




From Airplanes to Elephants



Chris Thomas

Sr. Network Engineer

UCLA Office of Information Technology

Chris.Thomas @ ucla.edu

March 12, 2007

Introduction



The PAC-10 Conference



- UCLA, USC, Stanford, Cal, Arizona, ASU
– CENIC (HPR)



- Oregon, OSU
– Oregon Gigapop, Abilene



- Washington, WSU
– PNW GP, PacificWave




Use of Video in Coaching Sports

- Essential to any modern sports program
 - All coaches in all sports use video
 - Football is largest user
- No In-person scouting (NCAA rule)
 - Each school tapes own games and then exchanges video with next opponent

Video Processing

- Video capture at stadium
 - Broadcast-quality video equipment
 - Video captured directly to disk @ 25 Mbit/sec (DV-25 standard – SD, not HD yet)
 - Typical game is 18 gigabyte file
- Loaded into video editing system
 - Manually mark start & end of each play
 - Intercut two camera angles
 - Add keywords (e.g., type of play)

Press Box View



Old Process

- Early each Sunday
 - Put all our games played so far into a box
 - Drive to airport (may not be close), send box to next opponent, wait until his box arrives
 - Return to office, process video, distribute to coaches' video stations
 - Coaches expect video by 5 PM Sunday
 - Repeat next weekend with next opponent

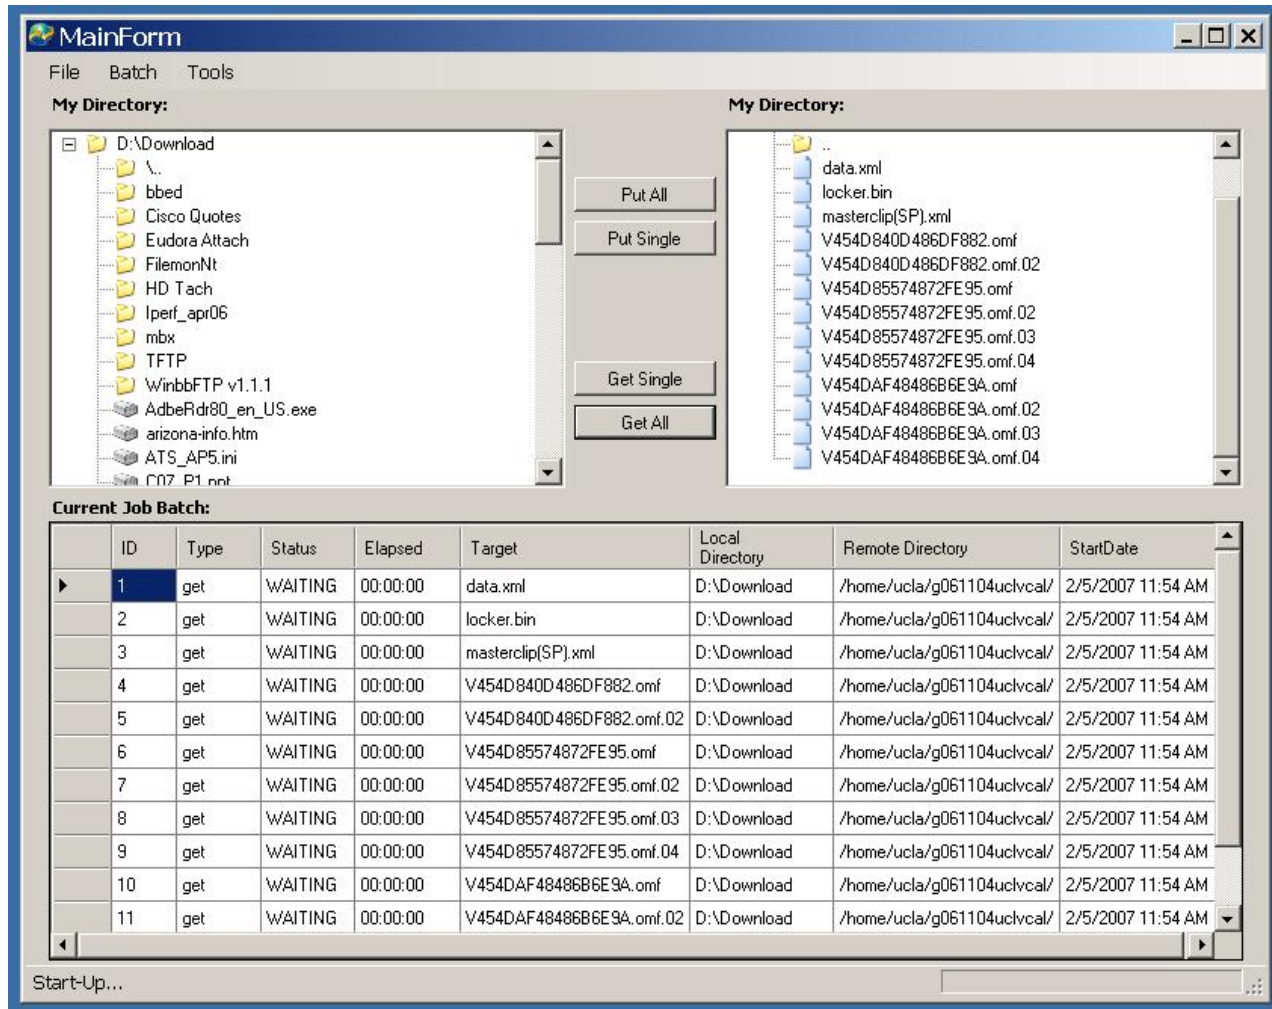
“The Challenge”

- Exchange multiple games between ten schools every weekend during season
 - Is sufficient bandwidth available?
 - 5 games/week = 95 GB/week, 950 GB total/wk
 - Timely and predictable?
- Using Windows-based systems
- Users more interested in sports than networks

Electronic Video Exchange

- Pilot for 2005 season
- Production for 2006 season
- Used bbFTP (freeware)
 - <http://doc.in2p3.fr/bbftp/>
 - UNIX server with UNIX & Windows clients
 - Selected on basis of performance bake-off on simulated network
- Central server model (server at UCLA)
 - Alternative is peer to peer

Jeff's bbFTP GUI



2006 Season Results

- All ten Pac-10 schools participated
- All conference games exchanged electronically – no courier shipments!
- Average server traffic was 1.4 terabytes per week
 - Largest source of CENIC traffic at UCLA
 - UCLA ↔ OSU made Abilene “Top 10 bulk flows” each week of season
- Users saw **NO** significant problems

How fast is typical transfer?

- 100 Mb/s NIC – 24-27 min per game
 - Typically 90 megabits/second
- Gigabit NIC – 6-9 min per game
 - Typically 250-400 megabits/second
- Highly repeatable
- Distance-independent (fairness)

2006 Results (cont)

- All 10 schools passionate about results
 - “FTP saved me an incredible amount of time”
 - “Best thing since non-linear video” – USC
 - “I got to spend Sunday mornings with my kids” – UCLA
- Saving \$6K/yr offsets a significant part of CENIC/I2 membership fee

Additional Benefits

- Elimination of courier costs makes participation by other sports feasible
 - Pac-10 Women's VB
- Solved network issues and built practical high-speed FTP experience at ten schools
 - Experience useable by other potential users

Performance Issues

- File exchange testing with ~25 schools
 - “Ubiquitous 100” mbit/sec isn’t here yet
 - Not even close; 30 mbit/sec typical
 - Getting worse, not better
- You’ll encounter other campus’ network performance issues
 - Slow departmental connections (10 mb/s)
 - Firewalls
 - Asymmetric performance (send vs. receive)

Performance Issues (cont)

- Highest traffic campus has seen
 - “Network police”
 - What is message to our user communities?
- Once you solve first/last mile (campus) problems, CENIC is excellent transport

Why *doesn't* FTP go fast?

- There are 4 main reasons

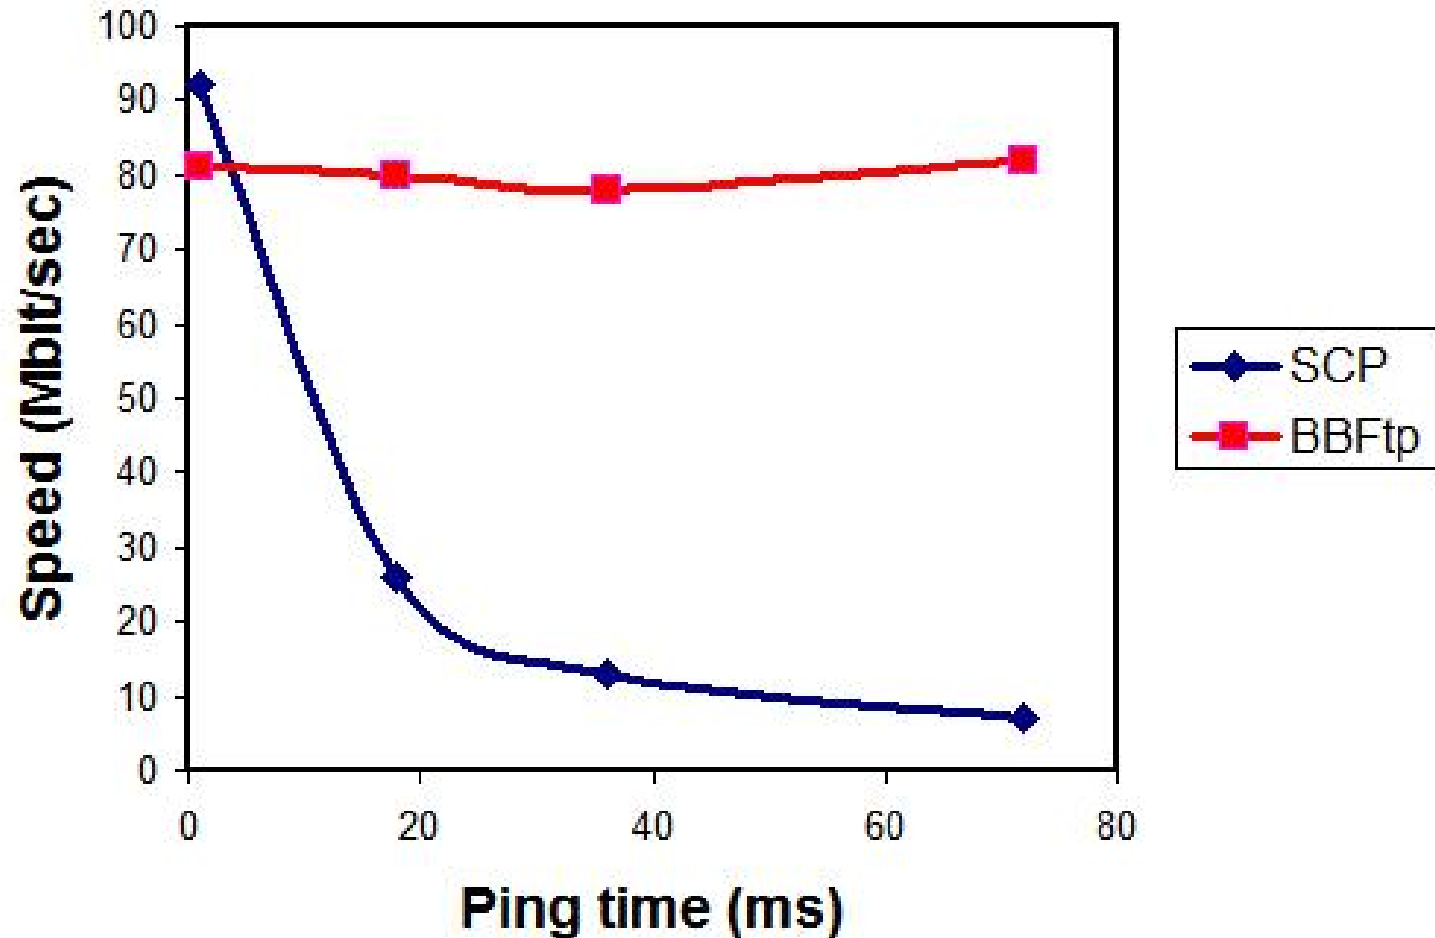
1 – Actual Network Speed

- Effective BW of “Internet2” is over-estimated for single user flows
- Slow departmental connections
- Slowest regional network links (OC3) are **100 times slower** than Abilene and fast regional networks like CENIC HPR
 - Slow (OC3) links are often heavily congested
 - 0.5% busy on 10G is 50% busy on 100M

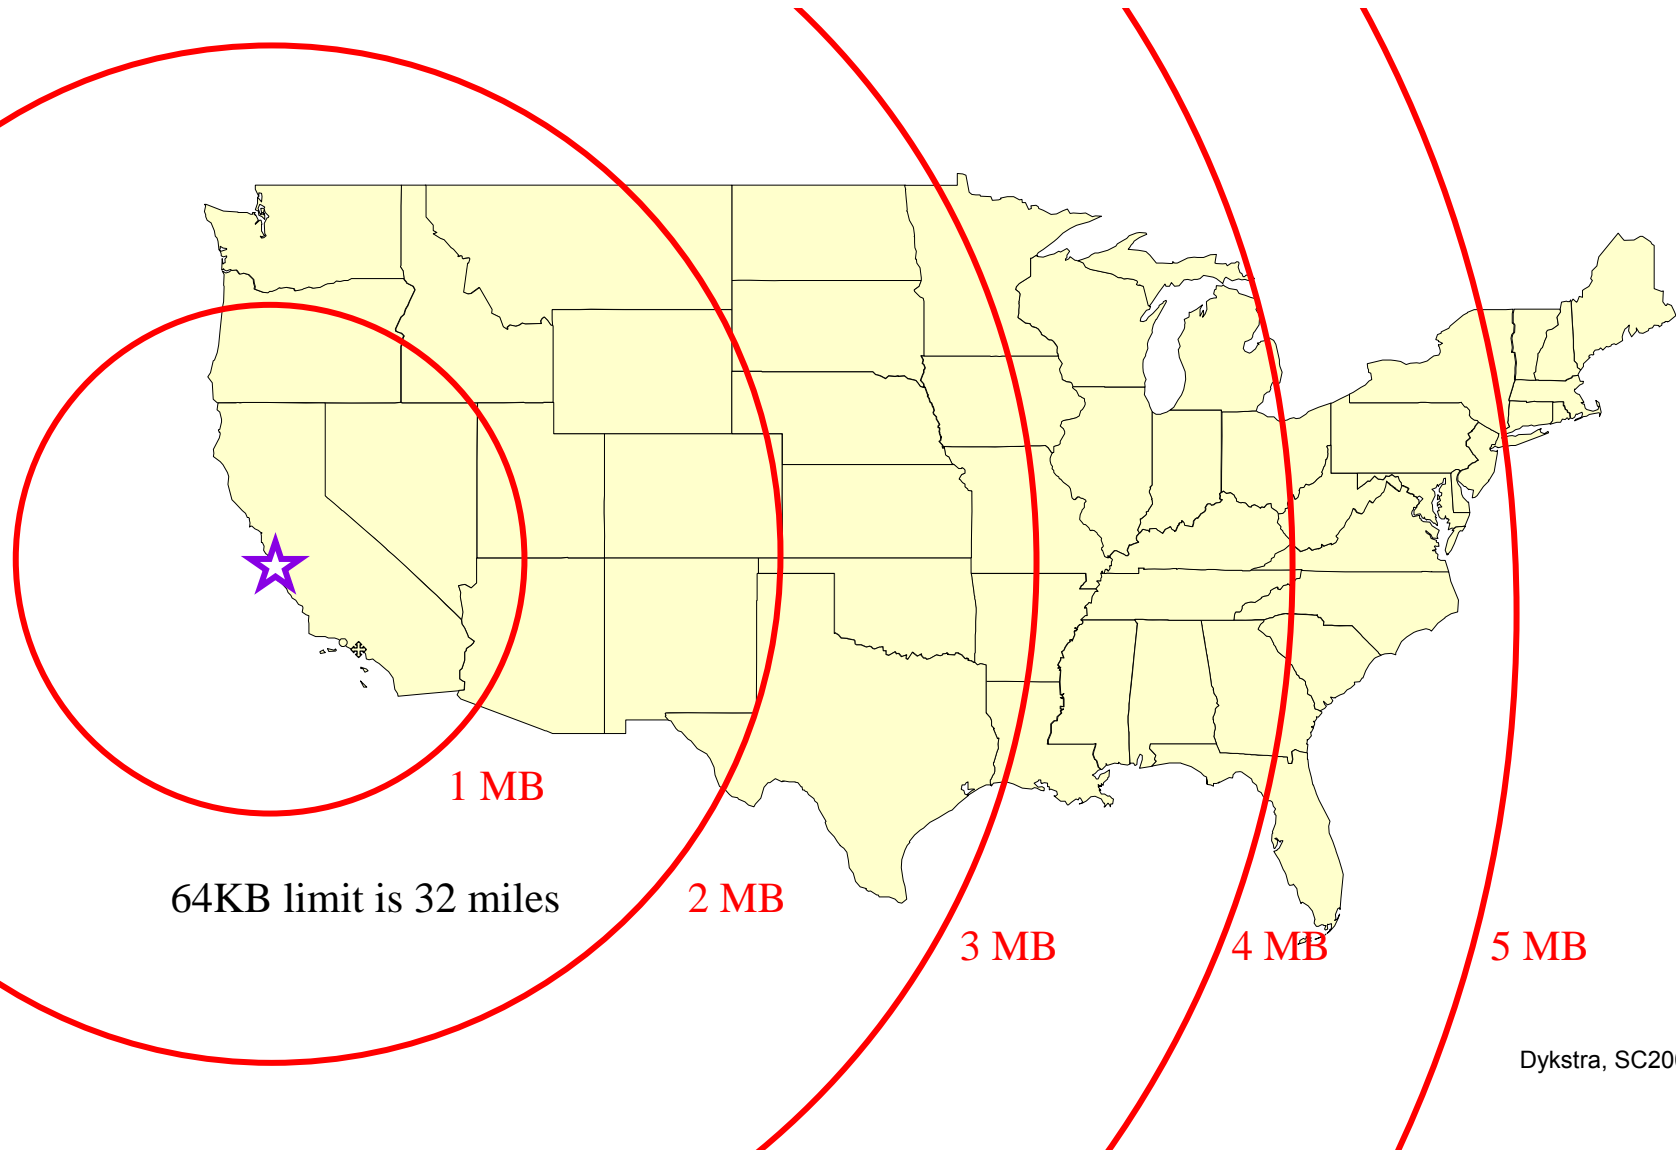
2 – TCP Design

- “Send and wait” characteristic of TCP
 - On a long path, can be 95% wait time
 - Distance-dependent transfer speed
 - Known solution

2 – Distance-dependent Speed



Receive Windows for 1 Gbps



3 – Congestion Control

- Original TCP lacked ability to deal effectively with congestion (other traffic, send/receive speed mismatch)
 - Caused network meltdowns
 - TCP Reno introduced congestion window
 - Turned out TCP Reno's fix was overkill
- Problem addressed in recent TCP work
- Selective ACK greatly reduces resent data

TCP Development

- High-speed TCP development is [ongoing](#) (mostly on UNIX)
 - TCP Reno unusable performance at 10G
 - Current UNIX has working high speed TCP
- Windows XP is TCP Reno
- Windows Vista has new TCP stack
 - See: [SLAC-TN-06-005](#) technical report

4 – PC Hardware Matters

- At 100 Mb/sec, recent PCs **can** keep up with network
- At gigabit speeds, average PCs **can't** go fast enough and become the bottleneck
- Use iperf to measure network performance
 - <http://dast.nlanr.net/Projects/Iperf/>
 - Doesn't include effects of disk speed

PC Disk Speed

- Disk speed is largest PC bottleneck
 - Single PC disks do about 30 – 40% of Gb network speed
 - <http://www.simplisoftware.com/Public/index.php?request=HdTach> (freeware)
 - Disk arrays are faster. Watch RAID-5 write speeds – can be 1/10 of read speed.
 - A pair of WD Raptor II in RAID-0 stripe will do Gb speed

PC Network Interface

- Network Card (NIC)
 - PCI-E bus
 - Large onboard buffers (e.g., 96KB)
 - TCP offload
 - Jumbo frames (9000 bytes) (2x speed)
 - Highly recommended: SysKonnnect SK-9S21 / SK-9E21 card (about \$100)
 - Even if you have a Gb port on the m'board, adding an SK card may help (2x)

PC Anti-Virus SW

- Anti-virus software typically has a **major impact** (~50%) on speed
 - Writes every file twice
 - Temporarily disable or configure to bypass exchange directory.

Parting Thoughts

- Know your own campus network
 - What file transfer rates are realistic from your campus to another?
- What happens when users call support with performance question?
 - Make sure your (potential) users have access to campus expertise
- “Be Galileo, not Aristotle”

Thank you!

Chris.Thomas @ UCLA.EDU