

BCN Simulation Environment

CN-SIM Ad-Hoc Team

CN Simulation Ad-Hoc Team

Subbarao

Hugh

Davide

Greg

Uri

Uri

Tanmay

Raj

Benny

Bruce

Gopi

John

Pat

Manoj

Jeff

Arumilli

Barrass

Bergamasco

Chesson

Cummings

Elzur

Gupta

Jain

Koren

Kwan

Sirineni

Terry

Thaler

Wadekar

Wise

What needs to be agreed

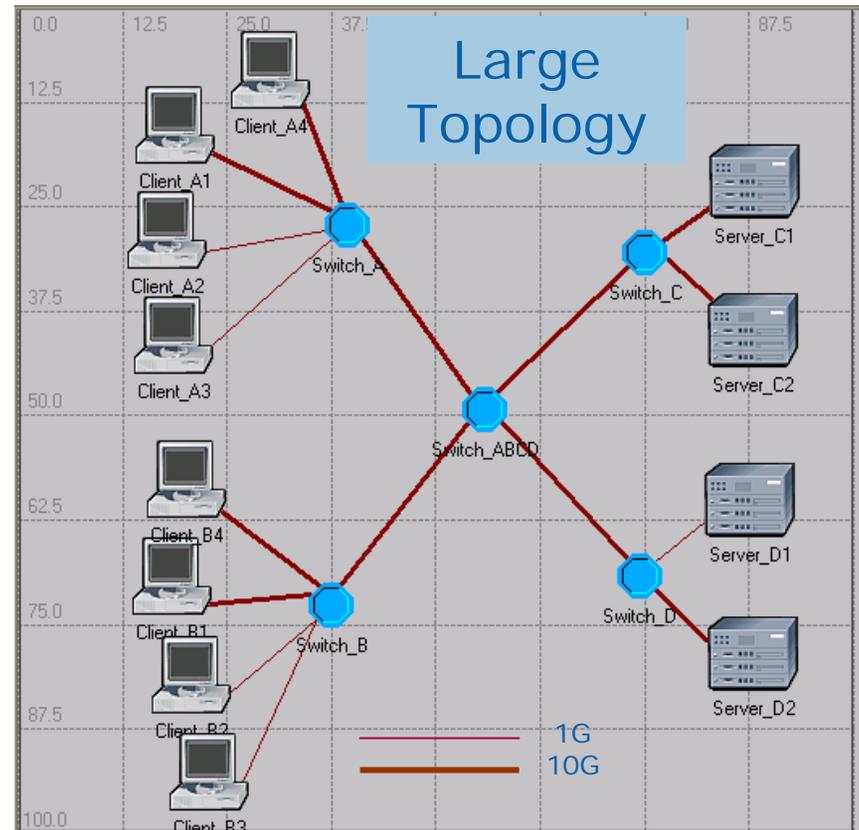
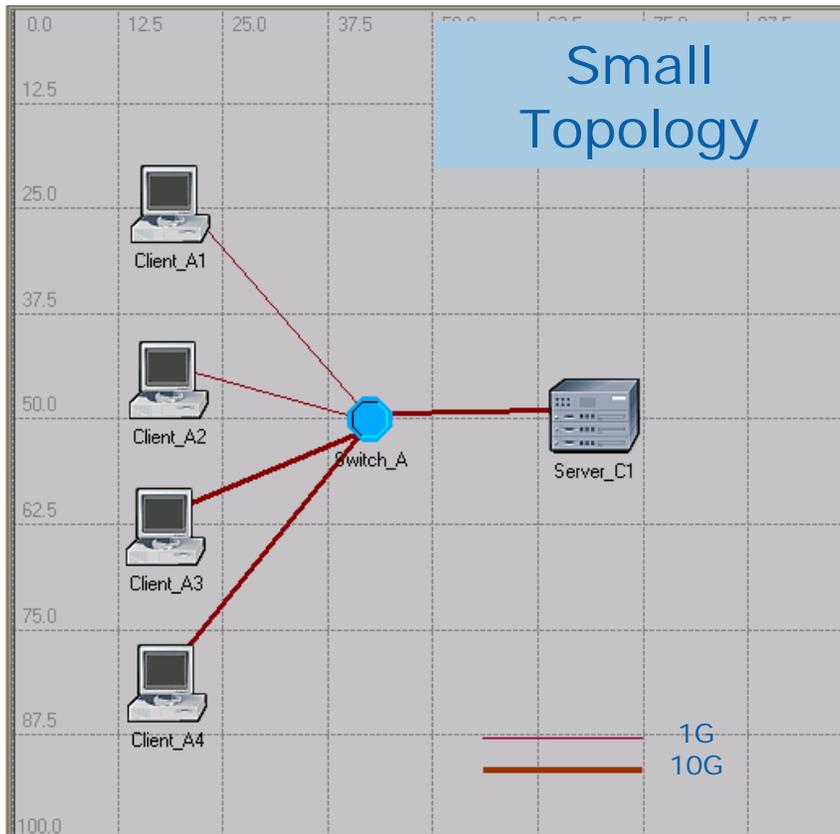
- Network Topologies/Scenarios
 - End Stations, Switches, Hops, Link Lengths (delay) etc.
- Simulation Workloads
 - Transport Layers
 - Application abstraction: Packet Size, distribution etc., Traffic Mix
 - Granularity of flows, number of flows etc.
- Measurement Metrics
 - Throughput (where – application, congested link etc.)
 - Latency (where – application, L2 etc.), Latency Jitter?
 - Buffer Utilization?
 - Fairness factor?
- CN Protocol
 - Davide's September Presentation AND
 - FAQ document to clarify details

What does not need to be agreed

- Simulation Tools/Methodology
- Switch or end station implementation details (? If required, how much be disclosed?)
- ??

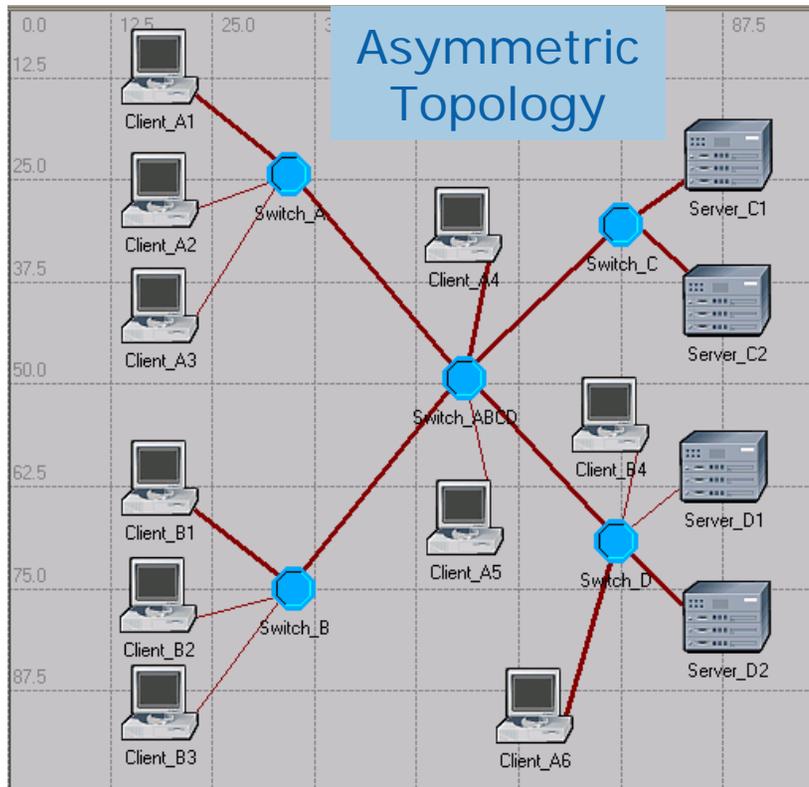
Simulation Framework Proposal

Topologies



- Mix of 10 GbE and 1 GbE links to create extreme congestion
- Less than 100m link lengths

Topologies ..contd



- Adding clients to intermediate, congested switches
- Provides different “distance from congestion” for sources

Workloads – Application Characterization

1. File Transfer Workload

- Large bursts of packets
- Data in large bursts, typically 64 Kbyte and rising
- a. File Transfer Workload (Read)
- b. File Transfer Workload (Write)
- c. File Transfer Workload (Mix)
 - 50% Read, 50% Write

2. Database Access

- Mix of large and small traffic
- Double peak : 256B and large packets
- a. Database Access (Read)
- b. Database Access (Write)
- c. Database Access (Mix)
 - 50% Read, 50% Write

Workloads Transport Details

- Transport Layers
 - TCP and UDP
- Traffic Mix
 - 80% TCP and 20% UDP
 - All traffic with same 802.1p priority
- Granularity of flows, number of flows etc.
 - Each client initiates 10 TCP connections and 1 UDP connection to each server
 - All flows are persistent long-lived flows

Simulation Scenarios

	File Transfer - Read	File Transfer - Write	File Transfer - Mixed	Database - Read	Database - Write	Database - Mixed
Small Topology	S-FT-R	S-FT-W	S-FT-M	S-D-R	S-D-W	S-D-M
Large Topology	L-FT-R	L-FT-W	L-FT-M	L-D-R	L-D-W	L-D-M
Asymmetric Topology	A-FT-R	A-FT-W	A-FT-M	A-D-R	A-D-W	A-D-M

Metrics

- Throughput
 - Granularity
 - Application level throughput (workload dependent)
 - Aggregate link throughput
 - Per flow throughput
 - Measured at
 - Most Congested Bottleneck link
 - Uncongested link
 - Fairness across flows contributing to congestion
 - Fairness definition required (i.e. Max-Min Fairness)
 - Jain's fairness index may be used to characterize CN protocol capabilities
- Latency
 - Mean, Min, Max, Variance

Metrics (contd..)

- Buffer Utilization
 - Measured at congestion detection point
 - Measure in units of bytes (not packets)
 - Mean, Max, Variance
- Packet Drop Probability (included from 6/2006 slide set)
 - Number of packets dropped in switching interconnect due to congestion
 - However, depends on buffer resources available in switching interconnect
- Time to Fairness
 - Time required to achieve a specific fairness goal following the introduction of a persistent congestion event.
 - Requires definition of target goal for fairness