loop [1]

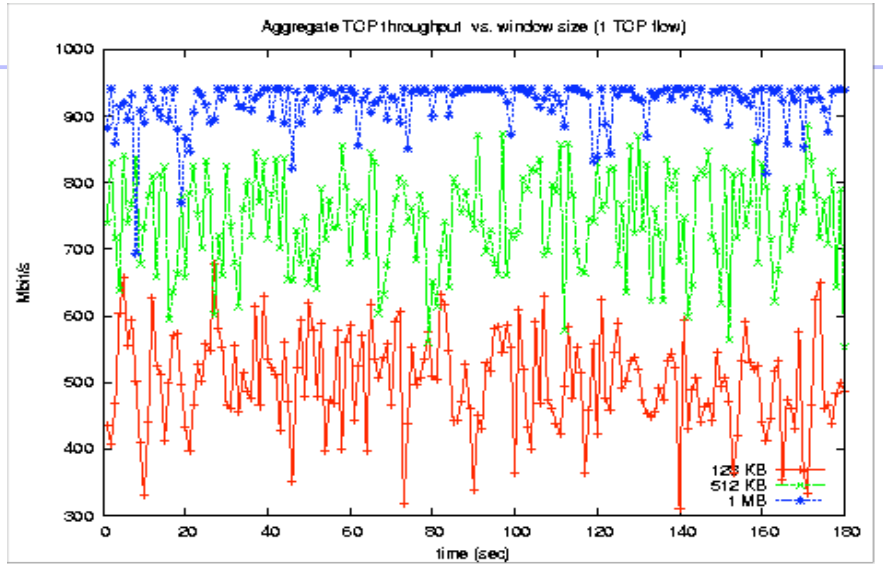

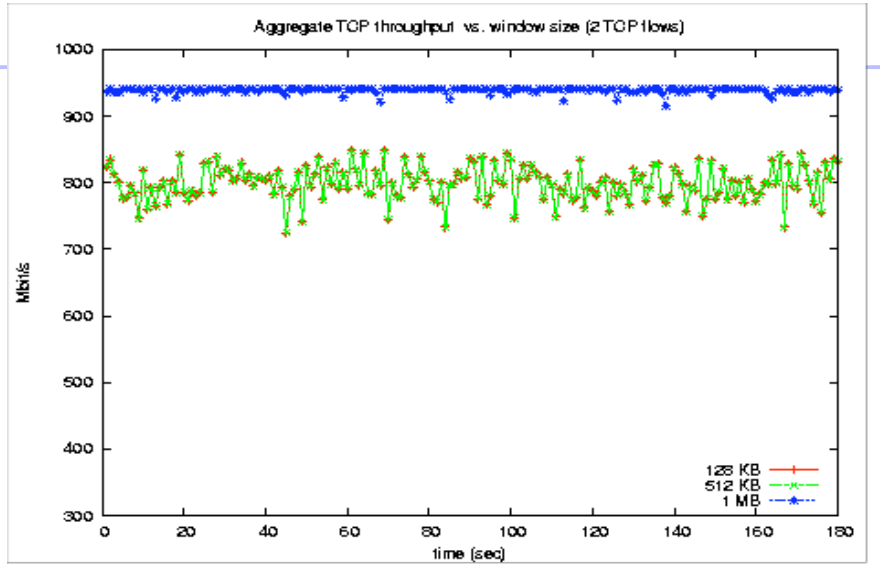

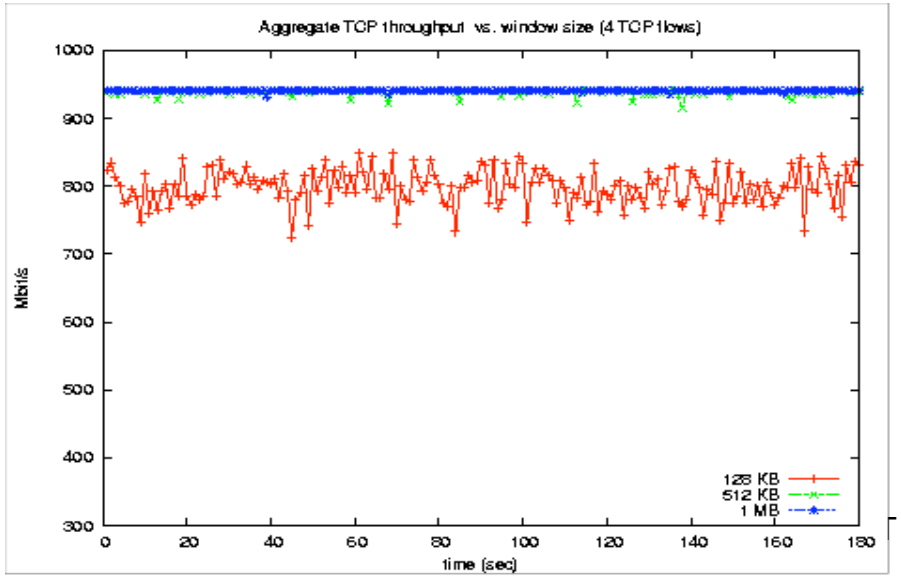

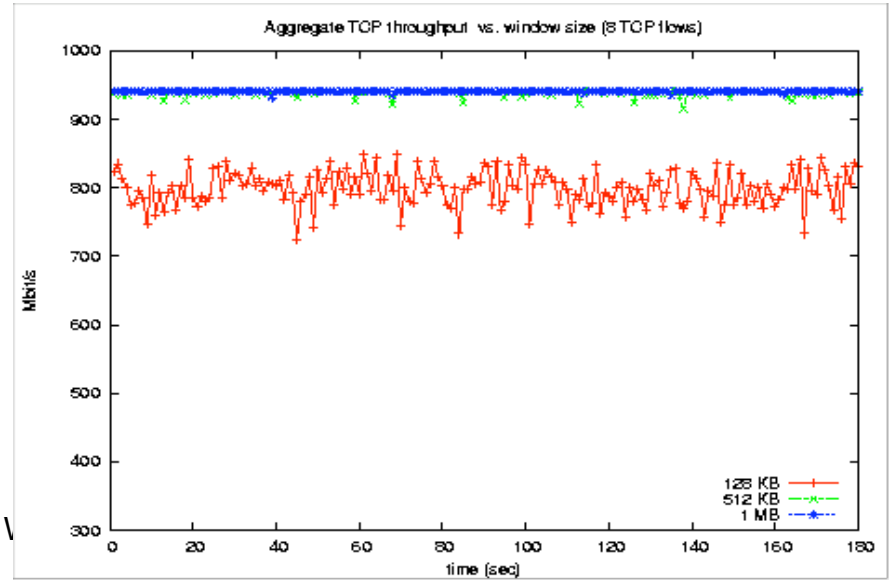

TCP Throughput, local-loop [2]

Multiple-flows, aggregate throughput

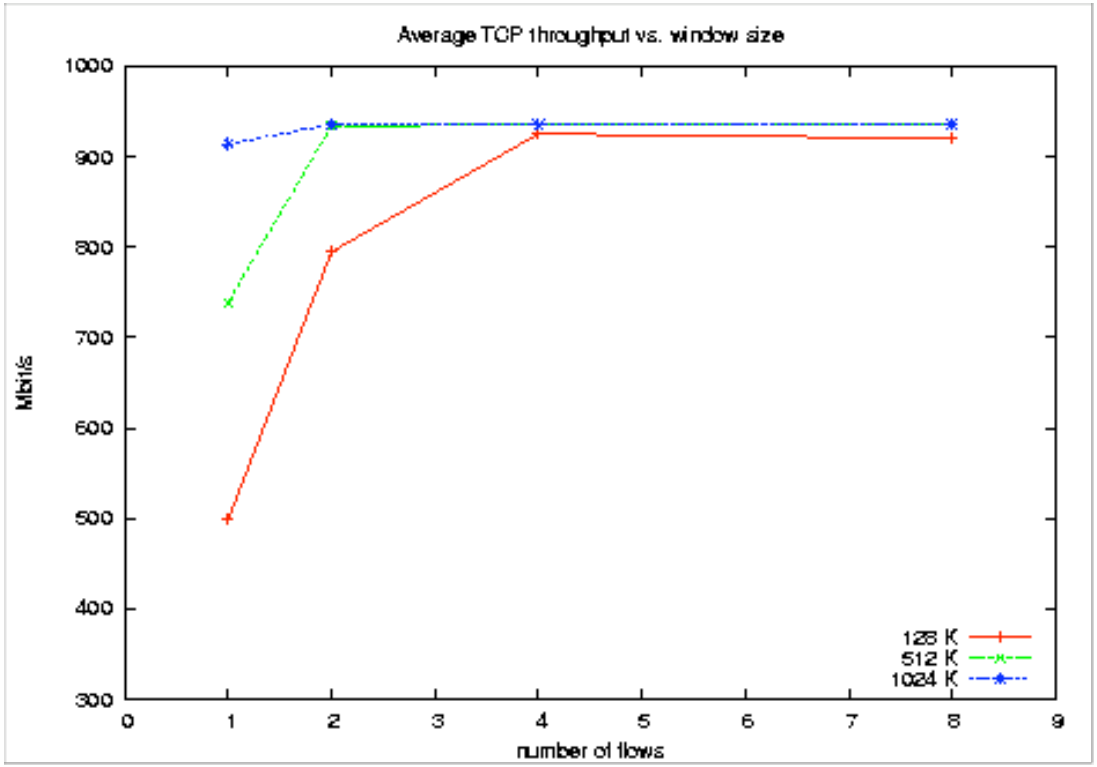

TCP Throughput, Chicago-loop [1]

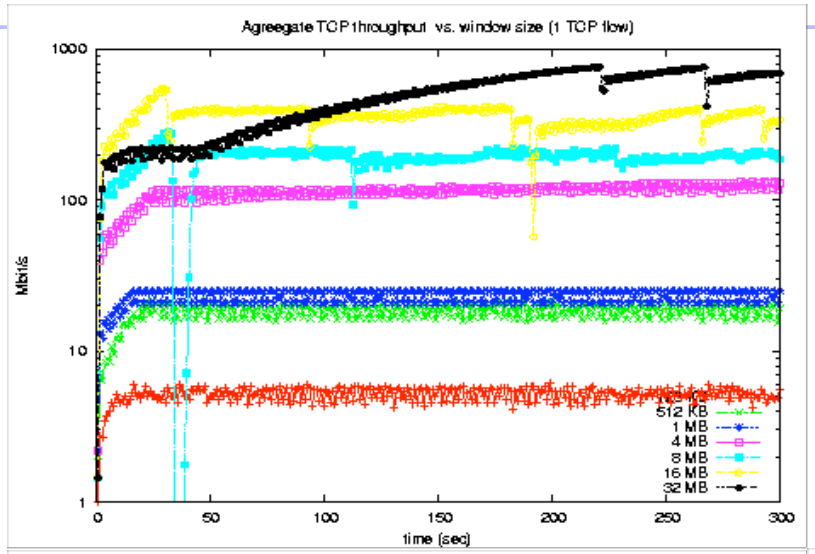

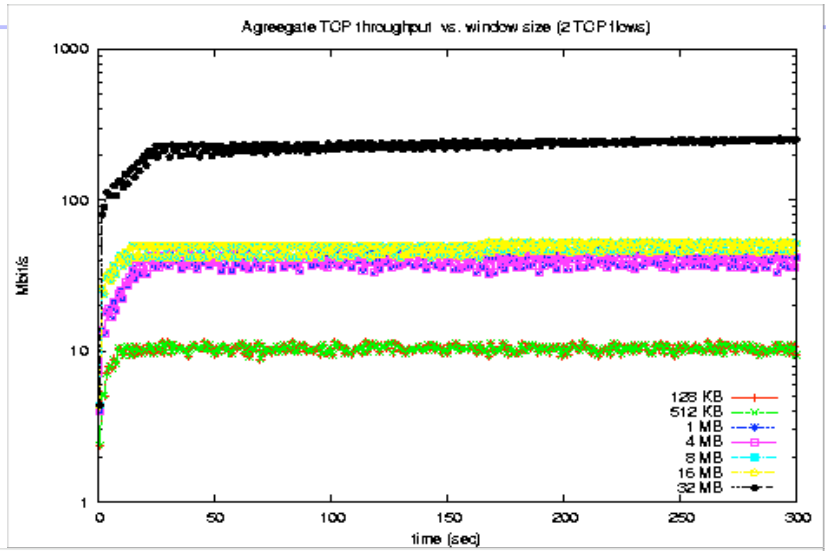

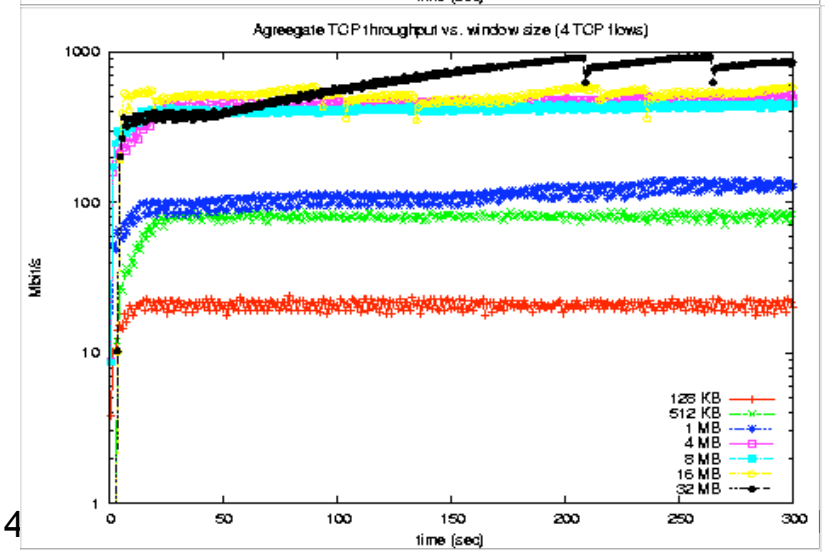

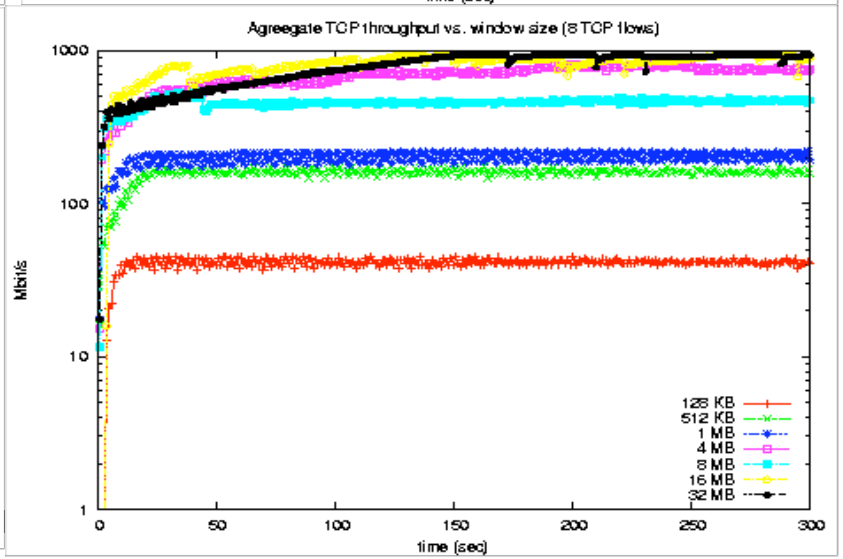

TCP Throughput, Chicago-loop [2]

Multiple-flows, aggregate throughput

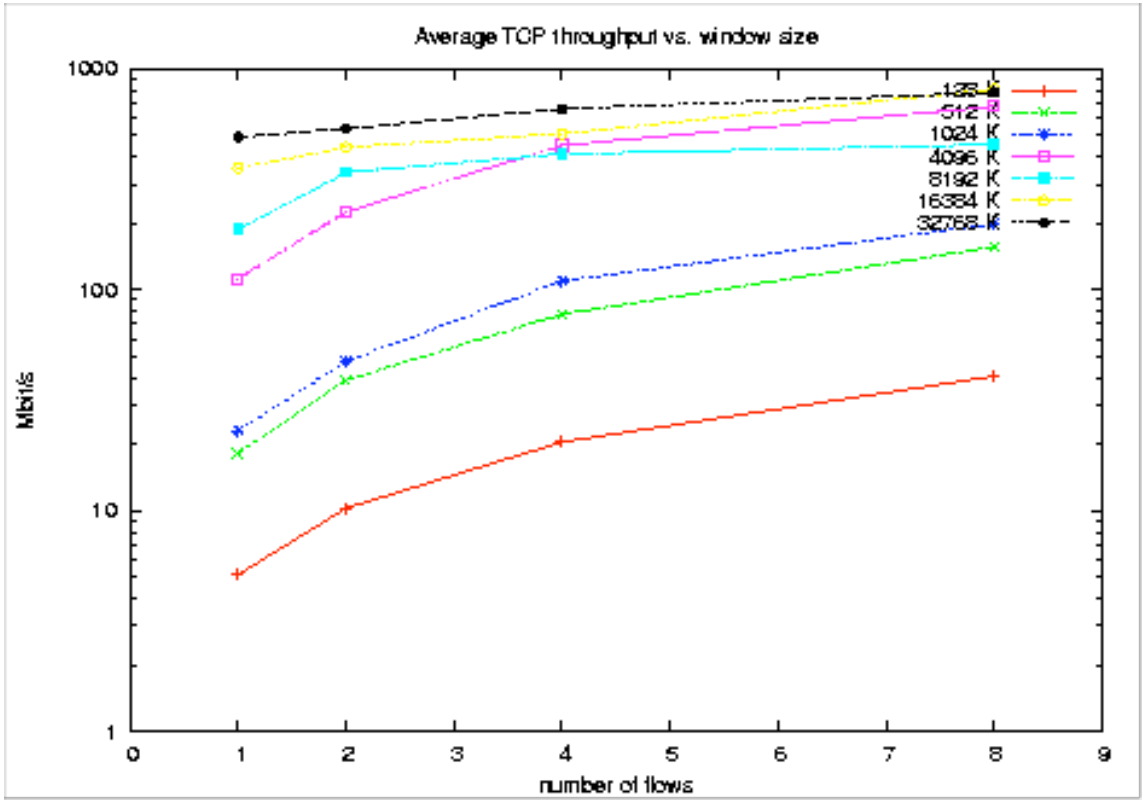

Observations

- Verification of known results
- Protocol performance very much dependant on several parameters (UDP buffer sizes, TCP window size and sender recv buffer size, CPU/HW speed)
  - For UDP, throughput depends on packet size
  - For TCP, window size is important for large bandwidth-delay products
- Tuning of protocol stack and/or application is therefore the most important performance factor

TCP ‘Fairness’ - **Very** Early Results

Matched bit-rate fairness (MBF)

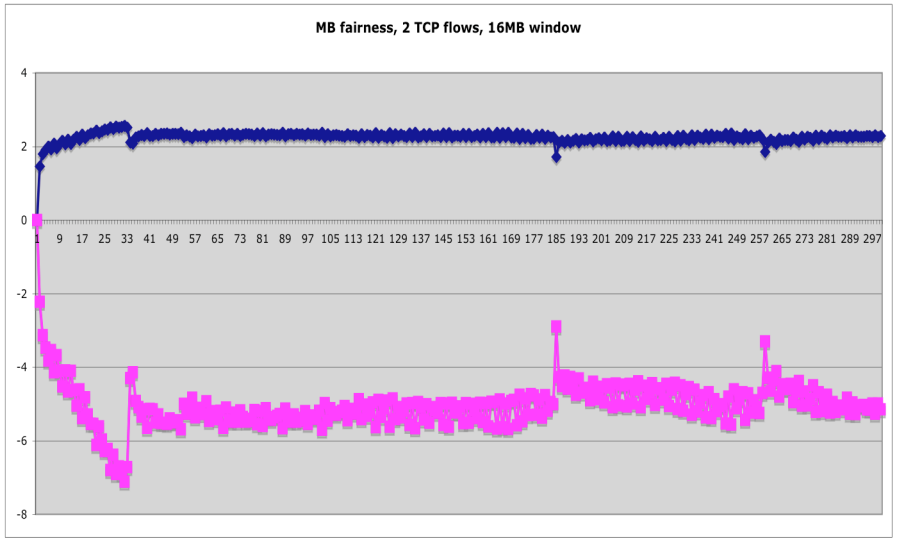

Chicago-loop, 16MB window

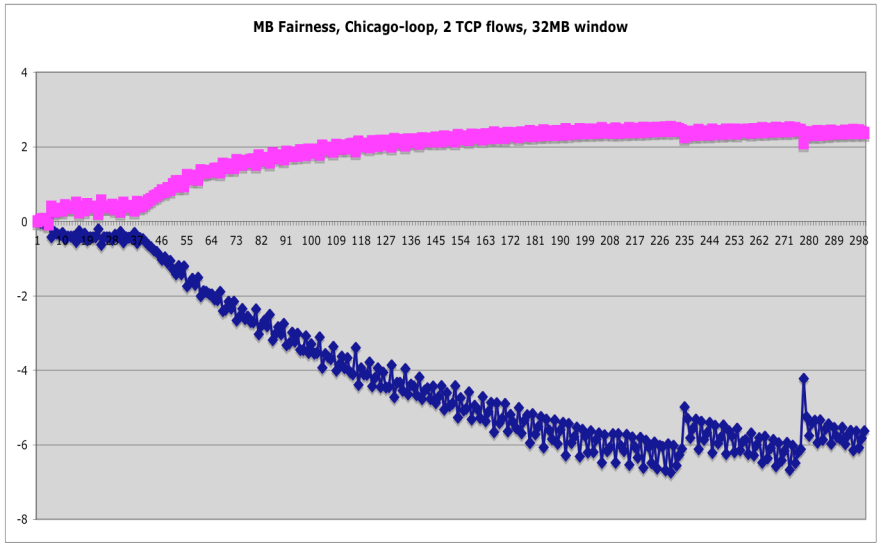

Chicago-loop, 32MB window

$$F_k(t) = 10\log_{10}\left(\frac{Nb_k(t)}{\sum_{n=1}^{N} b_n(t)}\right), \quad k = 1,\ldots,N$$

Future Work

46PaQ **future** experiments [1]

- Starting DCCP experiments
- Interaction with ECN and DIFFSERV
- Performance on IPv6
- Co-operation with other projects:
  - MASTS
  - ESLEA

46PaQ **future** experiments [2]

- Performance tests in the lab during 2005:
  - Loopback and site-to-site (sites to be chosen)
  - High-speed real-time traffic monitoring
  - Performance analysis and comparison for different protocols
- Move to 10Gb/s by Q1/2006:
  - 10Gb/s lab tests Q1/2006
  - **UCL: move to UKLIGHT 10Gb/s Q2/2006**
- **10Gb/s experiments on UKLIGHT during 2006**

Questions