



TCP/IP Modeling For Congestion Management

Tanmay Gupta
Manoj Wadekar
Gary McAlpine

Intel Corp.
July 2004



Agenda

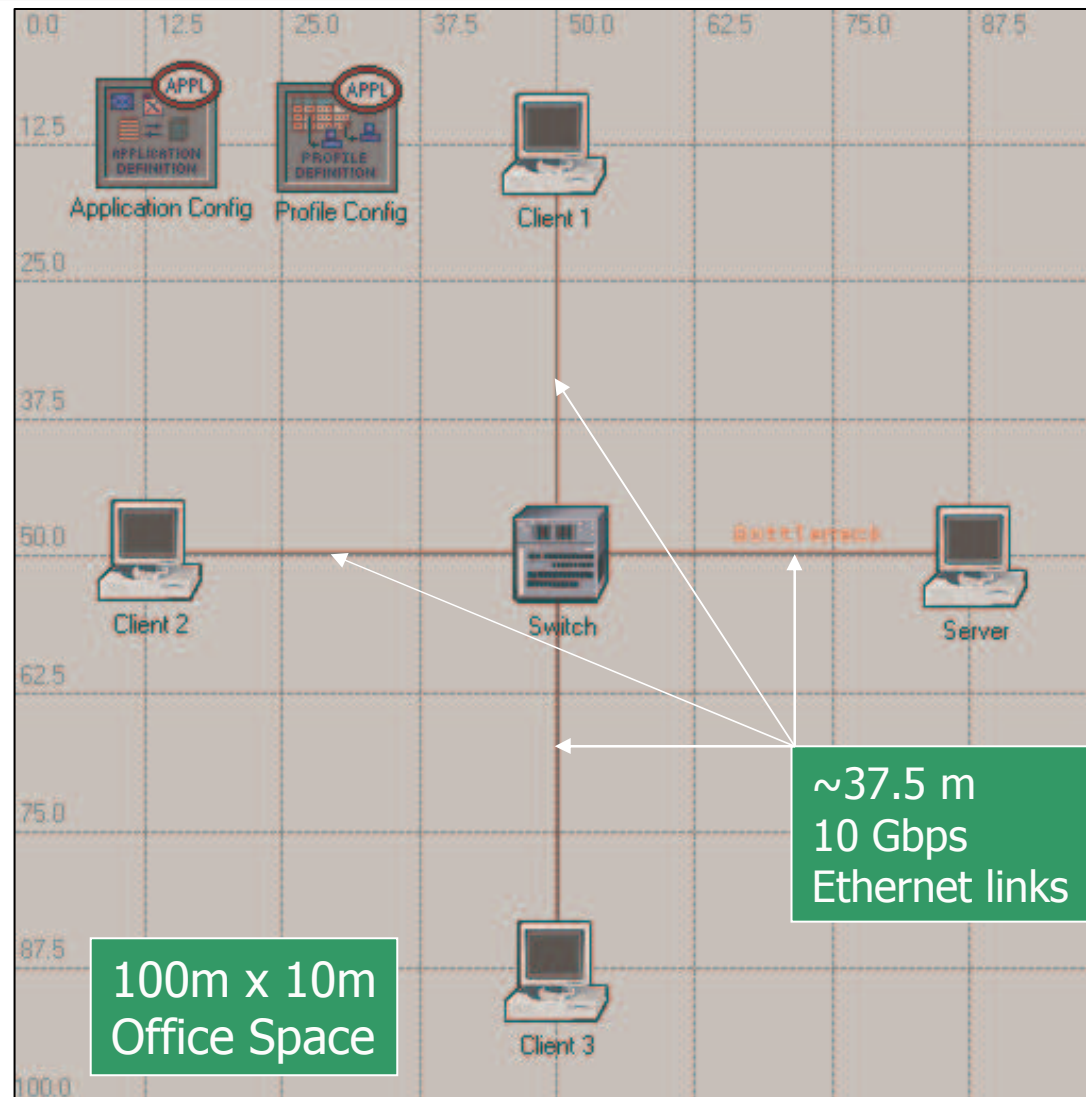
- Goals of Modeling
- Modeling Environment
- Preliminary Modeling Results
- Next Steps
- Summary



Goals of Modeling


- Study effect of Congestion on TCP/IP performance
 - Throughput
 - Latency (end-to-end, incl. retransmission)
 - Packet drops
- High speed (10Gbps), short link lengths (<100m)
- Study the impact of 802.3 Flow/Rate Control on TCP/IP performance

Modeling Topology



Modeling Workloads

- 3 Clients - 1 Server
 - TCP traffic
 - Upload only (client to server)
 - All traffic with same priority
 - All sources trying to generate application data at ~3.6 Gbps (excluding overheads)
 - Application data-size: Exponential (8000 bytes)





TCP Configuration

- Baseline WinXP TCP configuration with
 - Fast Retransmit
 - Fast Recovery (Reno)
 - Selective ACK
- Tuned parameters
 - 64K Receive Window
 - Retransmission Timers: Faster recovery from packet drops



Simulation Scenarios

- NoFC
 - No Flow Control
 - 48K Shared Memory
- 8023x
 - 802.3x (XON/XOFF)
 - 48K Shared Memory



Measurement Definitions



- Throughput
 - Measured at “Server” as total number of bytes/sec received
- Application Latency
 - Measured at the entry (from client App.) and exit (to server App.) interfaces for TCP
 - Includes time lost by TCP in retransmissions
- Packet Drop
 - Number of packets dropped at the switch due Shared Memory overflow
- Bandwidth Sharing
 - 3 applications compete for BW to send traffic to the same server. Comparison of how they share the bandwidth dynamically



Preliminary Simulation Results


TCP Application Throughput

At TCP (without overheads)




Packet Drop at the Switch

No Flow Control



Flow Control can eliminate packet drops


TCP Application Latency



802.3x can provide significantly lower jitter


BW Sharing between TCP Applications

No Flow Control



With no Flow Control, all the sources continuously compete for BW
Poor sharing with non-deterministic and uneven distribution between sources

BW Sharing between TCP Applications With Flow Control



Flow Control improves BW sharing



Next Steps

- Rate Control mechanisms
 - CM protocols with more granular control
- TCP enhancements
 - RED, ECN
 - Explore ongoing work in IETF to improve TCP performance in High Speed networks
 - High Speed TCP (Limited Slow Start with large Congestion Windows)
 - XCP



Summary

- For high speed, short link length networks:
 - Packet drops cause TCP performance degradation
 - Results in non-deterministic behavior
 - Flow Control provides significant improvement to TCP performance in congested conditions
 - Results in deterministic behavior
- Next step is to evaluate and simulate simple rate control enhancements

Will present simulation results at next SG meeting