Delay and Delay Variation Simulation Results for Multi-hop Conventional Ethernet Cases with Bursting/Bunching

Geoffrey M. Garner
gmgarner@comcast.net

SAMSUNG Electronics

IEEE 802.3 ResE SG
2005.08.19

Outline

- Introduction
- Simulation models and assumptions
- Simulation cases
- Simulation results
- Conclusions
- References

Introduction

- Reference [1] presented initial simulation results for transport of time-sensitive traffic over conventional Ethernet
  - Considered simple, one-hop, two-switch network
    - one switch-to-switch link with end-devices attached to each switch
- Results showed that unfiltered peak-to-peak delay variation for competing Constant Bit Rate (CBR) traffic streams whose rates differ slightly from nominal can be appreciable compared to the requirements for digital audio and video [2]
  - For 3 CBR streams, 50% link utilization, and 256 byte packets, unfiltered delay variation was almost 50 μs in one case
    - For sufficiently small frequency offsets (e.g., 1 ppm or less), phase-locked loop (PLL) filtering at the egress did not reduce the delay variation appreciably
  - This exceeds the requirements for uncompressed digital video and digital audio, and is close to the limit for compressed digital video (50 μs) [2]
  - For 6 CBR streams, 50 % link utilization, and 256 byte packets, unfiltered delay variation exceeded 100 μs in one case, and was reduced to just over 80 μs by filtering
    - This exceeds the limits for digital audio and video
  - Also considered adding a best-effort stream with maximum size packets, though this did not appreciably change the time-sensitive stream results

Introduction (Cont.)

- Discussion during the presentation of [1] indicated it would be of interest to show a worse case, with multiple hops
- Further discussion in a subsequent ResE SG conference call indicated it would be of interest for the multiple hop case to resemble the bursting/bunching scenarios described in [3] (see Annex F of [3] for details)
  - Combine N traffic sources from N locally-attached end devices at a switch, and transport over a link to a downstream switch
  - Replicate this configuration N times, so that the downstream switch has N incoming links
  - Drop the traffic from $N - 1$ of the sources from each incoming link at the downstream switch to locally-attached end devices (the number of locally attached end devices is therefore $N(N-1)$)
  - Transport the remaining N streams (one from each incoming link) over an outgoing link to a downstream switch
  - Repeat the above scenario at the downstream switch, i.e., replicate the above configuration N times
  - For k stages, the total number of sources at the ingress grows like N^k

Introduction (Cont.)

- The analysis in [3] is mainly qualitative, i.e., is carried out by graphically representing packets at various times
  - The analysis assumes worst-case arrival patterns (i.e., packets from competing time-sensitive streams always arrive simultaneously)
- The analysis in [3] considers both queueing at the input of each switch and queueing at the output of each switch (these are separate cases)
- It was felt it would be desirable to simulate this scenario as a case that is possibly worse than those considered in [1]
- In addition, it is of interest to consider total delay for multi-hop cases
  - While total delay was not explicitly discussed in [1], end-to-end delays for a path consisting of the ingress link, single switch-to-switch link, and egress link were on the order of at most 300 μs, and were this large only for the case that included a single best-effort stream with maximum size packets
    - The longest path through the network was between 100 and 200 m, and the propagation delay was 1.755×10^8 m/s (the default minimum propagation speed in Opnet, which assumes dispersion representative of the medium and configuration
    - With this assumption, propagation delay is of the order a few μs, and is therefore negligible (and would be negligible with no dispersion)

Simulation Models and Assumptions

- As in [1], OPNET simulation tool was used to simulate packet delays
  - OPNET contains models for full-duplex Ethernet MAC and for Ethernet bridges
  - Models were modified (as indicated in [1]) to include priority classes
    - Priority queueing is non-preemptive
- Considered basic topology as described in [3] and summarized in Introduction
  - At each stage, combine N previous stages
    - Each previous stage supplies N traffic streams to this stage
    - Drop the traffic from $N - 1$ of the traffic streams from each previous stage
      - Therefore, need $N(N - 1)$ end devices at this stage
    - Carry one traffic stream from each incoming link (from each previous stage) over an egress link to the next stage
      - Therefore, N traffic streams are carried to the next stage
- Assume all packets are maximum size
  - 1500 bytes (Opnet adds Ethernet overhead)
- Assume 100 Mbit/s Ethernet links
- Assume the maximum path length through the network is 100 – 200 m, and the propagation speed is as in [1], i.e., 1.755×10^8 m/s

Simulation Models and Assumptions (Cont.)

- Assume all time sensitive traffic streams have the same nominal rate, with a small frequency offset
  - Frequency offset is different for each competing stream, and is chosen on input
    - This captures the fact that Time-sensitive video and audio clients have specified nominal rates, but are allowed to differ from those nominal rates by specified frequency tolerances
  - Nominal rate is chosen based on the number of streams per switch at the network ingress and desired link utilization
    - Input rate (and time between packets) is constant
- OPNET model assumptions (same as in [1])
  - Two priority classes
    - Time-sensitive traffic gets high priority
    - Best-effort traffic gets low priority
    - Priority queueing is non-preemptive
    - Queueing is first-come, first-served (FCFS) within each priority class
  - OPNET model for full-duplex Ethernet MAC is used (with priorities added)
  - OPNET contains spanning tree and rapid spanning tree algorithms
    - Use rapid spanning tree algorithm here

Simulation Cases

- The above scenario can give rise to a very large number of traffic source
  - E.g., for 7 hops ($k = 7$) and 4 incoming links per switch ($N = 4$), get 4^7 = 16384 traffic sources
- Therefore, it was felt it would be desirable to look at more linear (i.e., rather than tree-like) configurations, with just a few traffic streams getting added and dropped at each node
- Case 1 – base case described above, with $N = 3$ and $k = 4$ (81 traffic streams)
  - All streams CBR (time-sensitive), with maximum size packets (1500 bytes + Ethernet overhead)
  - 3 of the 81 streams go through the entire network (4 hops)
  - Packet service time for 100 Mbit/s Ethernet is 0.12304 ms
  - Nominal packet rate = 1333.33 packets/s (nominal inter-packet time = 0.00075 s); 50% switch to switch link utilization
  - Actual rates differ from nominal by different amounts that are within ± 100 ppm (details in backup slides)

Simulation Case 1 - Stages 1 and 2

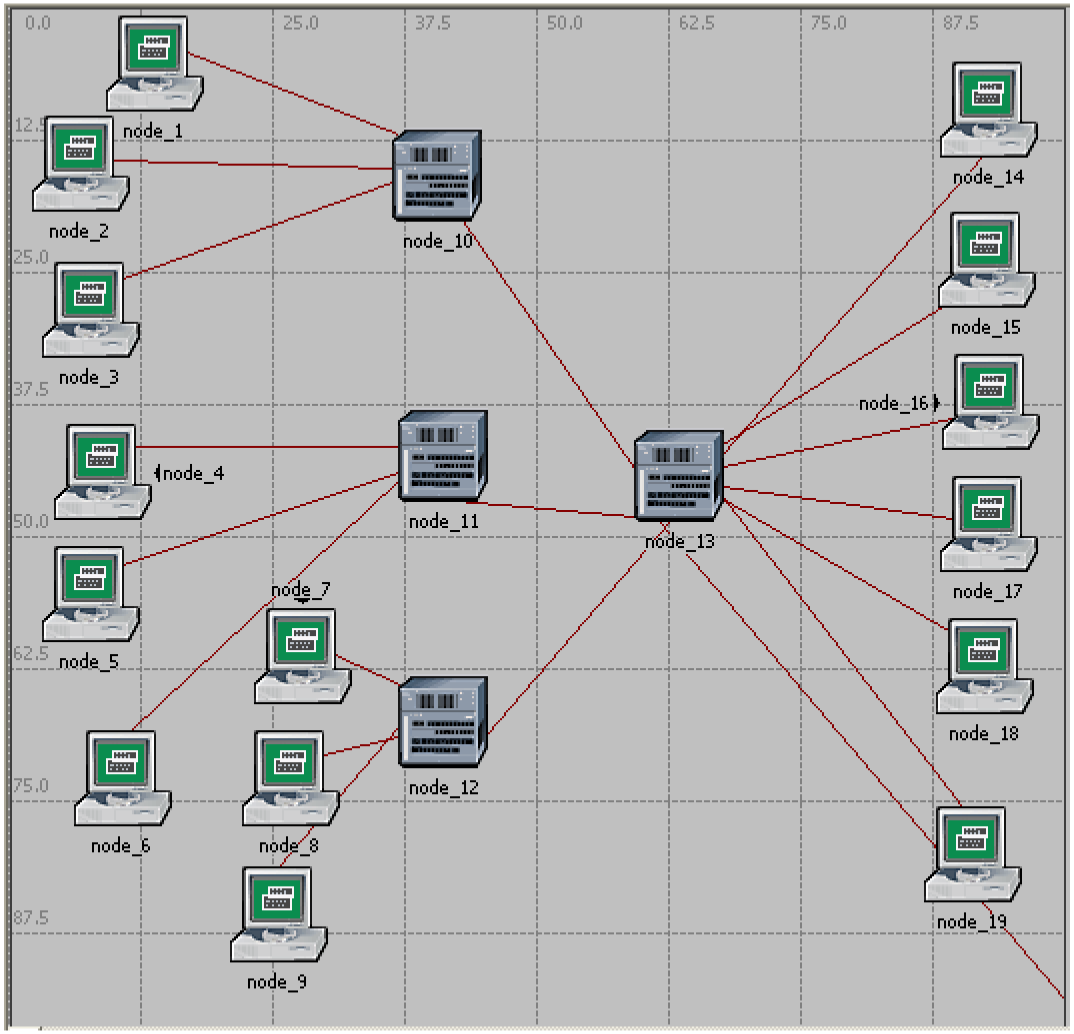

Switch to Switch link utilization = 50%

Simulation Case 1 - Stage 3

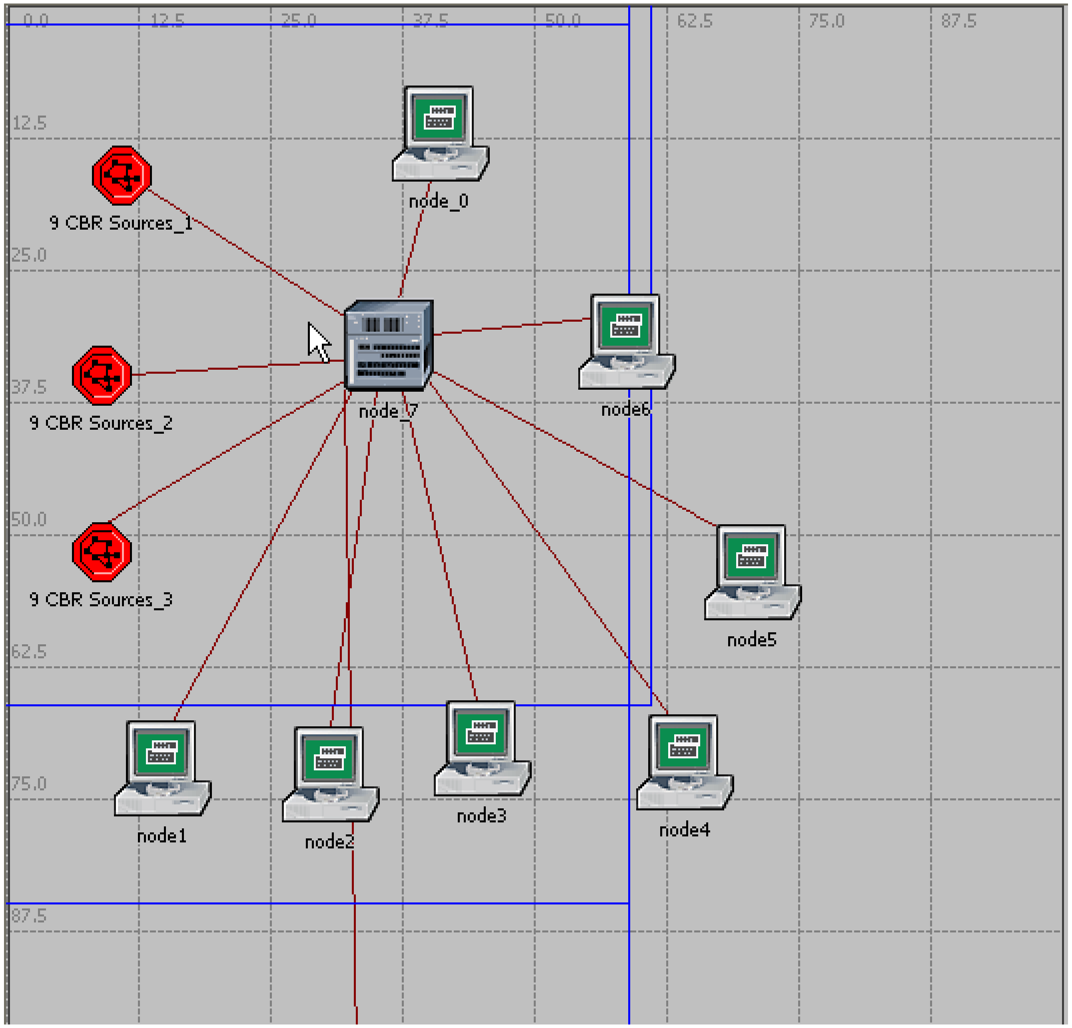

Switch to Switch link utilization = 50%

Each subnet (the red octagonal icons labeled 9 CBR Sources_1, _2, and _3 represent a stage 1 and 2 network as shown on the previous slide

Simulation Case 1 - Stage 4

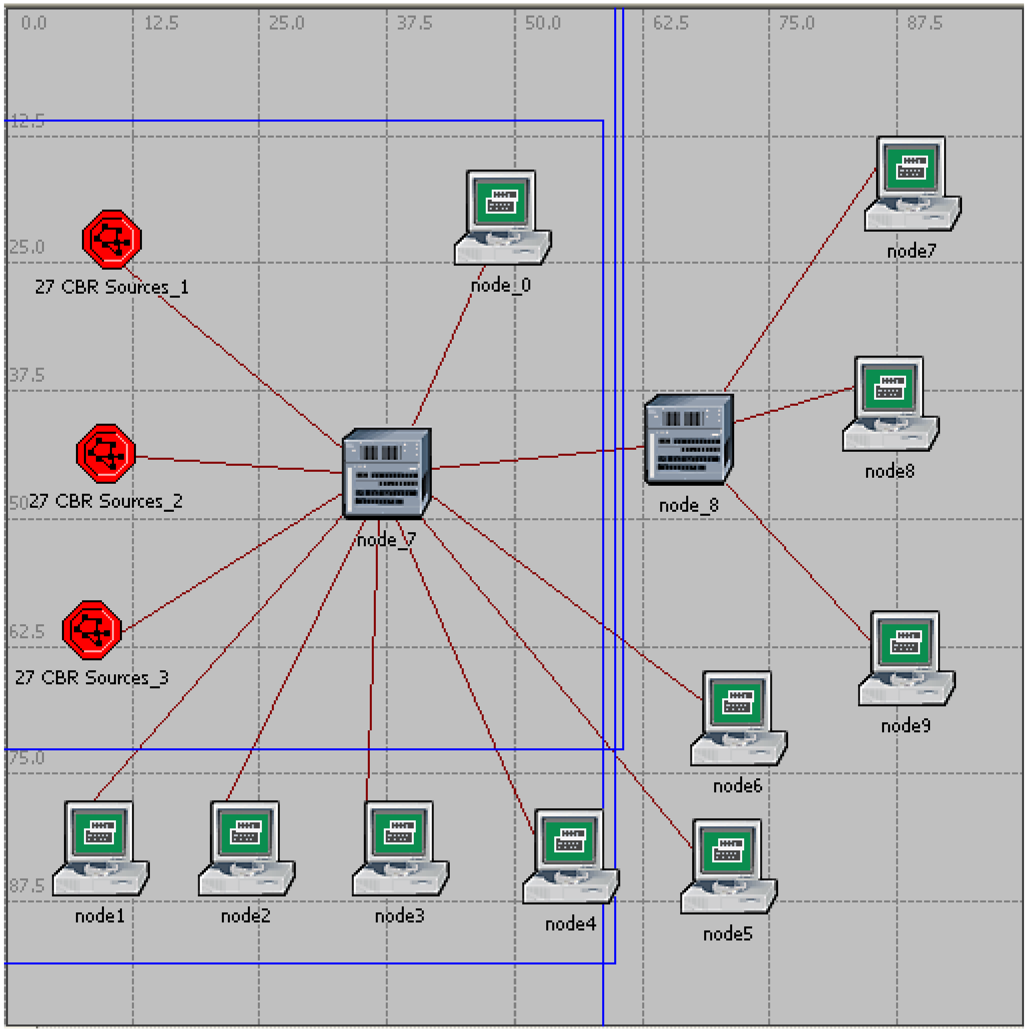

Switch to Switch link utilization = 50%

Each subnet (the red octagonal icons labeled 27 CBR Sources_1, _2, and _3 represent a stage 3 network as shown on the previous slide

Simulation Cases (Cont.)

❑Cases 2 and 5

- 7 switch-to-switch hops
- 1500 byte packets
- 9 time-sensitive traffic sources, with 2 (from nodes 1 and 2) going through all 7 hops and 7 getting added at successive nodes and dropped after 1 hop
- Same nominal rates as in case 1; 50 % switch to switch link utilization
- Case 2 – streams have frequency offsets within ± 100 ppm
- Case 5 – streams are synchronized in frequency and have phase relations chosen to produce worst-case delays
  - Node 2 and 3 sources turned on at 5 s
  - Node 1 source turned on at 5.000001 s
  - Node 4 source turned on at 5.000369 s
  - Node 6 source turned on at 5.000613 s
  - Node 8 source turned on at 5.000860 s
  - Node 10 source turned on at 5.001100 s
  - Node 12 source turned on at 5.001340 s
  - Node 14 source turned on at 5.001582 s

Simulation Cases 2 and 5

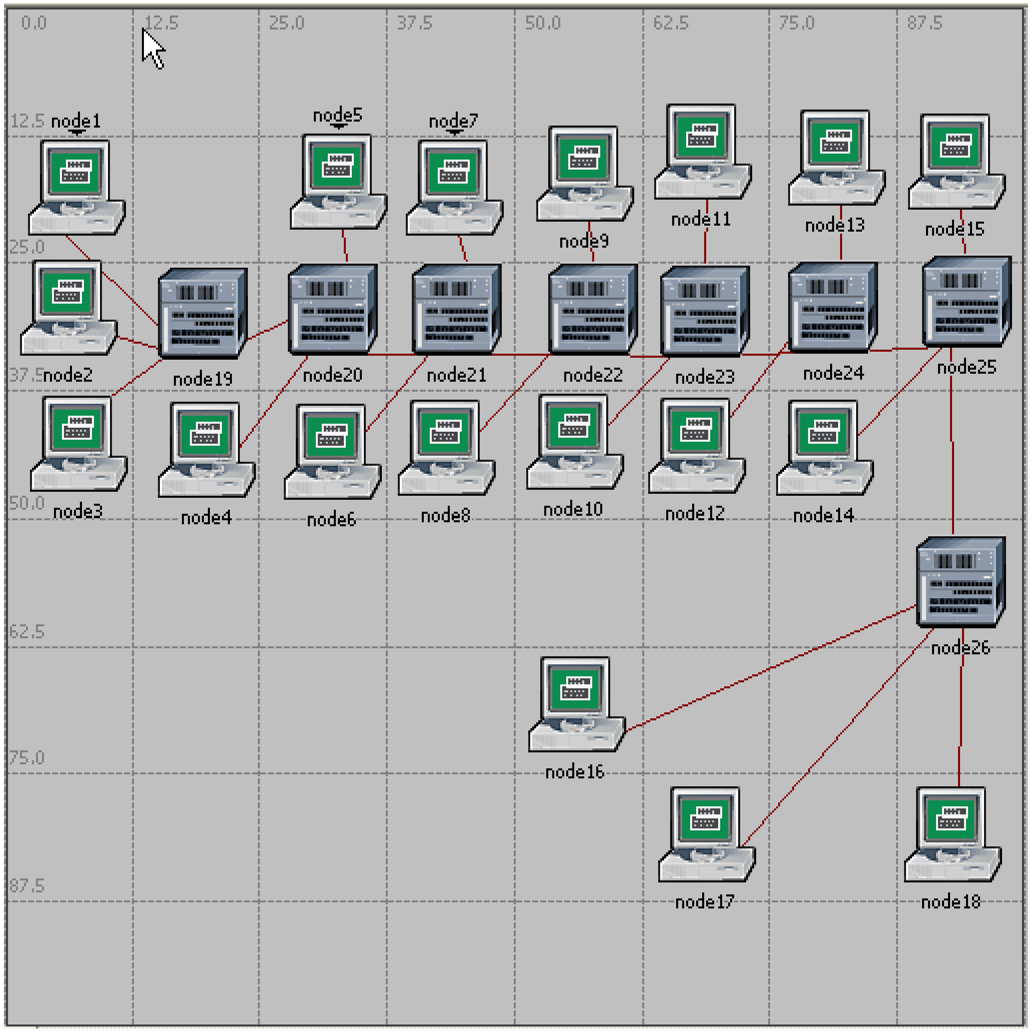

Switch to Switch link utilization = 50%

Simulation Cases (Cont.)

- Simulation Case 3
  - Similar to case 2, but now have 2 sources at each intermediate node getting added and then dropped after 1 hop
  - Only traffic from node 1 goes all the way through
  - 50 % switch to switch link utilization
- Simulation Case 4
  - Similar to case 2, but now have added best-effort streams at each switch that go one hop and are dropped
    - Best effort stream is Poisson, with mean arrival rate of 2083.33 packets/s (mean packet inter-arrival time = 0.00048 s)
    - 1500 byte packets
    - Switch to switch utilization= 75%
- See backup slides for more detail on description of simulation cases

Simulation Case 3

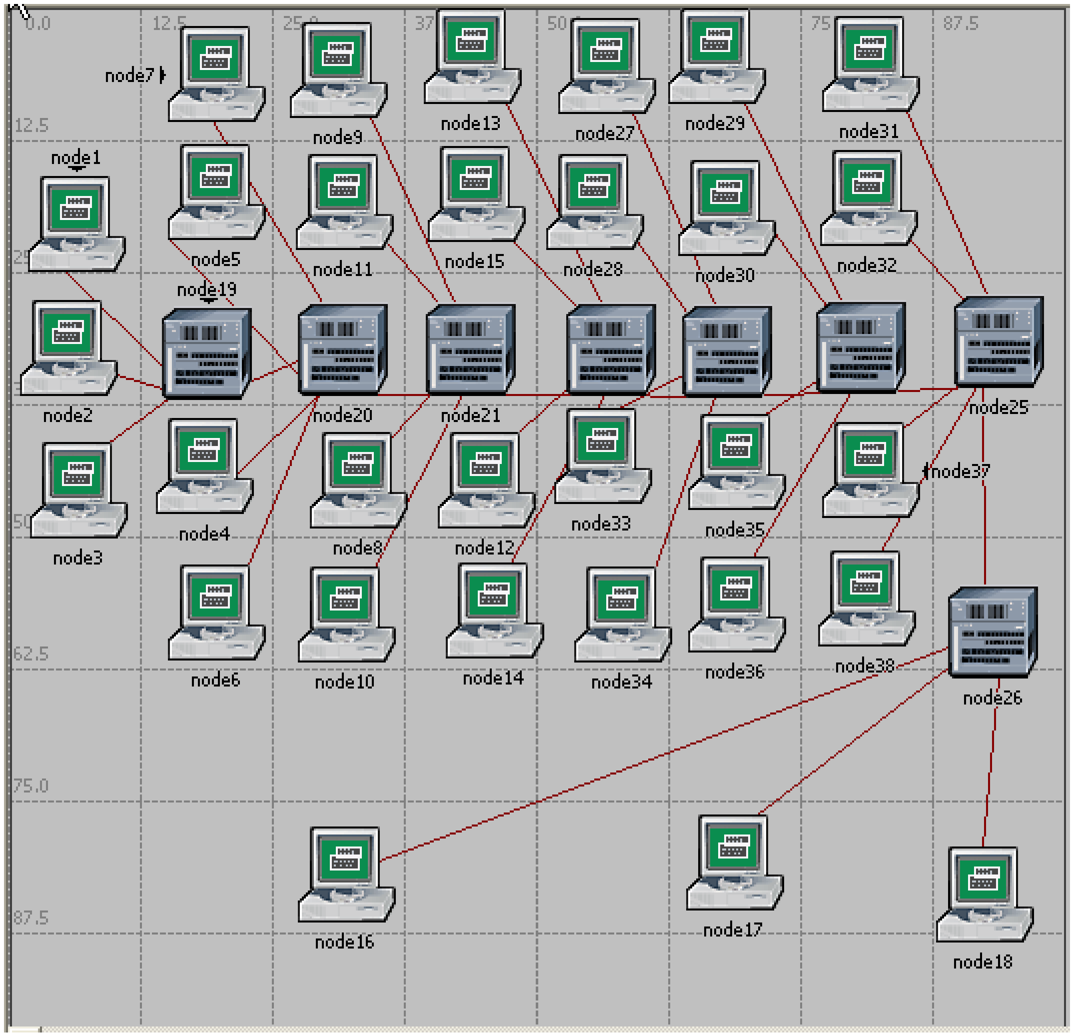

Switch to Switch link utilization = 50%

Simulation Case 4

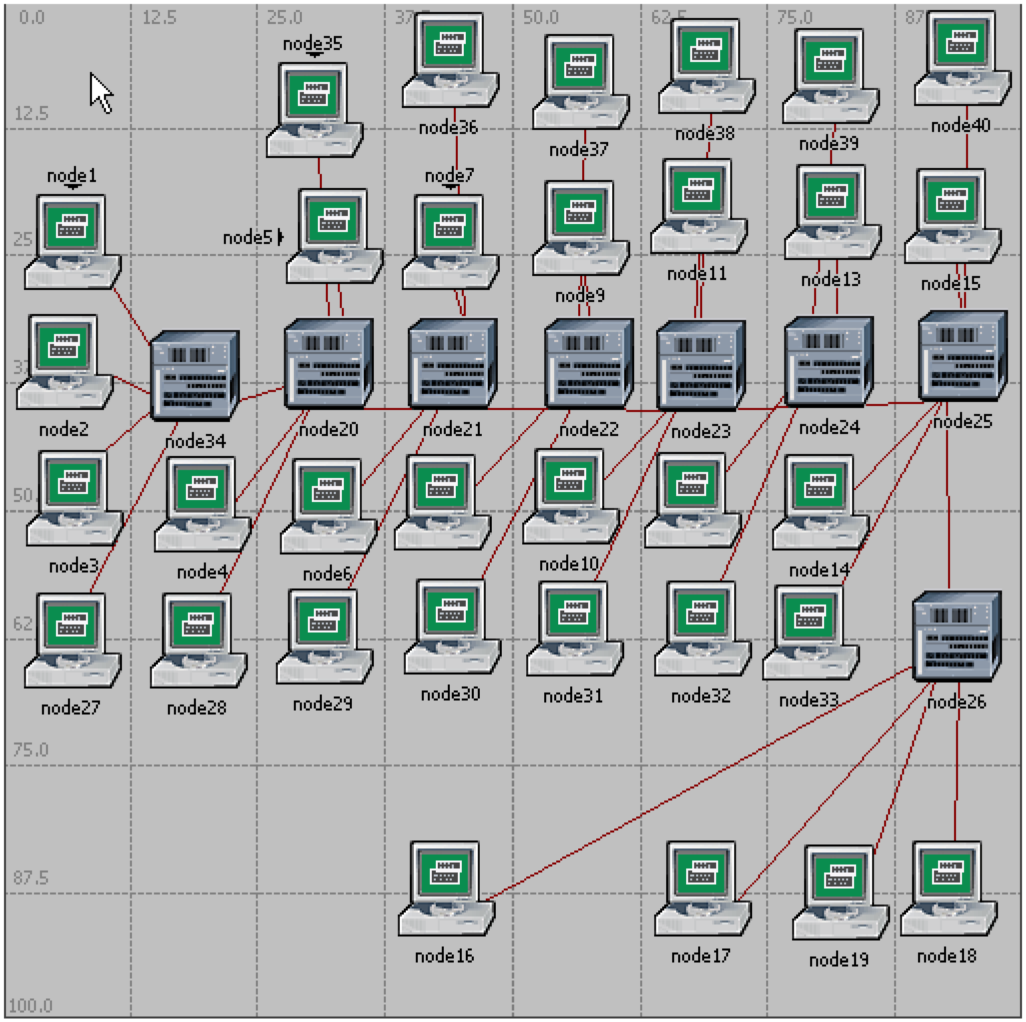

Switch to Switch link utilization = 75%

Simulation Results -- Cases 1 - 5

Case	Number of Hops	Peak Delay (ms)	Peak-to-Peak Delay Variation (μs)
1	1	0.62	250
1	2	1.0	500
1	3	1.1	510
1	4	1.6	850
2	1	0.63	260
2	7	1.75	650
3	1	0.62	250
3	7	1.89	790
4	1	0.73	350
4	7	2.2	1100
5	1	0.37	0
5	7	2.1	0

Conclusions

- The Case 1 results indicate the following rules of thumb may be used to estimate worst-case delay and worst-case peak-to-peak delay variation for an arbitrary N hop path through a network
  - Worst case end-to-end delay = $[(N + 2) + \Sigma_{j\,=\,1\text{ to number of switches}}(\{\text{number of incoming links at switch } j\} - 1)][\text{frame transmission delay}]$
    - Assumes propagation delay is negligible (must be added if it is not negligible)
  - Worst case peak-to-peak delay variation = $[\Sigma_{j\,=\,1\text{ to number of switches}}(\{\text{number of incoming links at switch } j\} - 1)][\text{frame transmission delay}]$
- Rules of thumb give worst-case delay and delay variation
  - Apply to cases with only time-sensitive (CBR) traffic, and cases with both CBR and best-effort traffic
- For link utilizations of 50 – 75%, worst-case delay and delay variation are reached after reasonable simulation time (e.g., up to a few hundred seconds) for a relatively small number of contending traffic streams (e.g., 5 or fewer)
  - For a larger number of contending traffic streams, it apparently takes a much longer simulation time to reach the theoretical maximum delay and delay variation predicted by the rules of thumb
- Can get worst-case delay predicted by rule of thumb for streams that are synchronized
  - Streams must be synchronized in frequency and have the right phase relationships

References

1. Geoffrey M. Garner and Felix Feng, *Delay Variation Simulation Results for Transport of Time-Sensitive Traffic over Conventional Ethernet*, Samsung presentation at July, 2005 IEEE 802.3 ResE SG meeting, San Francisco, CA, July 18, 2005.
2. Geoffrey M. Garner, *End-to-End Jitter and Wander Requirements for ResE Applications*, Samsung presentation at May, 2005 IEEE 802.3 ResE SG meeting, Austin, TX, May 16, 2005.
3. *Residential Ethernet (RE) (a working paper)*, Draft 0.136, maintained by David V. James and based on work by him and other contributors, August 10, 2005, available via http://www.ieee802.org/3/re_study/public/index.html

Backup Slides

Simulation Case 1

- Three sources at the ingress of each switch at the initial stage ($N = 3$)
- 4 stages ($k = 4$), i.e., a traffic stream that is not dropped at an intermediate stage traverses 4 switch-to-switch links (plus one ingress and one egress link)
  - Total of 81 traffic sources
    - We needed to restrict the numbers of hops and/or sources/switch to keep the total number of sources manageable (e.g., $N = 3$ and $k = 7$ would have produced 2187 traffic sources; $N = 6$ and $k = 7$ would have produced 279936 traffic sources.)
- All traffic is time-sensitive
  - Packet size is as given above (1500 bytes plus Ethernet overhead added by Opnet)
  - Nominal packet arrival rate for each stream is 1333.33 packets/s
    - Nominal time between packets is 0.00075 s
    - Resulting switch-to-switch link utilization is approximately 50% (results from 3 traffic streams)
      - Utilization per stream assuming nominal arrival rate and excluding Ethernet overhead is 16%
- Network topology is shown on the next 3 slides
  - It was convenient to use the subnet capability of Opnet, due to the large number of traffic sources and hierarchical structure of the network

Simulation Case 1 (Cont.)

- Simulate for 255 s, with traffic turned on at 5 s
  - Needed to add small amount of best-effort traffic in reverse direction to ensure each destination node would be in the forwarding database of each switch (otherwise get flooding and link utilizations that exceed 100%)
    - Node_0 sink in stages 3 and 4 is used for some of this reverse traffic

Simulation Case 1 - Stages 1 and 2

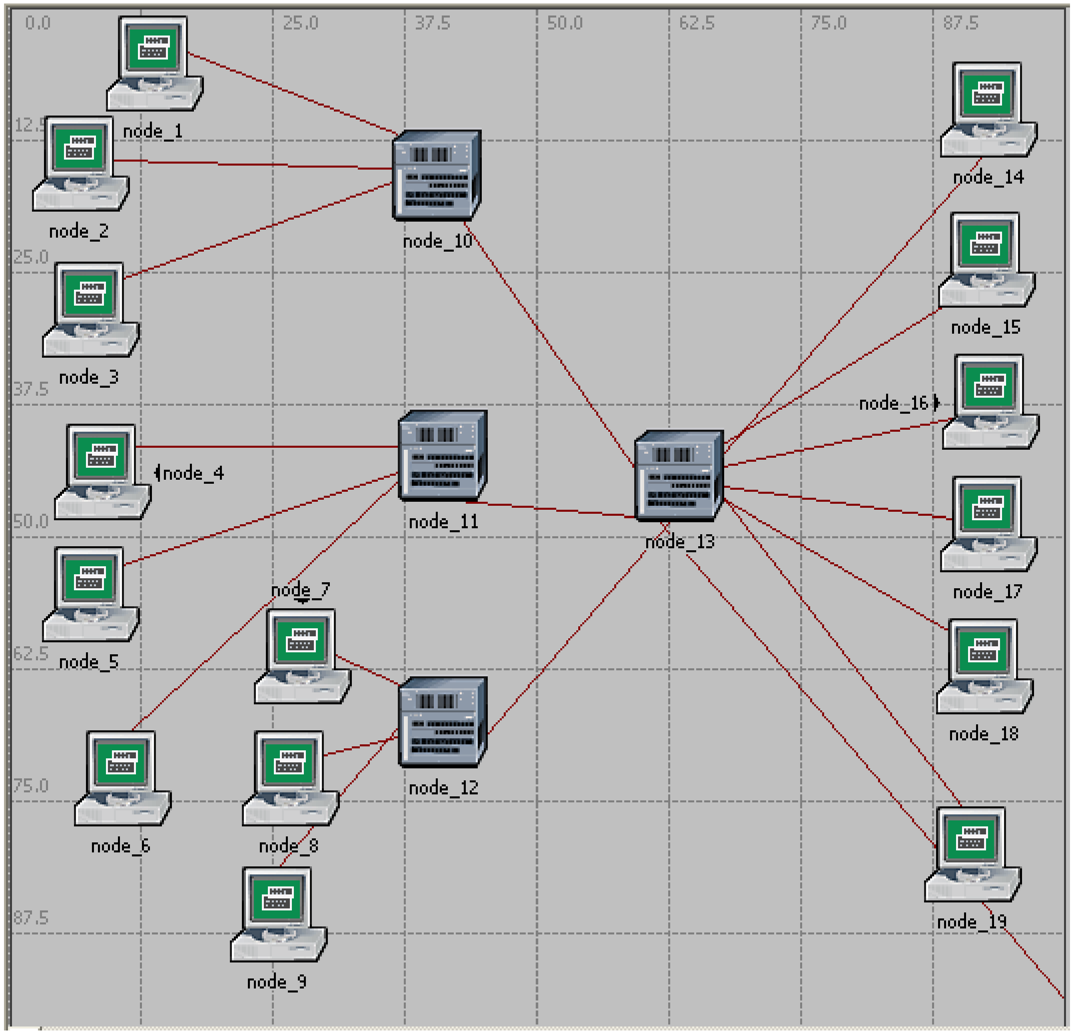

Switch to Switch link utilization = 50%

Simulation Case 1 - Stage 3

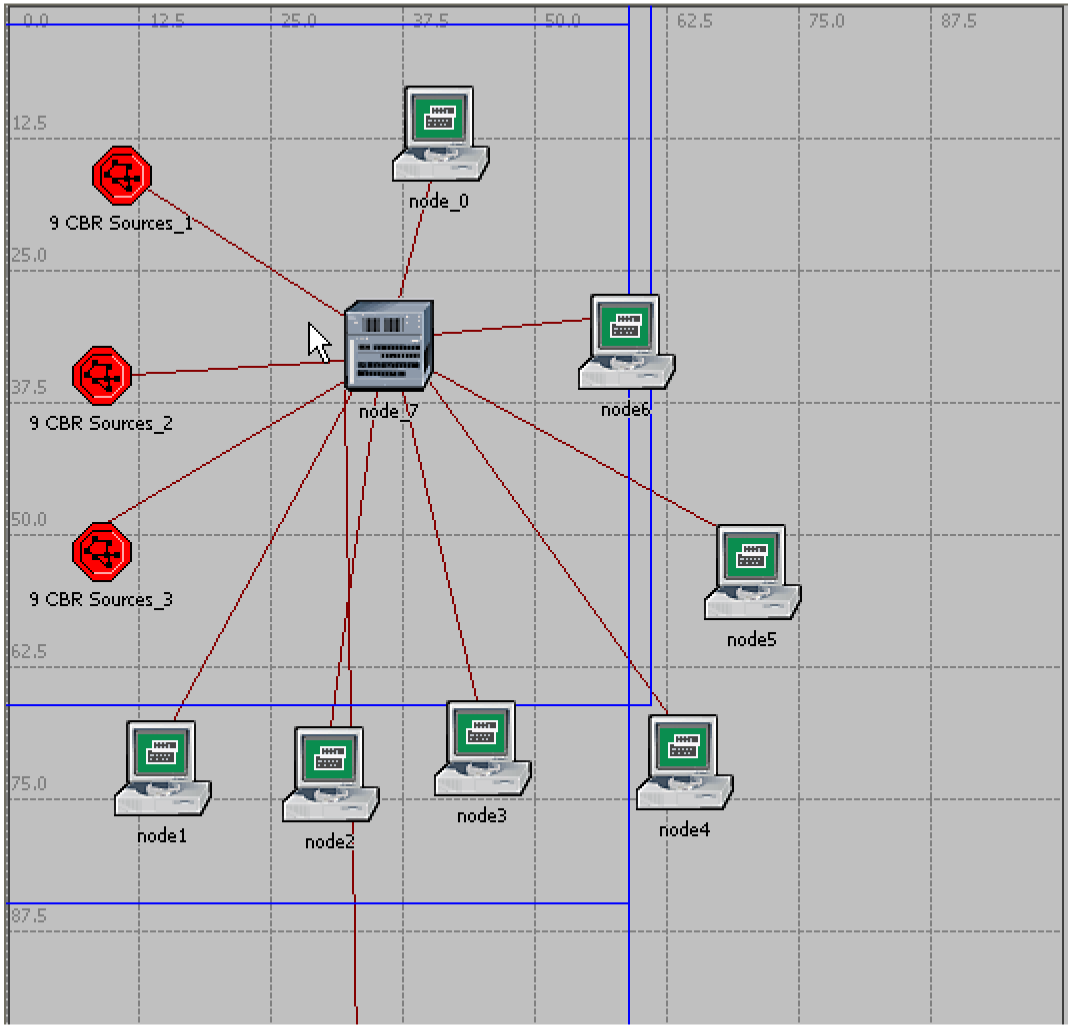

Switch to Switch link utilization = 50%

Each subnet (the red octagonal icons labeled 9 CBR Sources_1, _2, and _3 represent a stage 1 and 2 network as shown on the previous slide

Simulation Case 1 - Stage 4

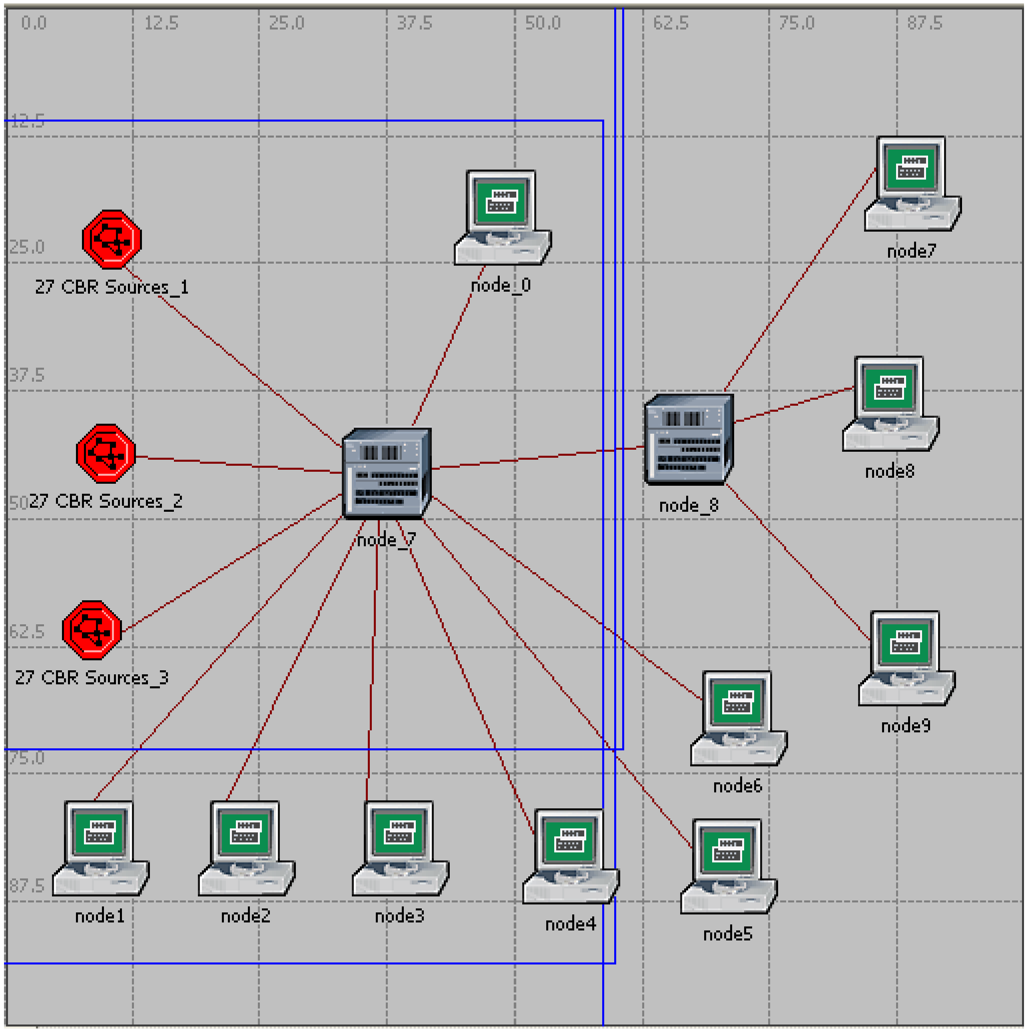

Switch to Switch link utilization = 50%

Each subnet (the red octagonal icons labeled 27 CBR Sources_1, _2, and _3 represent a stage 3 network as shown on the previous slide

Simulation Case 1 Traffic Streams – Stages 1 and 2

- Node 2 to node 14, rate offset by -100 ppm
- Node 3 to node 15, rate offset by +100 ppm
- Node 5 to node 16, rate offset by -50 ppm
- Node 6 to node 17, rate offset by +50 ppm
- Node 8 to node 18, rate offset by -75 ppm
- Node 9 to node 19, rate offset by +75 ppm
- Streams from nodes 1, 4, and 7 are transported on link to stage 3 (see following slides for details)

Simulation Case 1 Traffic Streams – Stage 3

- 27 CBR Sources_1 subnet
  - 9 CBR Sources_1 subnet, node 4 to node 1, rate offset by –10 ppm
  - 9 CBR Sources_1 subnet, node 7 to node 2, rate offset by +10 ppm
  - 9 CBR Sources_2 subnet, node 4 to node 3, nominal rate
  - 9 CBR Sources_2 subnet, node 7 to node 4, rate offset by +10 ppm
  - 9 CBR Sources_3 subnet, node 4 to node 5, rate offset by –10 ppm
  - 9 CBR Sources_3 subnet, node 7 to node 6, nominal rate
- 27 CBR Sources_2 subnet
  - 9 CBR Sources_1 subnet, node 4 to node 1, rate offset by –10 ppm
  - 9 CBR Sources_1 subnet, node 7 to node 2, rate offset by +10 ppm
  - 9 CBR Sources_2 subnet, node 4 to node 3, nominal rate
  - 9 CBR Sources_2 subnet, node 7 to node 4, rate offset by +10 ppm
  - 9 CBR Sources_3 subnet, node 4 to node 5, rate offset by –10 ppm
  - 9 CBR Sources_3 subnet, node 7 to node 6, nominal rate

Simulation Case 1 Traffic Streams – Stage 3 (Cont.)

- 27 CBR Sources_3 subnet
  - 9 CBR Sources_1 subnet, node 4 to node 1, rate offset by –10 ppm
  - 9 CBR Sources_1 subnet, node 7 to node 2, rate offset by +10 ppm
  - 9 CBR Sources_2 subnet, node 4 to node 3, nominal rate
  - 9 CBR Sources_2 subnet, node 7 to node 4, nominal rate
  - 9 CBR Sources_3 subnet, node 4 to node 5, rate offset by –10 ppm
  - 9 CBR Sources_3 subnet, node 7 to node 6, nominal rate
- Note that the 3rd and 4th streams are different in all three 27 CBR subnets

Simulation Case 1 Traffic Streams – Stage 4

- 27 CBR Sources_1, 9 CBR Sources_1 subnet, node 1 to node 7, nominal rate
- 27 CBR Sources_1, 9 CBR Sources_2 subnet, node 1 to node 1, rate offset by –10 ppm
- 27 CBR Sources_1, 9 CBR Sources_3 subnet, node 1 to node 2, rate offset by +10 ppm
- 27 CBR Sources_2, 9 CBR Sources_1 subnet, node 1 to node 8, rate offset by –10 ppm
- 27 CBR Sources_2, 9 CBR Sources_2 subnet, node 1 to node 3, rate offset by +10 ppm
- 27 CBR Sources_2, 9 CBR Sources_3 subnet, node 1 to node 4, nominal rate
- 27 CBR Sources_3, 9 CBR Sources_1 subnet, node 1 to node 9, rate offset by +10 ppm
- 27 CBR Sources_3, 9 CBR Sources_2 subnet, node 1 to node 5, rate offset by –10 ppm
- 27 CBR Sources_3, 9 CBR Sources_3 subnet, node 1 to node 6, nominal rate

Simulation Case 1 Stage 4 Results for 4-Hop Streams

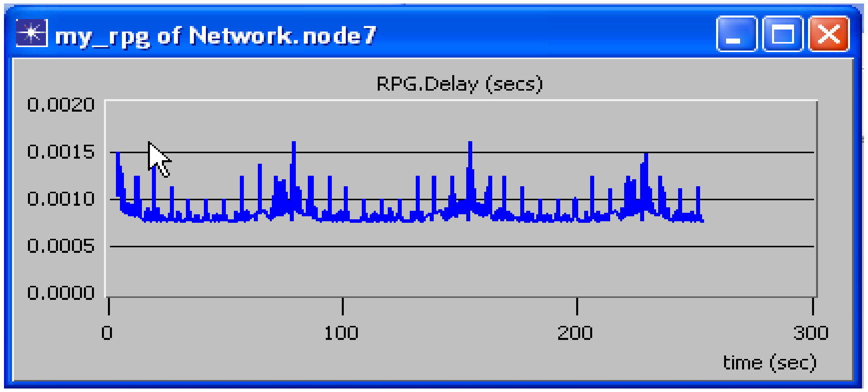

- ❑27 CBR Sources_1, 9 CBR Sources_1 subnet, node 1 to node 7, nominal rate

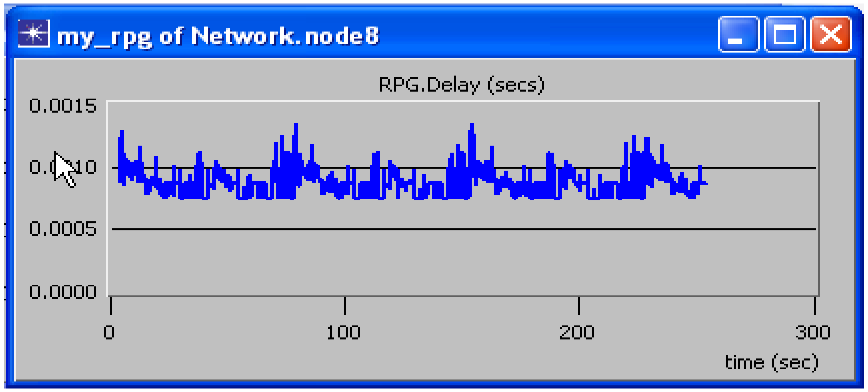

- ❑27 CBR Sources_2, 9 CBR Sources_1 subnet, node 1 to node 8, rate offset by –10 ppm

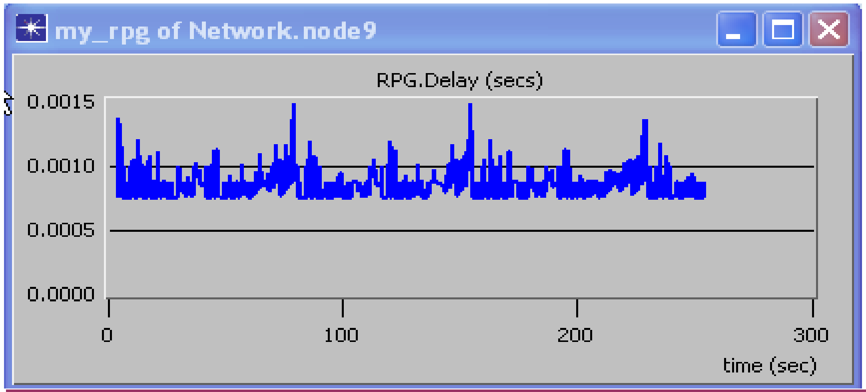

- ❑27 CBR Sources_3, 9 CBR Sources_1 subnet, node 1 to node 9, rate offset by +10 ppm

Simulation Case 1 Stage 3 Results for 3-Hop Streams

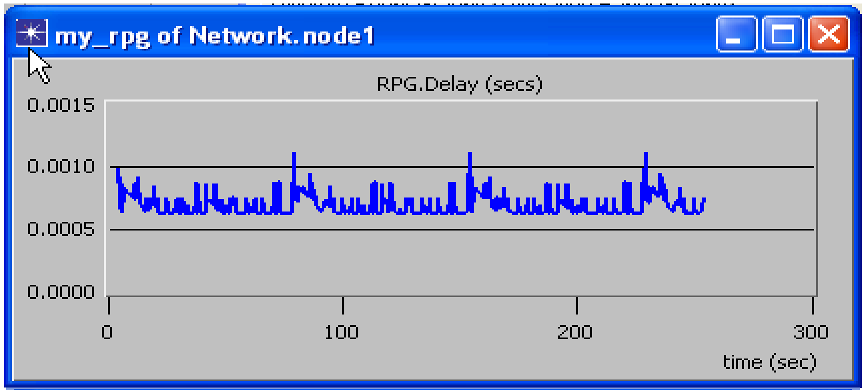

- ❑27 CBR Sources_1, 9 CBR Sources_2 subnet, node 1 to node 1, rate offset by –10 ppm

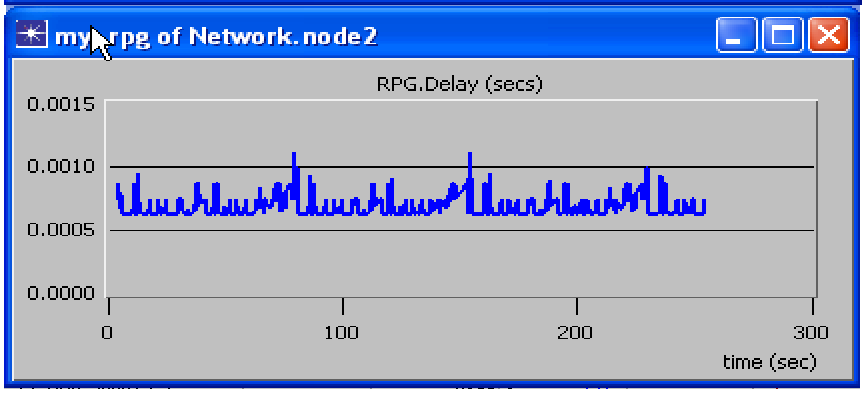

- ❑27 CBR Sources_1, 9 CBR Sources_3 subnet, node 1 to node 2, rate offset by +10 ppm

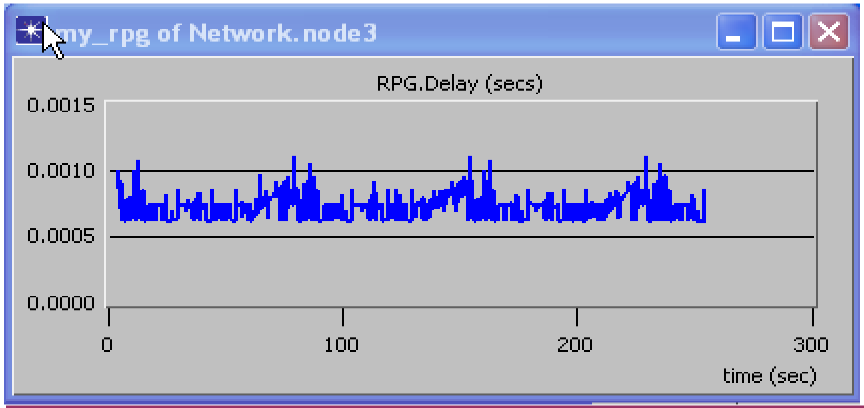

- ❑27 CBR Sources_2, 9 CBR Sources_2 subnet, node 1 to node 3, rate offset by +10 ppm

Simulation Case 1 Stage 3 Results for 3-Hop Streams

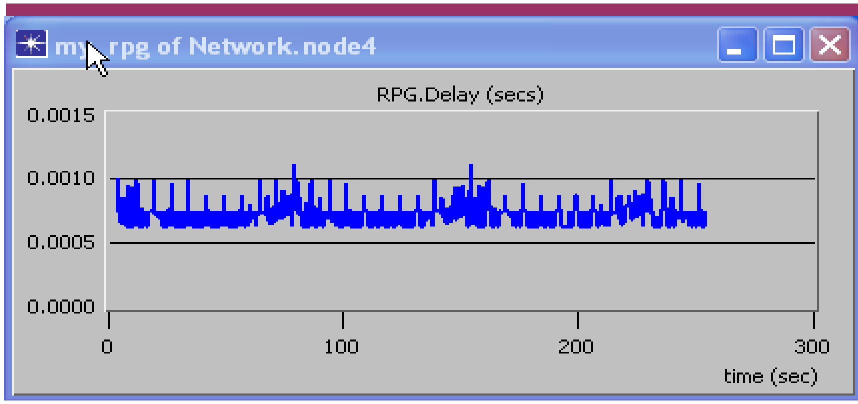

- ❑27 CBR Sources_2, 9 CBR Sources_3 subnet, node 1 to node 4, nominal rate

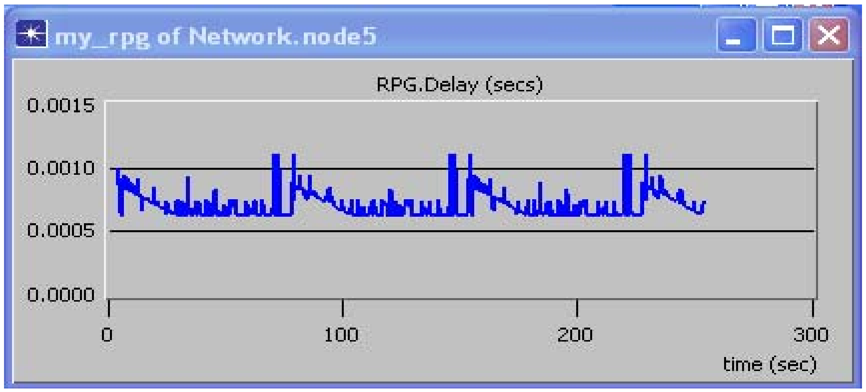

- ❑27 CBR Sources_3, 9 CBR Sources_2 subnet, node 1 to node 5, rate offset by –10 ppm

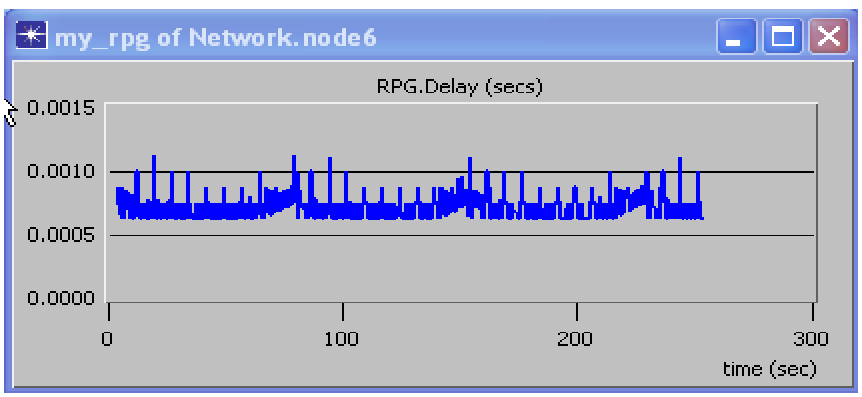

- ❑27 CBR Sources_3, 9 CBR Sources_3 subnet, node 1 to node 6, nominal rate

Simulation Case 1 Stage 2 Results for 2-Hop Streams

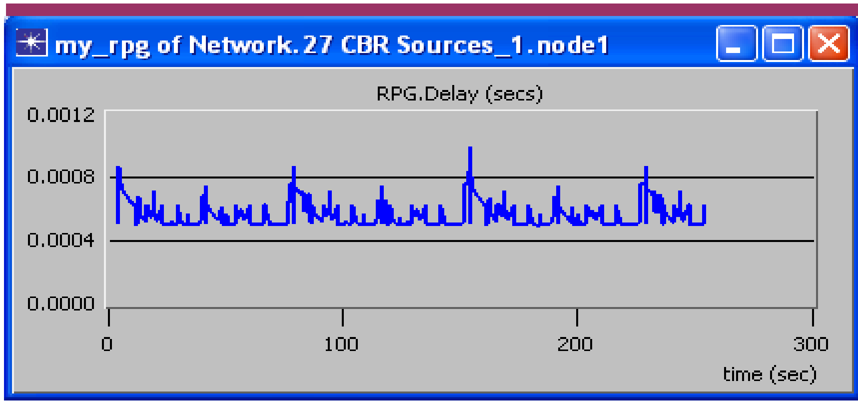

- ❑27 CBR Sources_1, 9 CBR Sources_1 subnet, node 4 to node 1, rate offset by –10 ppm

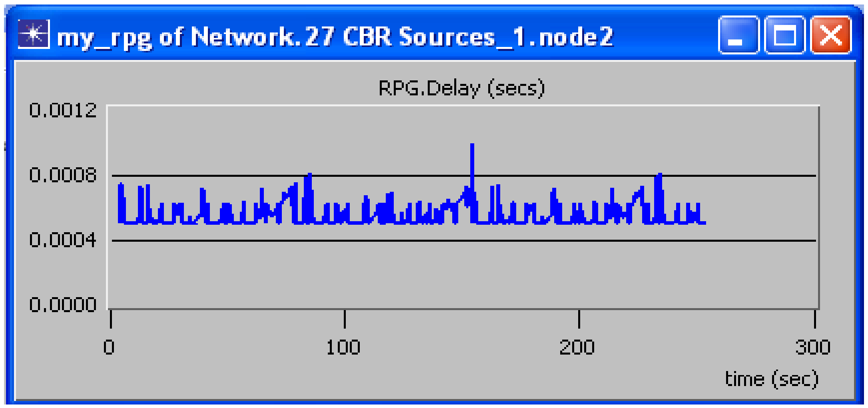

- ❑27 CBR Sources_1, 9 CBR Sources_1 subnet, node 7 to node 2, rate offset by +10 ppm

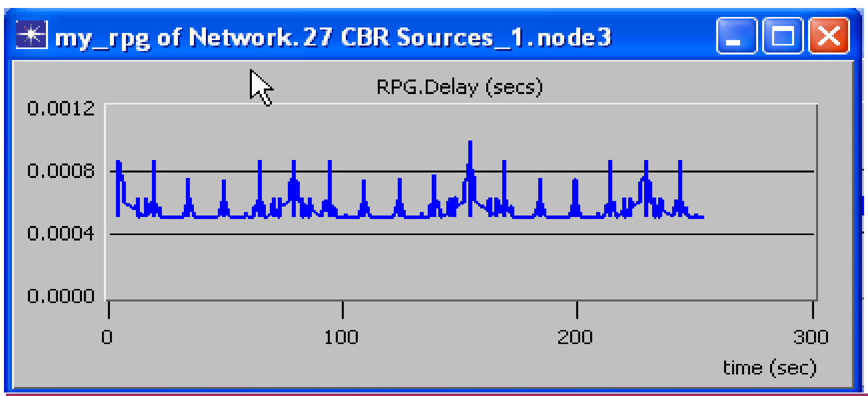

- ❑27 CBR Sources_1, 9 CBR Sources_2 subnet, node 4 to node 3, nominal rate

Simulation Case 1 Stage 2 Results for 2-Hop Streams

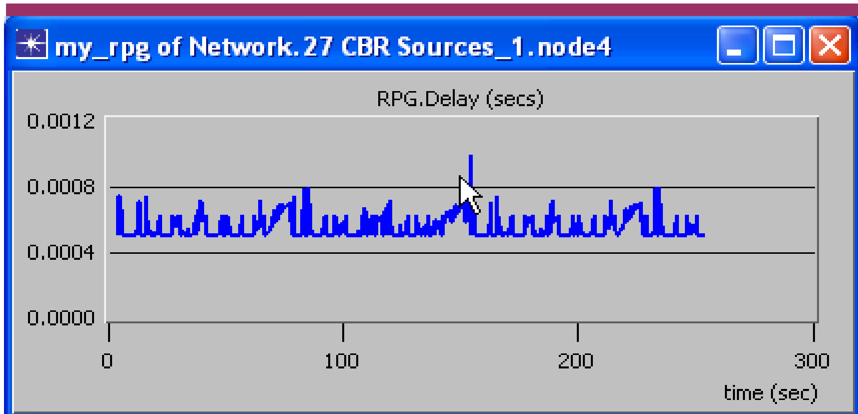

- 27 CBR Sources_1, 9 CBR Sources_2 subnet, node 7 to node 4, rate offset by +10 ppm

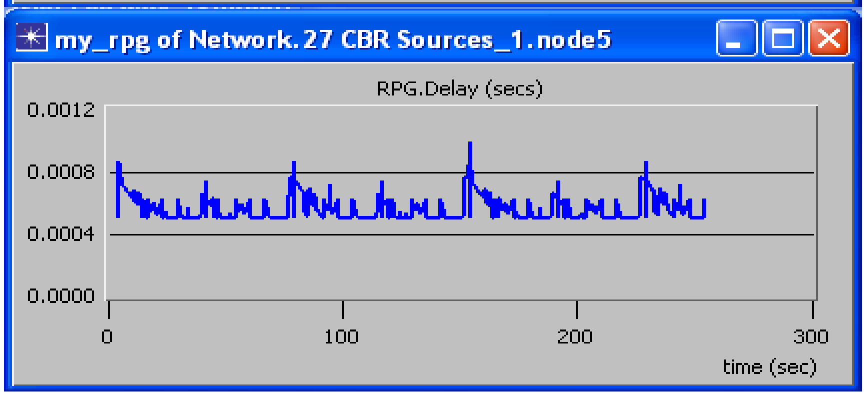

- 27 CBR Sources_1, 9 CBR Sources_3 subnet, node 4 to node 5, rate offset by –10 ppm

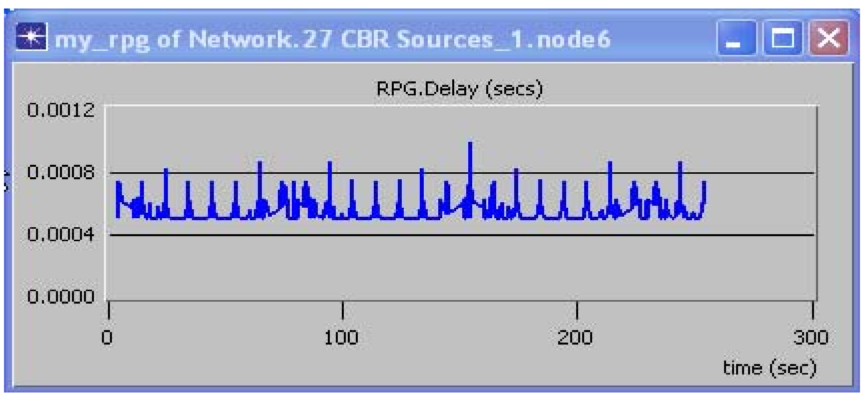

- 27 CBR Sources_1, 9 CBR Sources_3 subnet, node 7 to node 6, nominal rate

Simulation Case 1 Stage 2 Results for 2-Hop Streams

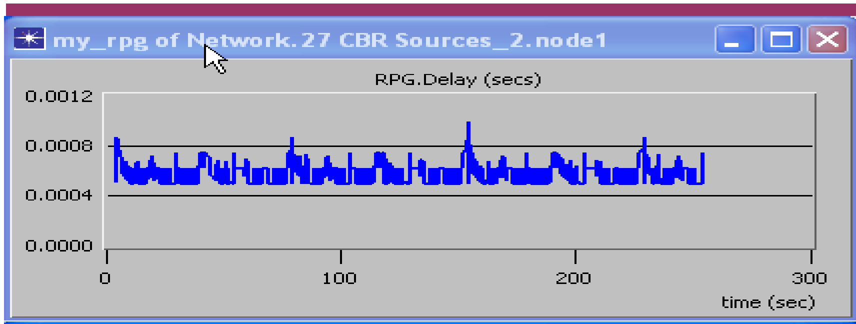

- ❑27 CBR Sources_2, 9 CBR Sources_1 subnet, node 4 to node 1, rate offset by –10 ppm

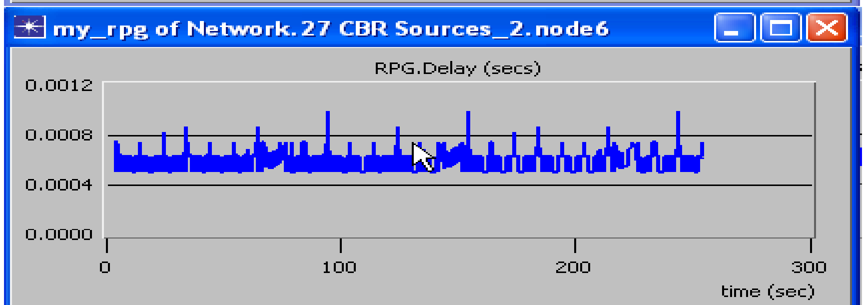

- ❑27 CBR Sources_2, 9 CBR Sources_3 subnet, node 7 to node 6, nominal rate

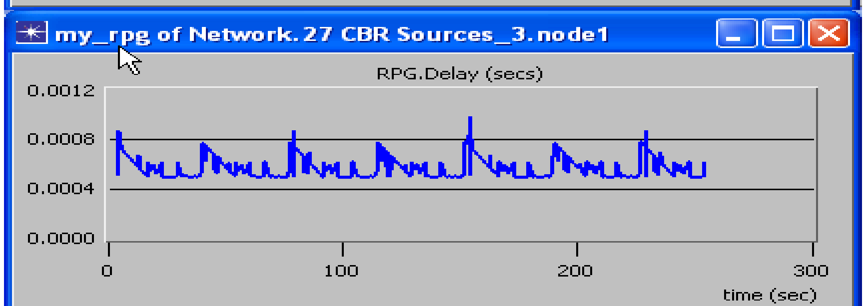

- ❑27 CBR Sources_3, 9 CBR Sources_1 subnet, node 4 to node 1, rate offset by –10 ppm

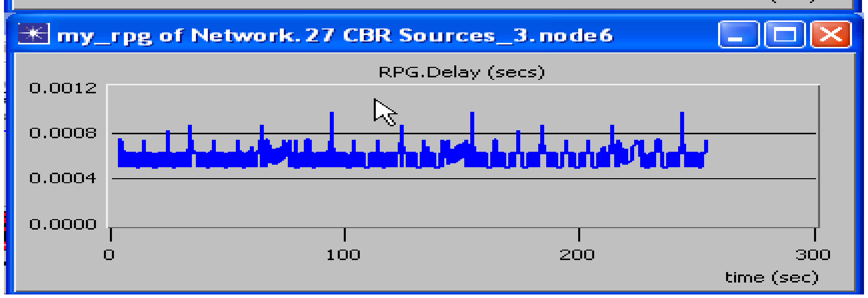

- ❑27 CBR Sources_3, 9 CBR Sources_3 subnet, node 7 to node 6, nominal rate

Simulation Case 1 Stage 1 Results for 1-Hop Streams

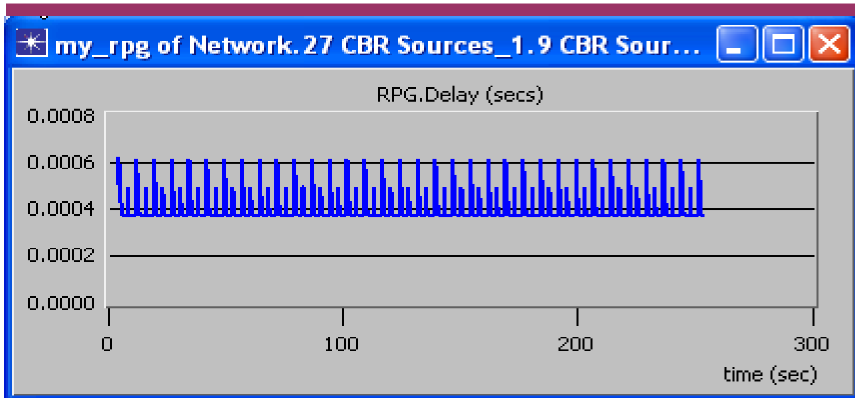

- 27 CBR Sources_1, 9 CBR Sources_1 subnet, node 2 to node 14, rate offset by –100 ppm

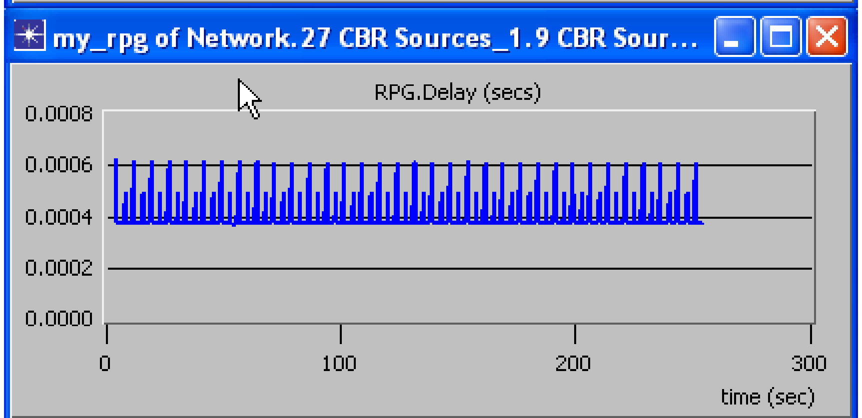

- 27 CBR Sources_1, 9 CBR Sources_1 subnet, node 3 to node 15, rate offset by +100 ppm

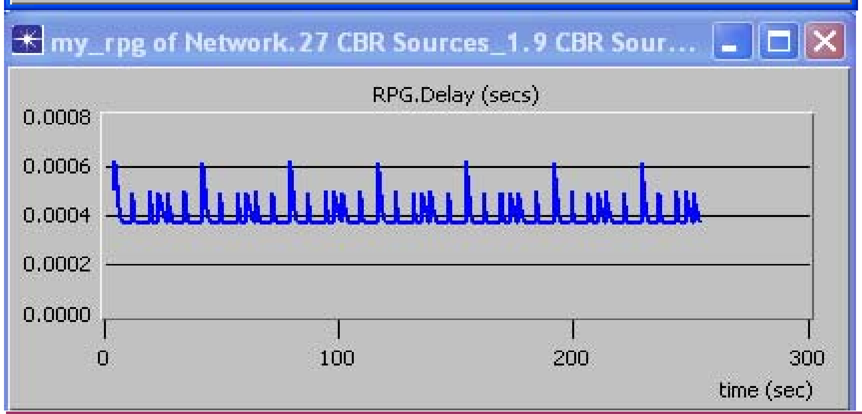

- 27 CBR Sources_1, 9 CBR Sources_1 subnet, node 5 to node 16, rate offset by –50 ppm

Simulation Case 1 Stage 1 Results for 1-Hop Streams

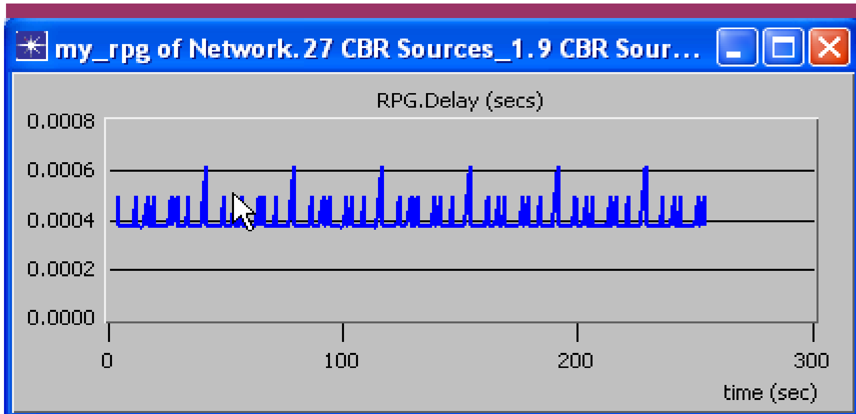

- ❑27 CBR Sources_1, 9 CBR Sources_1 subnet, node 6 to node 17, rate offset by +50 ppm

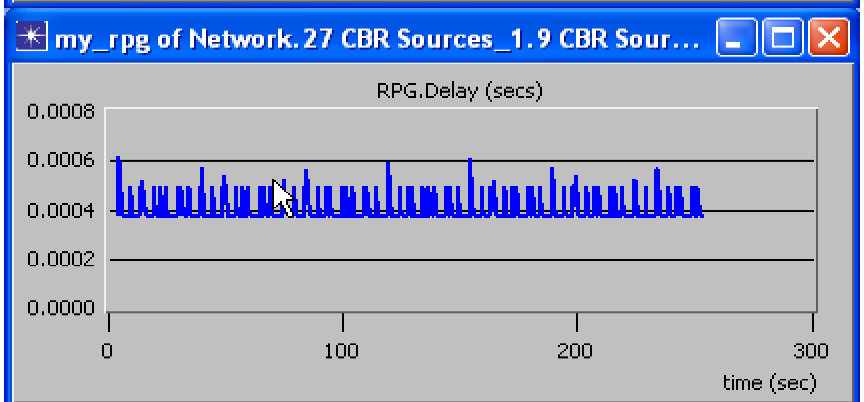

- ❑27 CBR Sources_1, 9 CBR Sources_1 subnet, node 8 to node 18, rate offset by –75 ppm

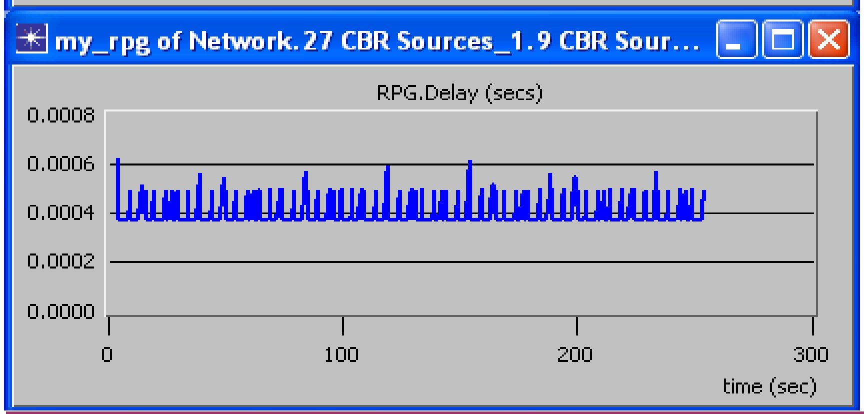

- ❑27 CBR Sources_1, 9 CBR Sources_1 subnet, node 9 to node 19, rate offset by +75 ppm

Simulation Case 1 Stage 1 Results for 1-Hop Streams

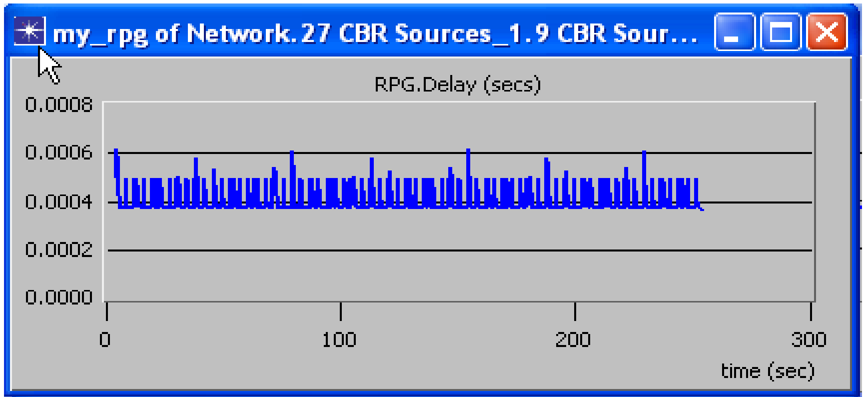

- ❑27 CBR Sources_1, 9 CBR Sources_2 subnet, node 2 to node 14, rate offset by –100 ppm

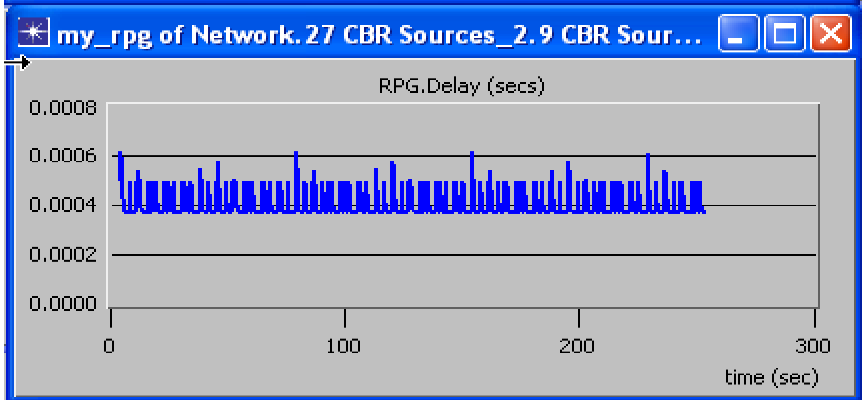

- ❑27 CBR Sources_2, 9 CBR Sources_2 subnet, node 2 to node 14, rate offset by –100 ppm

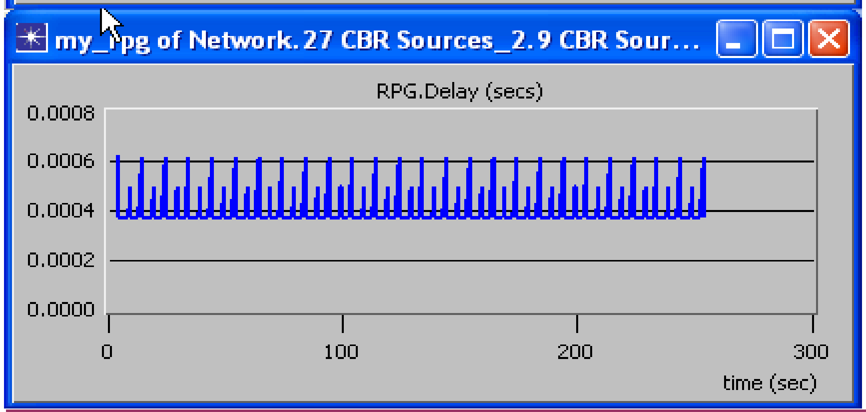

- ❑27 CBR Sources_2, 9 CBR Sources_2 subnet, node 9 to node 19, rate offset by +75 ppm

Simulation Case 1 Stage 1 Results for 1-Hop Streams

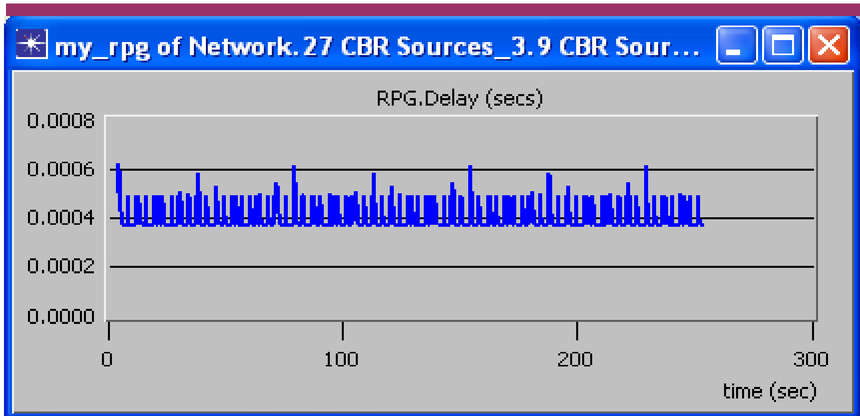

- ❑27 CBR Sources_3, 9 CBR Sources_2 subnet, node 2 to node 14, rate offset by –100 ppm

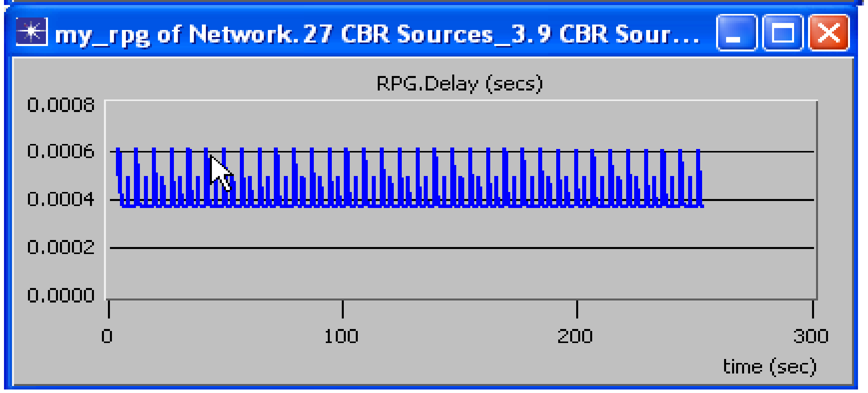

- ❑27 CBR Sources_3, 9 CBR Sources_2 subnet, node 9 to node 19, rate offset by +75 ppm

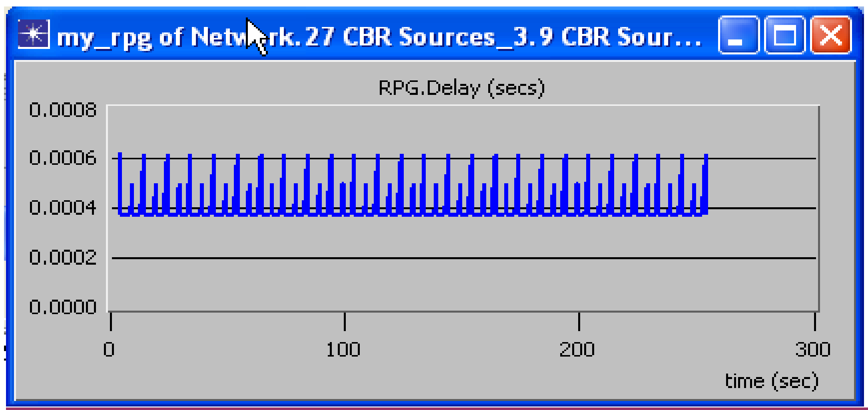

- ❑27 CBR Sources_3, 9 CBR Sources_3 subnet, node 2 to node 14, rate offset by –100 ppm

Summary of Peak Delay and Peak-to-Peak Delay Variation

Number of Hops	Peak Delay (ms)	Peak-to-Peak Delay Variation (μs)
1	0.62	250
2	1.0	500
3	1.1	510
4	1.6	850

Case 1 – Discussion of Results

- Peak delay and peak-to-peak delay variation increase with number of hops, as expected
  - Peak delay reaches 1.6 ms after 4 hops
  - Peak-to-peak delay variation reaches 850 μs after 4 hops
- Peak delay is slightly below worst-case that would be obtained for this 4-hop case
  - Worst-case for 3 contending CBR streams at a switch occurs when 2 frames are queued when a frame arrives (for link utilization < 100%)
  - Then, for contention occurring at 4 switches (in a 4-hop path), the delay due to contention is (4)(2)(frame transmission delay) = 8(frame transmission delay)
  - Also have transmission delay for the frame itself on the 4 switch-to-switch links plus the two access links
  - Then total delay due to transmission and queueing, in worst case, is 14(frame transmission delay)
  - Then worst-case delay (neglecting propagation delay since it is much smaller) is
    - $14(12000+8(38)\text{ bits})/10^8\text{ bits/s} = 1.72 \times 10^{-3}\text{ s} = 1.72\text{ ms}$

Case 1 – Discussion of Results (Cont.)

- If number of contending traffic streams is increased from 3 to 6 (still for a 4-hop case), would expect worst-case delay due to contention to increase by (4)(3)(frame transmission delay) = 12(12000+8(38) bits)/10^8 bits/s) = 1.48 ms
  - Worst-case total delay in this case would be 1.72+1.48 ms = 3.2 ms
- If number of hops is increased from 4 to 7 (still for 3 contending traffic streams), would expect worst-case delay to increase to [7(2)+9](frame transmission delay] (i.e., 2 contending frames at each of 7 switches plus 9 total transmission delays (switch-to-switch plus access links))
  - Worst case total delay in this case would be 23(12000+8(38) bits)/10^8 bits/s) = 2.83×10^{-3} s = 2.83 ms

Case 1 – Discussion of Results (Cont.)

- Peak-to-peak delay variation is slightly below worst-case that would be obtained for this 4-hop case
  - Worst-case peak-to-peak delay variation is equal to the worst-case delay due to contention at the switches, as this is the component of delay that is not always present
  - Then worst-case peak-to-peak delay variation is
    - 8(12000+8(38) bits)/10^8 bits/s) = 9.8×10^{-6} s = 980 μs
- Note that the amount by which the worst case peak delay exceeds the actual peak delay (1.72 ms – 1.6 ms = 0.12 ms) and the amount by which the worst case peak-to-peak delay variation exceeds the actual peak-to-peak delay variation (980 μs – 850 μs = 130 μs = 0.13 ms) are approximately equal, as expected
- Results obtained for 1, 2, and 3 hops are consistent with similar worst-case analyses for these cases (with 3-hop results below worst-case results by approximately 0.1 ms)
- While peak delay does not exceed 2 ms for 3 contending traffic streams and 4 hops, the results indicate that it will exceed 2 ms for 6 contending traffic streams with 4 hops, and 3 contending traffic streams with 7 hops

Case 1 – Discussion of Results (Cont.)

- Worst-case peak-to-peak delay variation is just below 1 ms (i.e., 980 μs) for 3 contending traffic streams with 4 hops
  - For 6 contending streams with 4 hops, this increases to 5(4)(frame transmission delay) = 2.46 ms
  - For 4 contending streams with 7 hops, this increases to 3(7)(frame transmission delay) = 2.58 ms
- Therefore, while peak-to-peak delay variation does not exceed 2 ms for 3 contending traffic streams and 4 hops, the results indicate that it will exceed 2 ms for 6 contending traffic streams with 4 hops, and 3 contending traffic streams with 7 hops

Rules of Thumb for Delay and Delay Variation

- The results indicate the following rules of thumb may be used to estimate worst-case delay and worst-case peak-to-peak delay variation for an arbitrary N hop path through a network
  - Worst case end-to-end delay = $[(N + 2) + \Sigma_{j\,=\,1\text{ to number of switches}}(\{\text{number of incoming links at switch } j\} - 1)][\text{frame transmission delay}]$
    - Assumes propagation delay is negligible (must be added if it is not negligible)
  - Worst case peak-to-peak delay variation = $[\Sigma_{j\,=\,1\text{ to number of switches}}(\{\text{number of incoming links at switch } j\} - 1)][\text{frame transmission delay}]$

Simulation Case 2

- 7 switch to switch hops
- 100 Mbit/s
- 9 traffic sources
- All sources are time sensitive, CBR traffic with nominal rate of 1333.33 packets/s (nominal time between packets = 0.00075 s)
- Sources have various different frequency offsets that are all within ±100 ppm
- Maximum size packets (1500 bytes plus Ethernet overhead)
  - Switch to switch link utilization ≈ 50%
  - Packet service time (including Ethernet overhead and inter-frame gap = (1500+38)(8)(10^{-8}) s = 0.12304 ms
- Network topology shown two slides following
  - 3 sources at first switch (nodes 1 – 3)
  - Traffic from 2 of these sources go to final switch (nodes 16 and 18)
  - Traffic from 3rd source (node 3) is dropped at 2nd switch
  - At switches 2 – 7 (nodes 20 – 25 in figure), traffic added from single CBR source, carried 1 hop, and dropped

Simulation Case 2 (Cont.)

- Simulate for 2405 s, with traffic turned on at 5 s
  - Needed to add small amount of best-effort traffic in reverse direction to ensure each destination node would be in the forwarding database of each switch (otherwise get flooding and link utilizations that exceed 100%)

Simulation Case 2

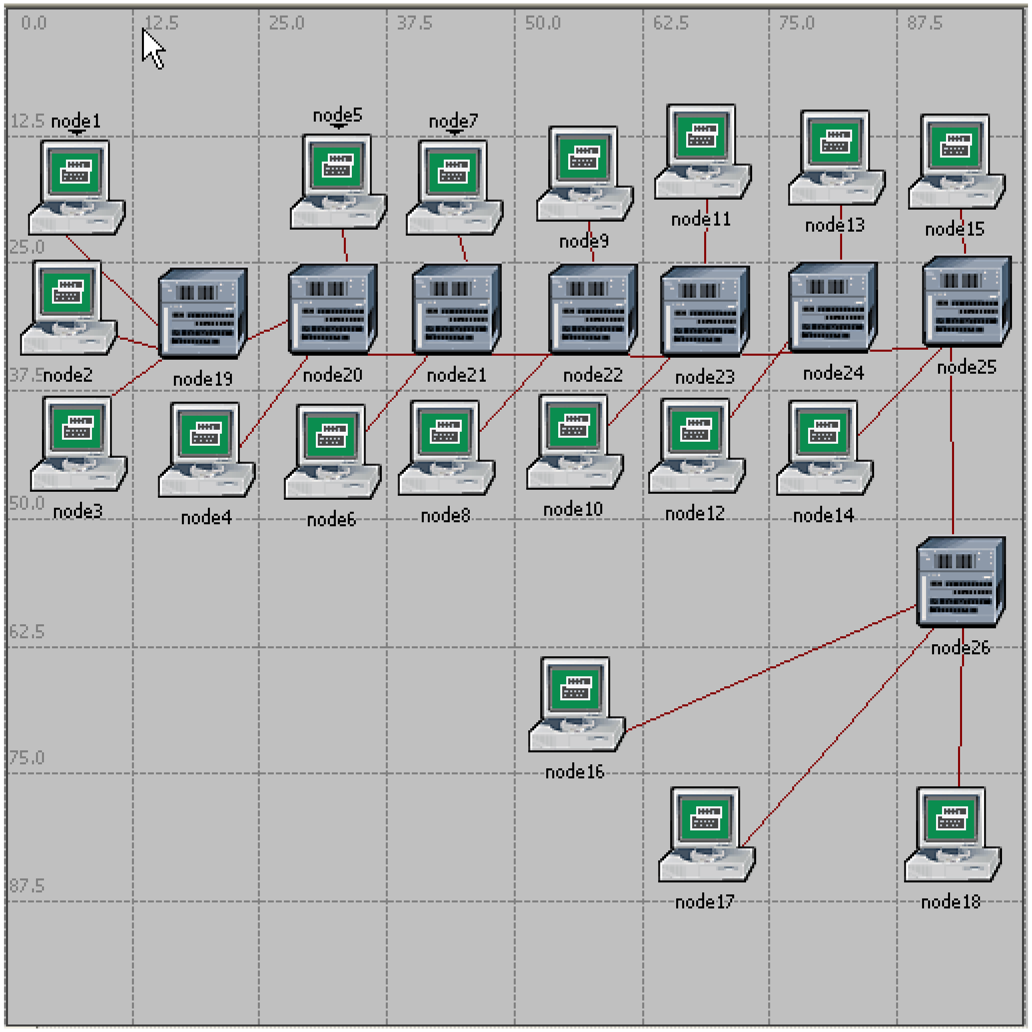

Switch to Switch link utilization = 50%

Case 2 - Results for Through Traffic Streams

Node 1 – 18

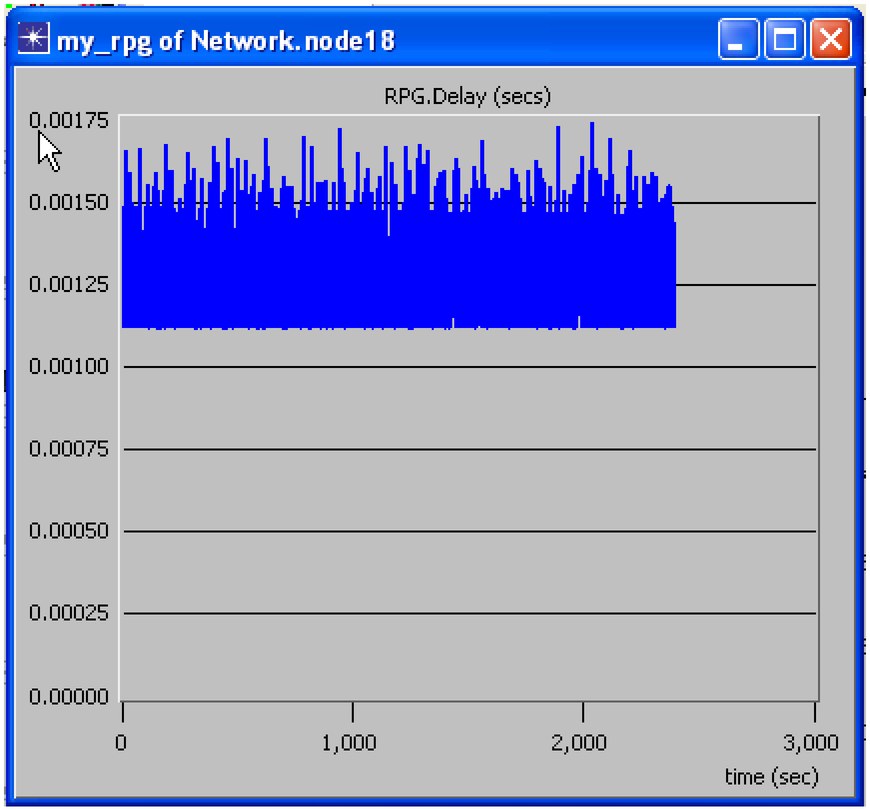

Node 2 – 16

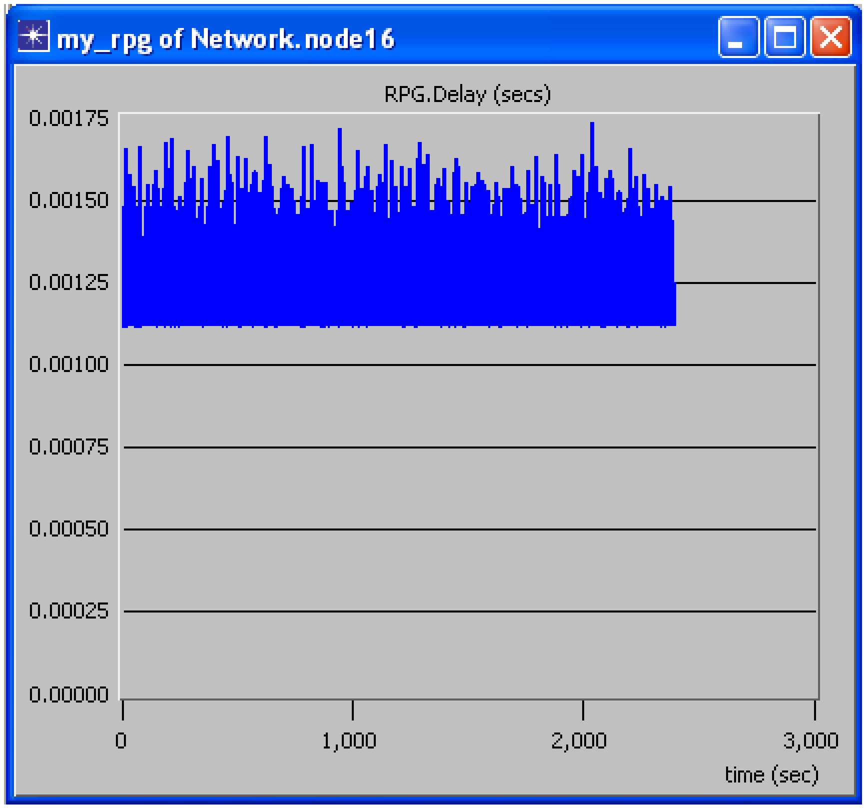

Case 2 - Results for 1-Hop Streams

Node 3 – 5

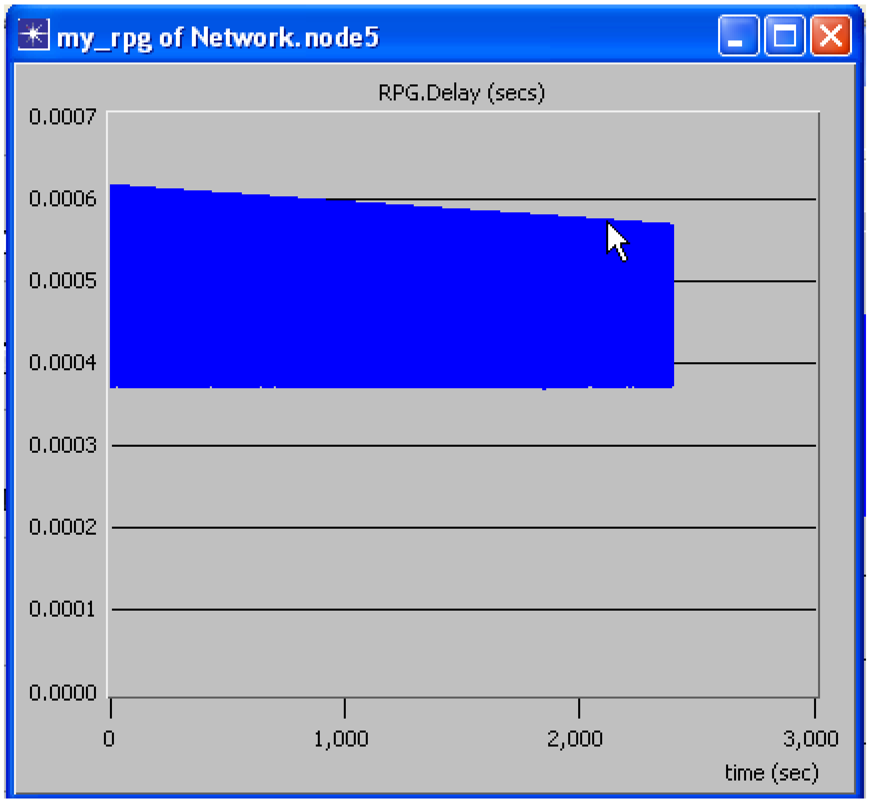

Node 12 – 15

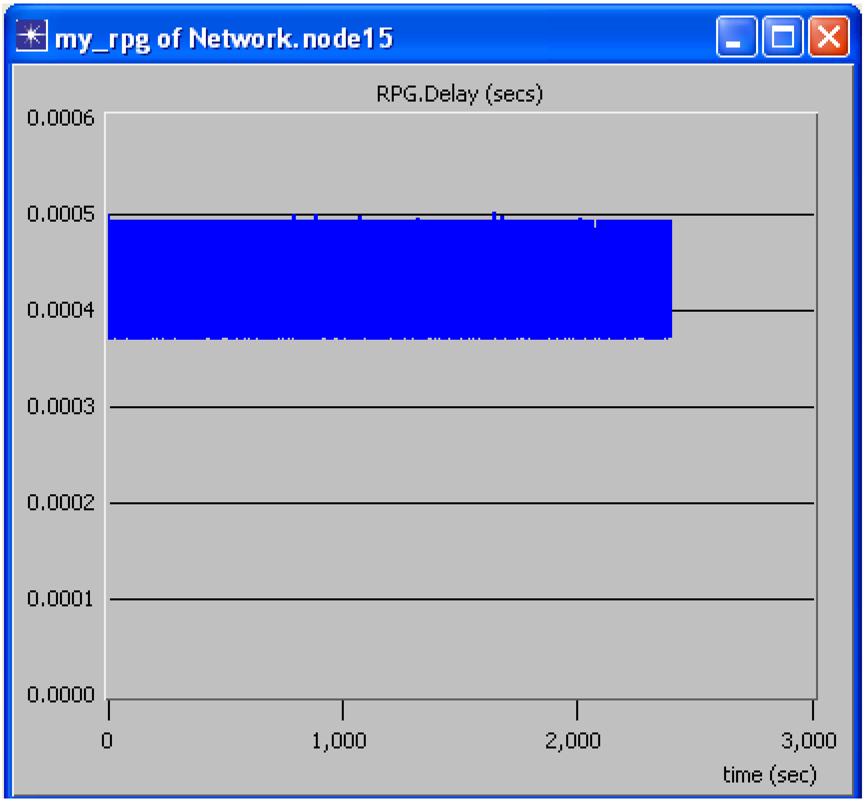

Case 2 Results

- Minimum delay for through streams = 1.1 ms
  - Approximately 9(packet service time) = 9(0.12304) ms as expected (7 interswitch hops plus 2 access hops)
- Maximum delay over 2400 s for through streams is approximately 1.75 ms
  - This is less than would be predicted by rule of thumb
  - Total of 8 contending streams
    - Rule of thumb would give (9+8)(0.12304 ms) = 2.09 ms
  - Other simulations with through stream from node 2 going to various intermediate nodes showed that maximum delay predicted by rule of thumb was not reached when number of contending streams reached 5 or 6
- Minimum delay for 1-hop streams = 0.37 ms (consistent with 3 hops)
- Maximum delay for Node 12-15 1-hop stream = 0.49 ms (consistent with 1 contending stream)
- Maximum delay for Node 3-5 1-hop stream = 0.63 ms
  - Roughly consistent with 2 contending streams, but not clear from figure if steady state has been reached
- Appears that must simulate for much longer to see full delay predicted by rule of thumb

Case 2 Results (Cont.)

- Nominal time between packets = 0.75 ms
- Packet service time = 0.12304 ms
  - Packet service time is 16% of interpacket interval for 1 stream
- To get worst-case delay predicted by rule of thumb, packets on all 8 contending streams must arrive at approximately same time
  - E.g., if they arrive within a time window equal to 5% of the packet time, the amount the actual delay will be less than the worst case delay will be at most 8(0.05 packet times) = 0.4 packet times
  - This means they must all arrive within a window of approximately (0.16)(0.4) inter-packet times = 0.064 inter-packet times
  - Due to the frequency offsets of the CBR streams, the packets of different streams are gaining/receding relative to each other
  - If we assume that, over a long time, the packet arrivals of any stream fall in all possible locations within the interpacket time of any other stream, and also assume that at any given time the location is random, then the probability that the packets of 8 streams all line up within 0.4 packet times is (0.064)^8 = 2.8×10^{-10}
  - Nominal number of packets observed for 1 stream over 2400 s = 2400/0.00075 = 3.2×10^{6}

Case 2 Results (Cont.)

- Then approximate probability of observing the packets of 8 streams lining up after 2400 s = $(3.2 \times 10^{6})(2.8 \times 10^{-10}) = 8.96 \times 10^{-4}$
- Would need to simulate for a time on the order of 2400 s/8.96×10^{-4} = 2.68×10^{6} s = 31 days to have a reasonable chance of observing such an event
- **Note:**
  - Previous results, for bursting and bunching case, also had 8 contending streams (2 contending streams at each of 4 switches)
    - Simulated delay (1.6 ms) was slightly less than theoretical maximum predicted by rule of thumb (1.72 ms)
    - Simulated peak-to-peak delay variation (850 μs) was slightly less than theoretical maximum predicted by rule of thumb (980 μs)
    - Previous simulated results for 3 hops (6 contending streams) were also slightly less than theoretical maximum predicted by rule of thumb
      - 1.1 ms
      - 510 μs versus 738 μs for delay variation

Simulation Case 3

- 7 switch to switch hops
- 100 Mbit/s
- 15 traffic sources
- All sources are time sensitive, CBR traffic with nominal rate of 1333.33 packets/s (nominal time between packets = 0.00075 s)
- Sources have various different frequency offsets that are all within ±100 ppm
- Maximum size packets (1500 bytes plus Ethernet overhead)
  - Switch to switch link utilization ≈ 50%
  - Packet service time (including Ethernet overhead and inter-frame gap = (1500+38)(8)(10^{-8}) s = 0.12304 ms
- Network topology shown two slides following
  - 3 sources at first switch (nodes 1 – 3)
  - Traffic from 1 of these sources (node 1) goes to final switch (node18)
  - Traffic from 2^{nd} and 3^{rd} source (nodes 2 and 3) is dropped at 2^{nd} switch (nodes 5 and 7)
  - At switches 2 – 7 (nodes 20 – 25 in figure), traffic added from 2 CBR sources, carried 1 hop, and dropped
- Simulate for 605 s, with traffic turned on at 5 s

Simulation Case 3

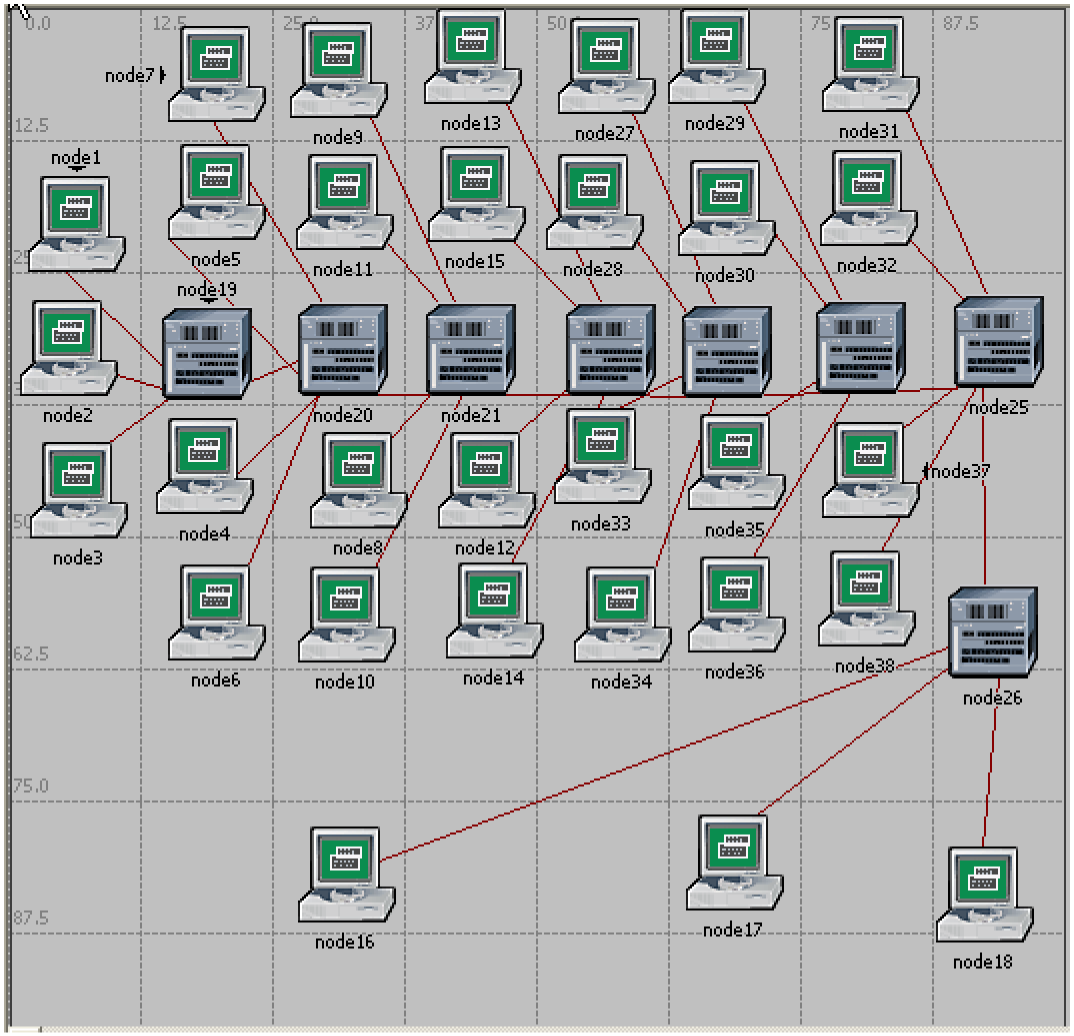

Switch to Switch link utilization = 50%

Case 3 - Results for Through Traffic Stream

Node 1 – 18

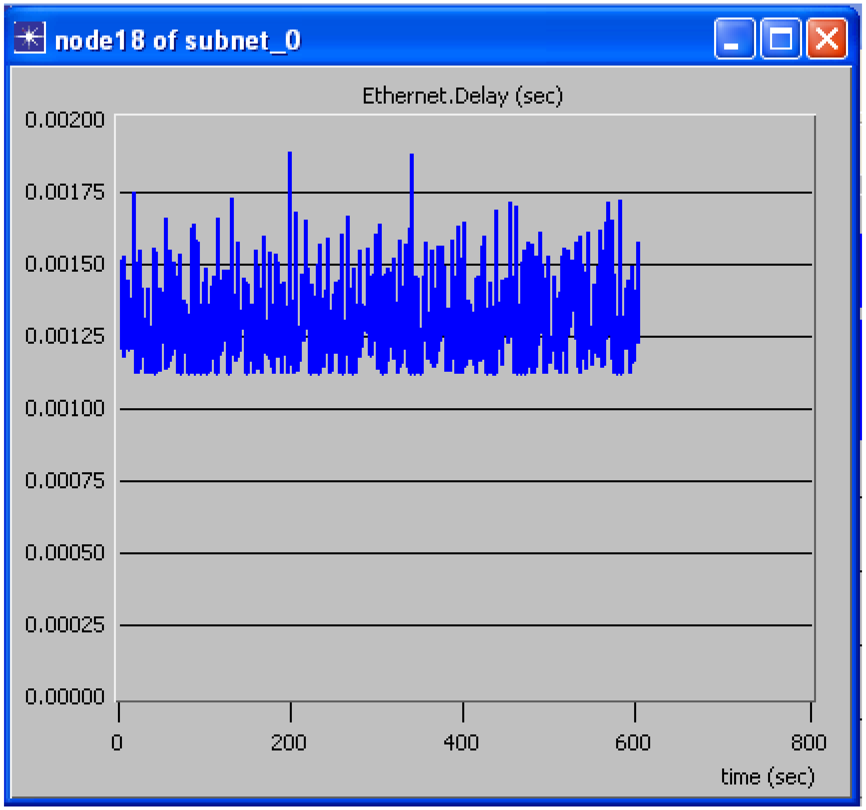

Case 3 - Results for 1-Hop Streams

Node 37 – 16

Node 38 – 17

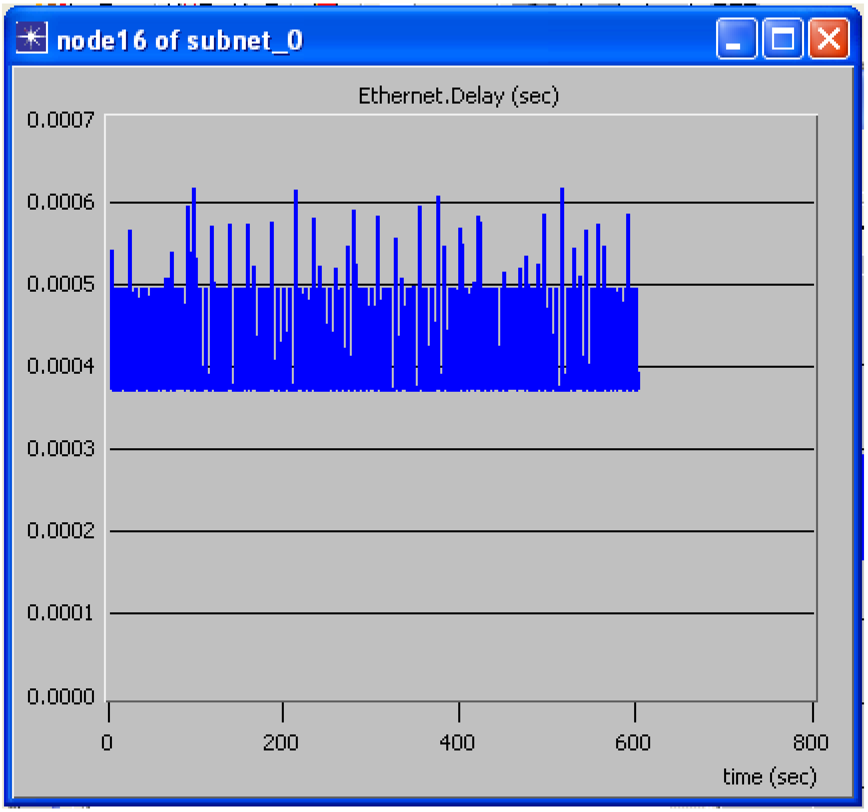

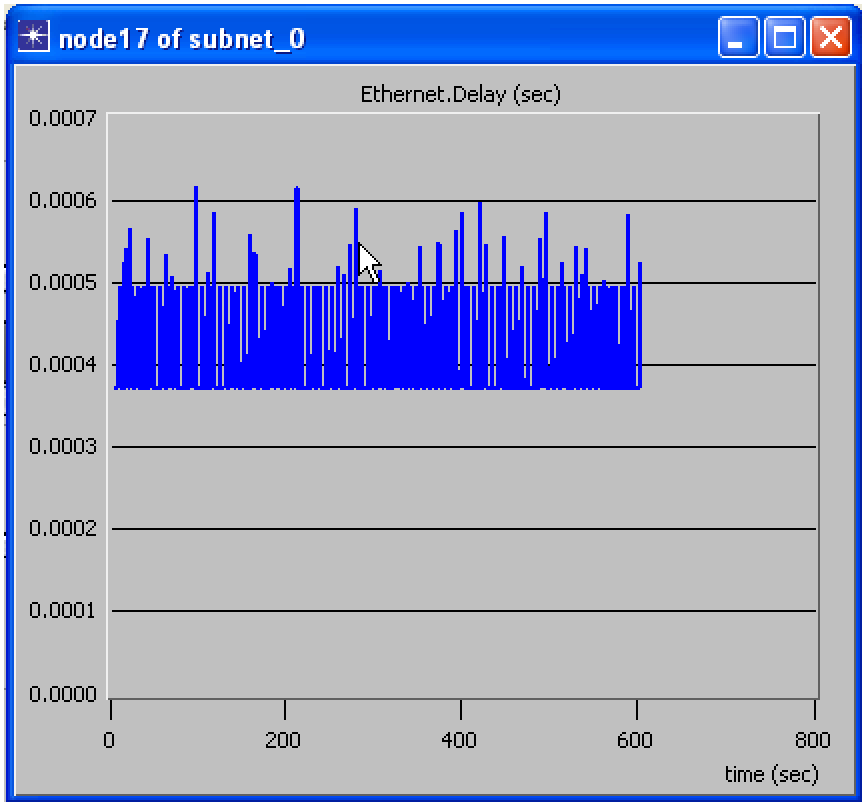

Case 3 Results

- Minimum delay for through streams = 1.1 ms
  - Approximately 9(0.12304) ms as expected (7 interswitch hops plus 2 access hops)
- Maximum delay over 600 s for through streams is approximately 1.89 ms
  - This is less than would be predicted by rule of thumb
  - Total of 14 contending streams
    - Rule of thumb would give (9+14)(0.12304 ms) = 2.8 ms
  - Note that this does exceed maximum delay for Case 1 (and for 2400 s there), which had two contending streams per switch instead of three
- Minimum delay for 1-hop streams = 0.37 ms (consistent with 3 hops)
- Maximum delay for Node 12-15 1-hop stream = 0.62 ms (consistent with 2 contending streams)
- Appears that must simulate for much longer to see full delay predicted by rule of thumb

Simulation Case 4

- 7 switch to switch hops
- 100 Mbit/s
- 9 time sensitive, CBR traffic sources, with nominal rate of 1333.33 packets/s (nominal time between packets = 0.00075 s)
- 7 best effort (Poisson) traffic sources with average rate of 2083.33 packets/s (mean time between packets = 0.00048 s)
- CBR sources have various different frequency offsets that are all within ±100 ppm
- Maximum size packets (1500 bytes plus Ethernet overhead)
  - Switch to switch link utilization ≈ 75%
    - 50% for time-sensitive; 25% for best-effort
  - Packet service time (including Ethernet overhead and inter-frame gap = $(1500+38)(8)(10^{-8})$ s = 0.12304 ms

Simulation Case 4 (Cont.)

- Network topology shown two slides following
  - 3 time-sensitive (nodes 1-3) and 1 best effort source (node 27) at first switch (node 34)
  - Traffic from 2 of these time-sensitive sources (nodes 1 and 2) go to final switch (nodes16 and 18)
  - Traffic from 3rd time-sensitive source (node 3) is dropped at 2nd switch (node 5)
  - Traffic from best effort source is dropped at 2nd switch (node 35)
  - At switches 2 – 7 (nodes 20 – 25 in figure), traffic added from 1 time-sensitive and 1 best effort source, carried 1 hop, and dropped
- Simulate for 305 s, with traffic turned on at 5 s

Simulation Case 4

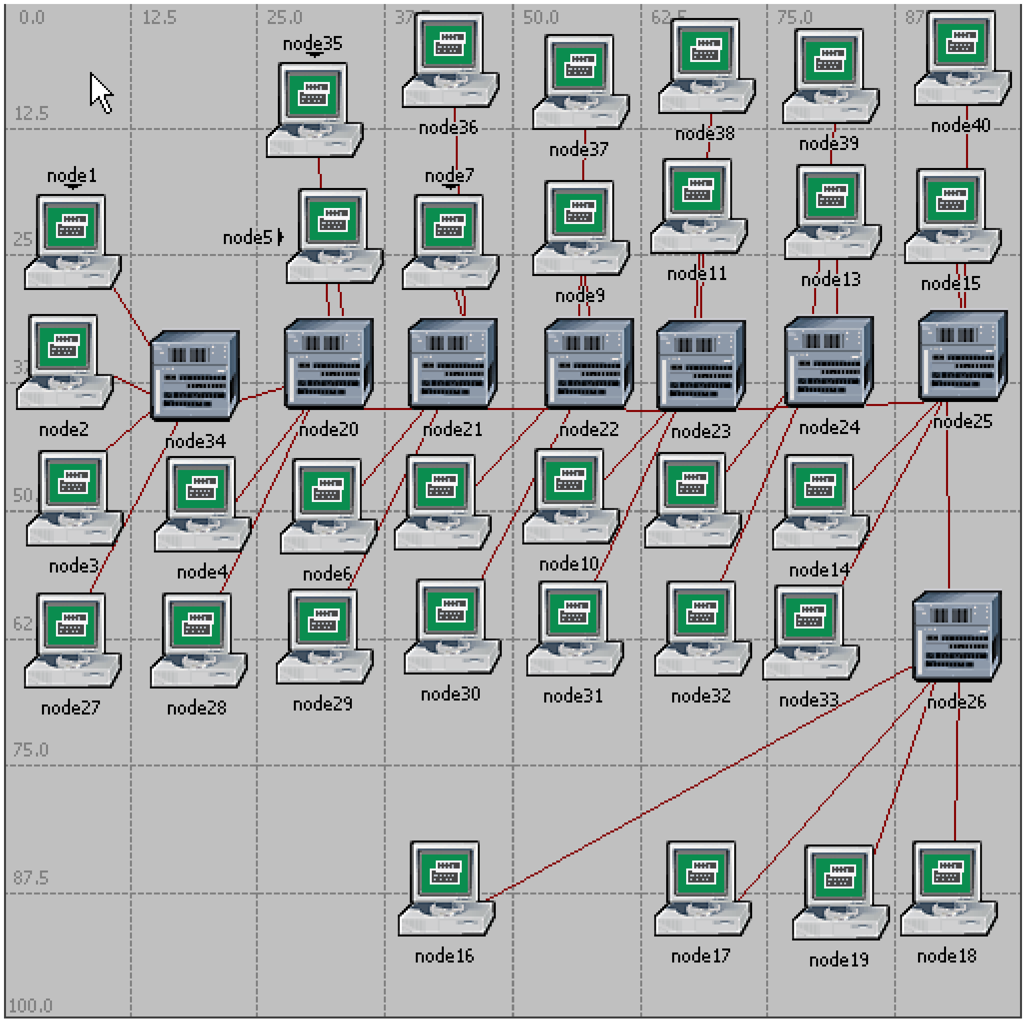

Switch to Switch link utilization = 75%

Case 4 - Results for Through CBR Traffic Streams

Node 1 – 18

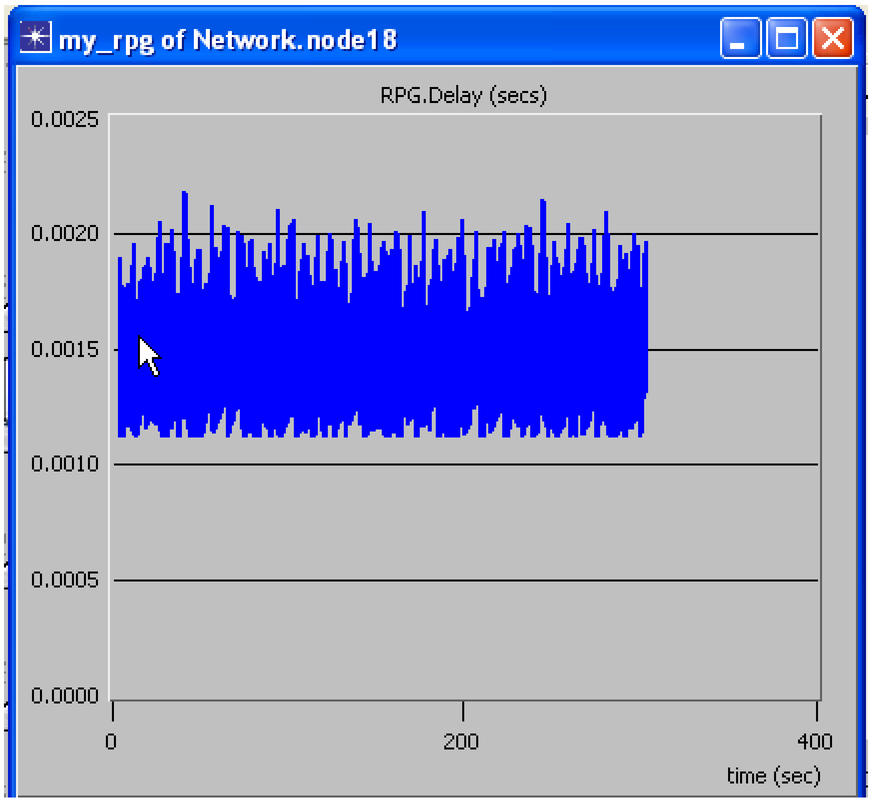

Node 2 – 16

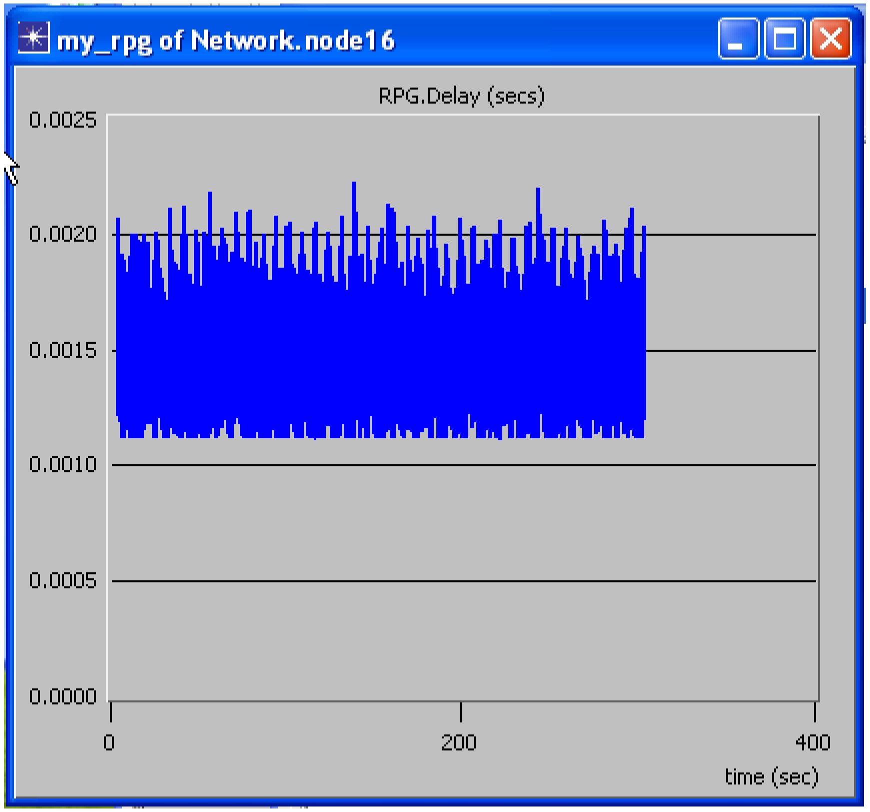

Case 4 - Results for 1-Hop CBR Traffic Streams

Node 3 – 5

Node 14 – 17

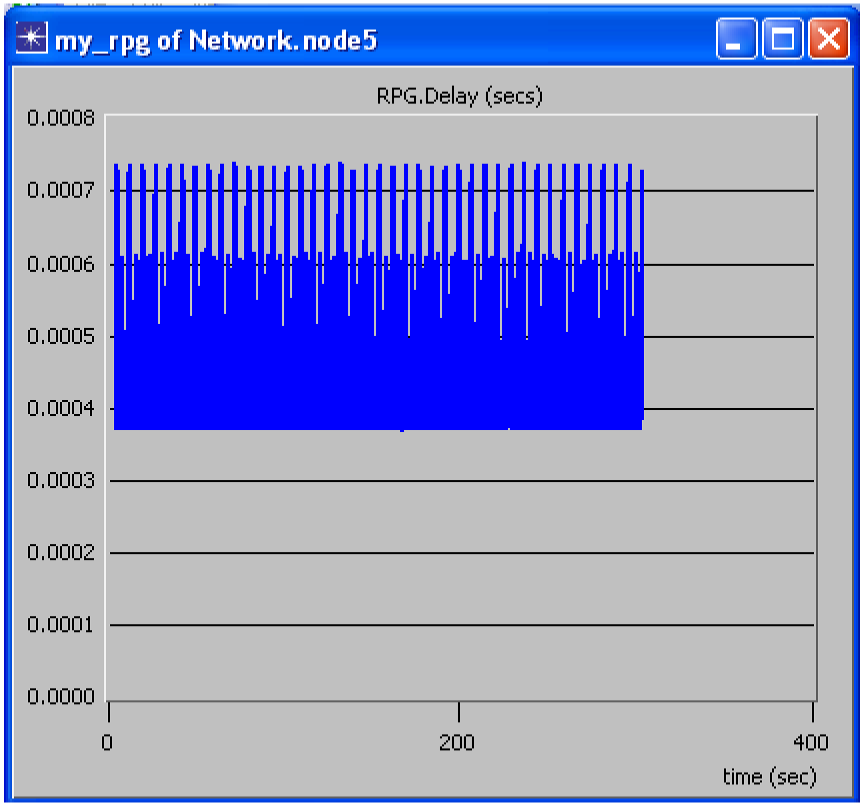

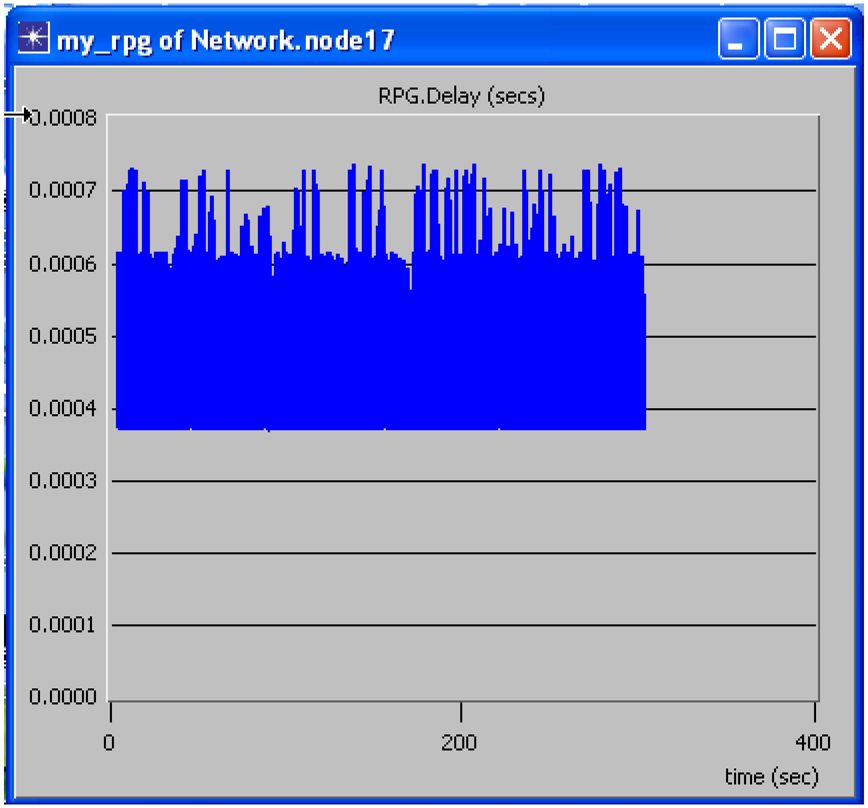

Case 4 - Results for 1-Hop CBR Traffic Streams (Cont.)

Node 23 – 15

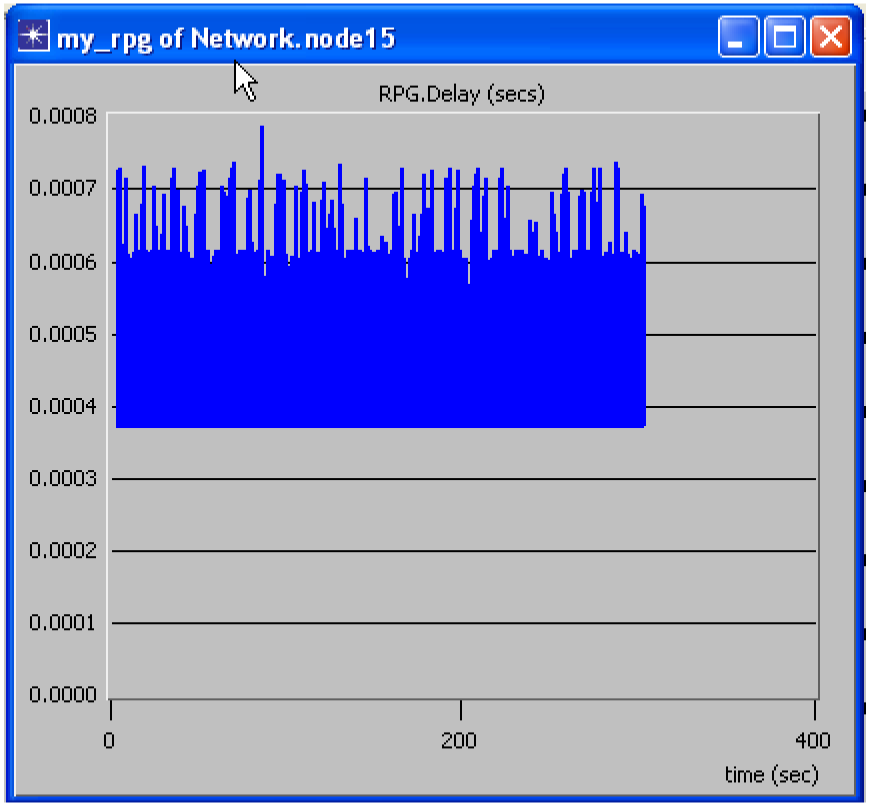

Case 4 - Results for 1-Hop VBR Traffic Streams

Node 32 – 40

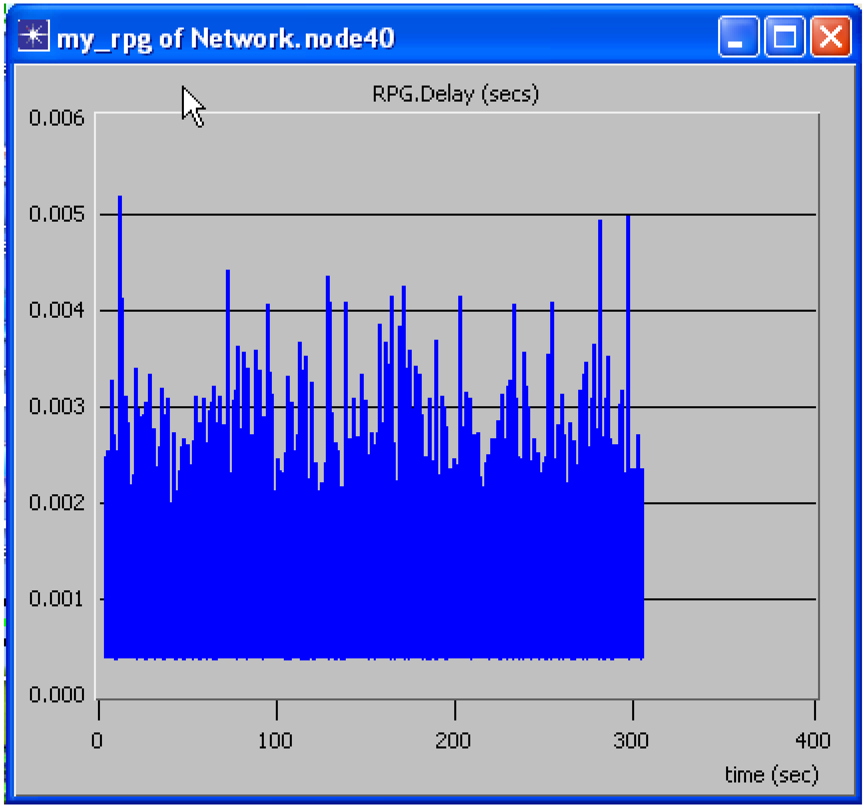

Case 4 Results

- Minimum delay for through CBR streams = 1.1 ms
  - Approximately 9(0.12304) ms as expected (7 interswitch hops plus 2 access hops)
- Maximum delay over 300 s for through streams is approximately 2.2 ms
  - This is less than would be predicted by rule of thumb
  - Total of 15 contending streams
    - Rule of thumb would give (9+15)(0.12304 ms) = 2.95 ms
  - Note that this does exceed the 2 ms limit for end-to-end delay
  - Maximum delay for Node 23-15 1-hop stream = 0.62 ms (consistent with 2 contending streams)
- Appears that must simulate for much longer to see full delay predicted by rule of thumb

Simulation Case 5

- Same as Case 2, except now all the sources are synchronized in time and frequency
  - 7 switch to switch hops, 100 Mbit/s, 9 traffic sources
    - Network topology shown on slide 35; repeated on next slide for convenience
  - Maximum size packets (1500 bytes plus Ethernet overhead)
    - Packet service time (including Ethernet overhead and inter-frame gap = $(1500+38)(8)(10^{-8})$ s = 0.12304 ms
- All sources are time sensitive, CBR traffic with rate of 1333.33 packets/s (time between packets = 0.00075 s for all sources)
  - Switch to switch link utilization ≈ 50%
- Sources at nodes 2 and 3 turned on at 5 s
- Source at node 1 turned on at 5.000001 s

Simulation Case 5 (Cont.)

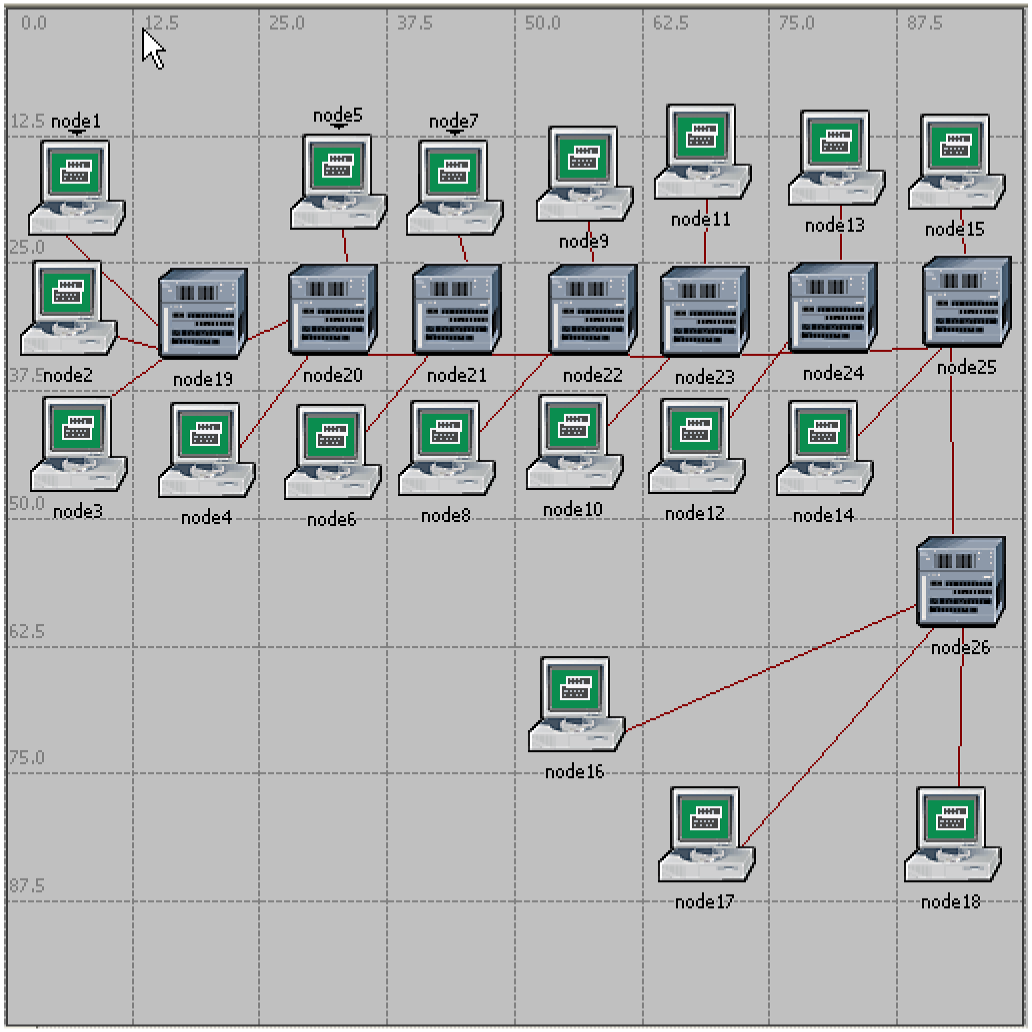

Switch to Switch link utilization = 50%

Simulation Case 5 (Cont.)

- Sources at nodes 4, 6, 8, 10 12, and 14 turned on at times such that packets arrive just before node 1 packets (i.e., to produce maximum contention)
  - Node 4 source turned on at 5.000369 s
  - Node 6 source turned on at 5.000613 s
  - Node 8 source turned on at 5.000860 s
  - Node 10 source turned on at 5.001100 s
  - Node 12 source turned on at 5.001340 s
  - Node 14 source turned on at 5.001582 s
- Simulate for 10 s, with traffic turned on at 5 s
  - As before, needed to add small amount of best-effort traffic in reverse direction to ensure each destination node would be in the forwarding database of each switch (otherwise get flooding and link utilizations that exceed 100%)

Case 5 - Results for Through Traffic Streams

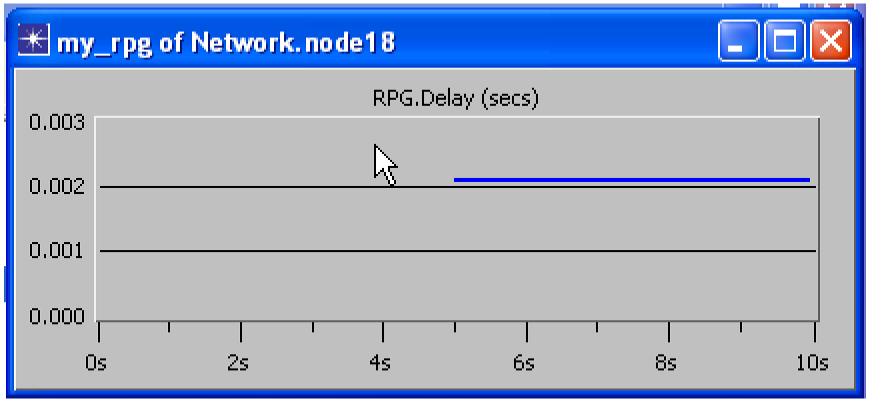

Node 1 – 18

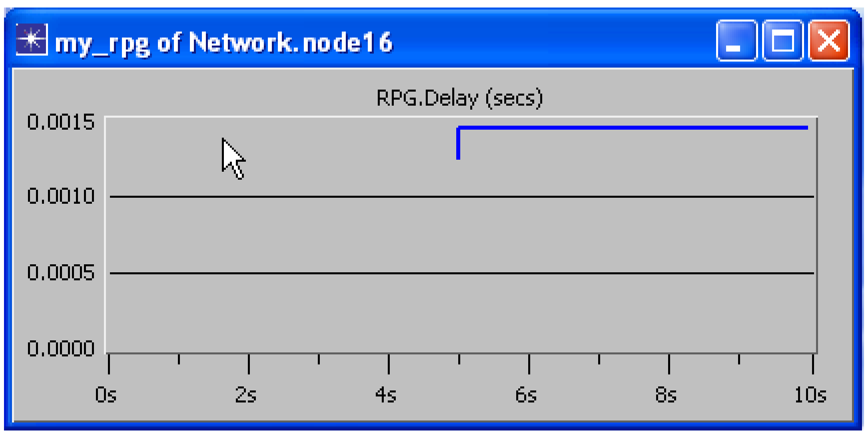

Node 2 – 16

Case 5 - Results for 1-Hop Traffic Stream

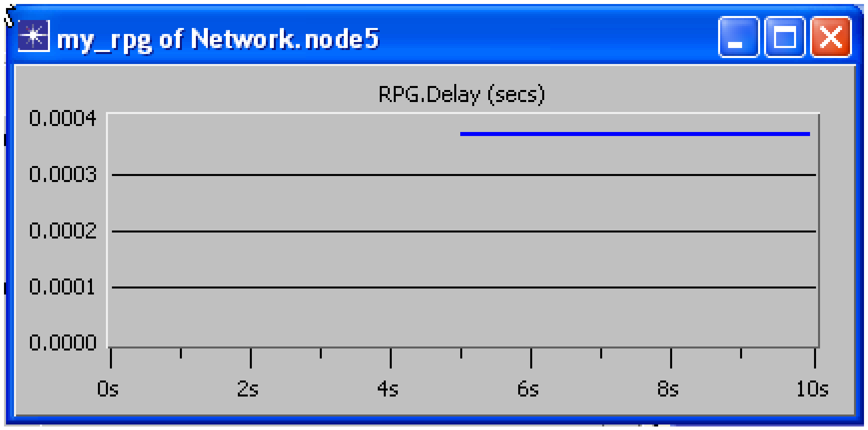

Node 3 – 5

Case 5 – Results

- End-to-end delay for node 1 – 18 traffic stream is 0.002075 s
  - This is 0.002075 s/0.00012304 s = 16.9 ≈ 17 transmission delays
  - Number on links (including access links) = 9
  - Number of contending traffic streams = 8
  - Therefore, result is consistent with rule of thumb
    - Expected this result, because the frequencies and phases of the arrivals were chosen to achieve it
- End-to-end delay for node 2 – 16 is 0.0014488 s
  - This is 11.8 ≈ 12 transmission delays
  - Result is less than rule of thumb predicts, because don't have maximal contention for this stream
- End-to-end delay for node 3 – 5 is 0.0003704 s
  - This is 3 transmission delays
  - Result is less than rule of thumb predicts (i.e., 5 transmission delays), because don't have maximal contention for this stream
- Peak-to-peak delay variation for all streams is zero
  - The delays of the successive packets on a stream are the same because all the streams are synchronized

Conclusions

- Rules of thumb give worst-case delay and delay variation
  - Apply to cases with only time-sensitive (CBR) traffic, and cases with both CBR and best-effort traffic
- For link utilizations of 50 – 75%, worst-case delay and delay variation are reached after reasonable simulation time (e.g., up to a few hundred seconds) for a relatively small number of contending traffic streams (e.g., 5 or fewer)
  - For a larger number of contending traffic streams, it apparently takes a much longer simulation time to reach the theoretical maximum delay and delay variation predicted by the rules of thumb
- Can get worst-case delay predicted by rule of thumb for streams that are synchronized
  - Streams must be synchronized in frequency and have the proper phase relationships