

Delay Variation Simulation Results for Transport of Time-Sensitive Traffic over Conventional Ethernet

Geoffrey M. Garner
gmgarner@comcast.net
Felix Feng
Feng.fei@samsung.com

SAMSUNG Electronics

IEEE 802.3 ResE SG
2005.07.18

Outline

- ❑ Introduction
- ❑ Simulation models and assumptions
- ❑ Case 1
 - Scenarios
 - Results
- ❑ Case 2
 - Scenarios
 - Results
- ❑ Conclusions

Introduction

- ❑ One main goal of Residential Ethernet (ResE) is the carrying of time-sensitive traffic with acceptable jitter, and wander performance
 - Jitter and Wander requirements for uncompressed and compressed (MPEG-2) digital video applications and for digital audio applications are summarized in [1]
- ❑ It is of interest to determine what performance can be expected if multiple time-sensitive traffic streams are transported using current Ethernet with priorities
 - Time sensitive traffic would get high priority
 - Best-effort traffic would get low priority
 - Timing for a time-sensitive traffic stream would be recovered at the network egress via filtering (e.g., using Phase-Locked-Loop (PLL))
- ❑ This presentation contains simulation results for several scenarios of time-sensitive traffic transport over current Ethernet using priorities

Introduction (Cont.)

□ Note on buffering

- If PLL filtering is used to recover application timing, must still buffer an amount of data on the order of the unfiltered phase peak-to-peak phase variation
 - Analogous to the PLL phase detector error
- For the time-sensitive traffic stream cases considered here, this is on the order of tens of microseconds to 100 – 200 microseconds
- In previous discussions, an alternative using a free-running clock at the egress to create the recovered application timing has been suggested, i.e., instead of PLL filtering
 - In this approach, must buffer enough data to prevent buffer underflow or overflow for the duration of the audio or video application
 - For ± 100 ppm clocks and video or audio applications on the order of hours, this would imply buffering some number of seconds worth of data (e.g., a minimum of 2.2 s for a 3 hour application)
 - This amount of delay would be added to the application end-to-end delay
 - The delay would be present at startup, e.g., when changing channels for video

Simulation Models and Assumptions

❑ OPNET simulation tool was used to simulate packet delays

- OPNET contains models for full-duplex Ethernet MAC and for Ethernet bridges
- Models were modified to include priority classes
 - Priority queueing is non-preemptive
- OPNET produces delays of successive packets, for each traffic stream
 - For a constant-rate packet ingress, the packet delay history differs from the unfiltered phase error history at the egress node by a constant
 - The constant is equal to the average network delay (see next slide)

❑ OPNET packet delays were input to a stand-alone C program (run under Cygwin) that implements a 2nd order, linear filter with 20 dB/decade roll-off

- Model details are described in Subclause VIII.2.2 of [2], and also in Section 2.1.2 of [3]
 - Exact integrating factor for the filter is obtained from state equations
- 1 Hz bandwidth
- 0.1 dB gain peaking

❑ Since the packet delay history and unfiltered phase error history at the egress differ by a constant, the low-pass filtering of each will produce the same steady-state peak-to-peak phase variation

Simulation Models and Assumptions (Cont.)

□ Relation between packet delay history and phase error history at egress, for a constant rate packet ingress

- Let $t_{1,k}$ be the time the k^{th} packet enters the network
- Let $t_{2,k}$ be the time the k^{th} packet leaves the network
- Let d_k be the delay for the k^{th} packet
 - Then $d_k = t_{2,k} - t_{1,k} = d_{av} + v_k$
 - d_{av} = average delay
 - v_k = delay variation
- For a constant rate stream, $t_{1,k} = kT$, where T is the time between packets at the ingress
- If the network did not impose any delay variation, then the delay for all the packets would be d_{av} and we would have $t_{2,k, \text{no delay variation}} = kT + d_{av}$
- The unfiltered phase error x_k at the egress is the difference between the time the packet arrives and the time it would have arrived had there been no delay variation. Then
 - $x_k = t_{2,k} - (kT + d_{av}) = t_{2,k} - t_{1,k} - d_{av} = v_k$
- Therefore, the unfiltered phase error at the egress is equal to the variable portion of the delay
 - The unfiltered phase error at the egress and the delay differ by the average delay (a constant)

Simulation Models and Assumptions (Cont.)

□ OPNET model assumptions

- Considered two types of traffic mixes:
 - Time sensitive traffic only
 - Both time-sensitive traffic and best-effort traffic
- Ethernet links are 100 Mbit/s (FE)
- 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 model contains spanning tree and rapid spanning tree algorithms
 - Same result is obtained with either for simple network cases considered here

Simulation Models and Assumptions (Cont.)

□ Time-sensitive traffic assumptions

- Packet size is a constant, equal to 256 bytes (2048 bits)
- Time between packets at source is a constant (chosen for each case to achieve desired link utilization)
- All the Time-sensitive streams have the same nominal rate, but differ slightly (within a frequency tolerance)
 - 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

□ Best-effort traffic assumptions

- Packet size is a constant, equal to 1538 bytes (12304 bits)
- Time between packets is exponentially distributed (i.e., Poisson packet arrivals) with mean inter-arrival time chosen to achieve desired link utilization

□ Two main simulation cases were run, each with several scenarios (sub-cases)

- Described on following slides

Simulation Case 1

☐ All traffic is time-sensitive

- Packet size is as given above
- Nominal packet arrival rate for each stream is 8000 packets/s
 - Nominal time between packets is 0.000125 s

☐ 3 talker nodes connected to one Ethernet switch

☐ 3 listener nodes connected to a second Ethernet switch

☐ The two Ethernet switches are connected together

- See figure on next slide

☐ Stream 1

- talker_1 to listener_1

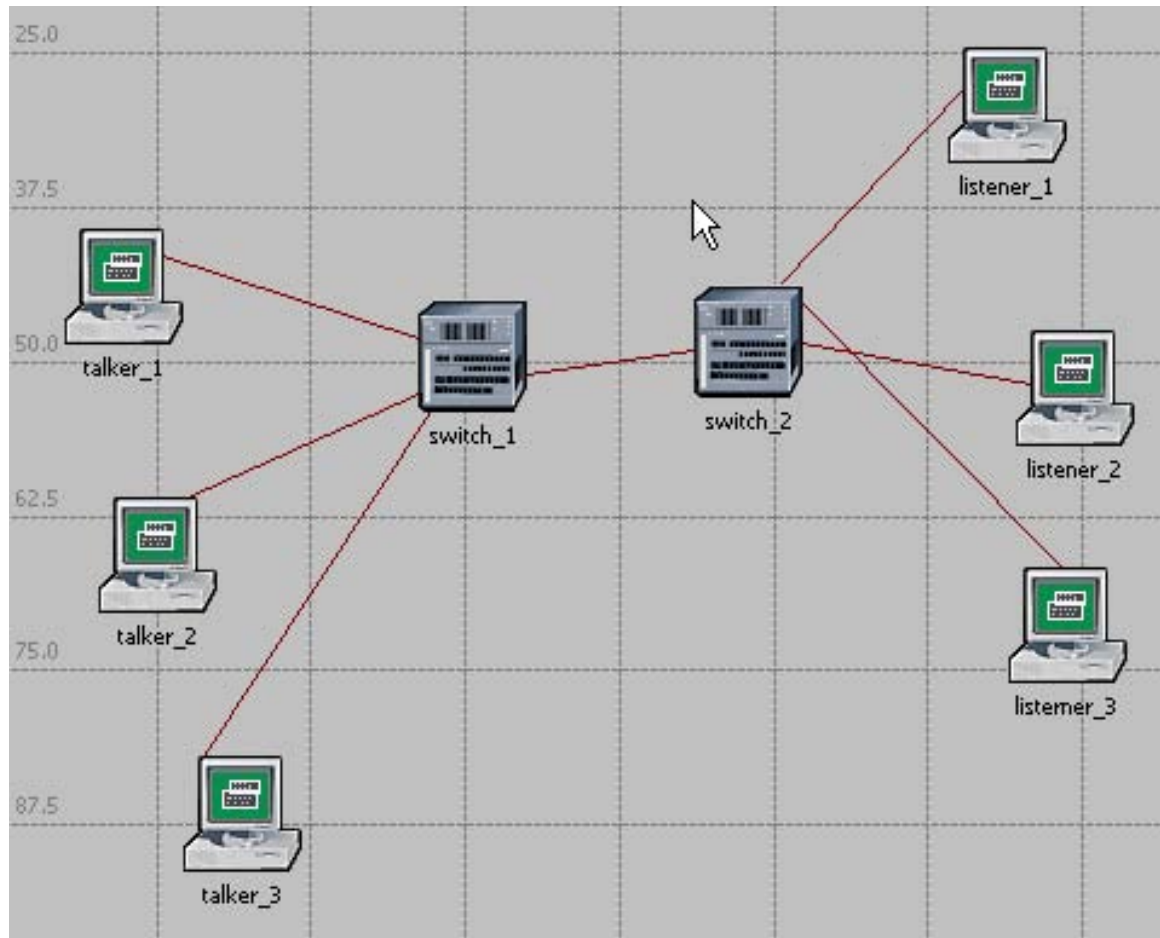
☐ Stream 2

- talker_2 to listener_2

☐ Stream 3

- talker_3 to listener_3

Simulation Case 1 (Cont.)



Switch_1 to Switch_2 link
utilization = 54%

Simulation Case 1, Scenario 1

❑ Stream 1

- talker_1 to listener_1, nominal rate

❑ Stream 2

- talker_2 to listener_2, rate offset by -100 ppm

❑ Stream 3


- talker_3 to listener_3, rate offset by +100 ppm

❑ Simulated for 105 s, with traffic turned on at 5 s


- Plots measure time from when traffic is turned on

Simulation Case 1, Scenario 1 Results

Case 1, Scenario 1, talker_1 -> listener_1
Unfiltered delay variation




Case 1, Scenario 1, talker_1 -> listener_1
Filtered delay variation




Simulation Case 1, Scenario 1 Results (Cont.)

Case 1, Scenario 1, talker_1->listener_1
Unfiltered delay variation
Detail of first 20 s




Simulation Case 1, Scenario 1 Results (Cont.)

Case 1, Scenario 1, talker_2->listener_2
Unfiltered delay variation




Case 1, Scenario 1, talker_2->listener_2
Filtered delay variation



Simulation Case 1, Scenario 1 Results (Cont.)

Case 1, Scenario 1 stream MTIE curves
and comparison with video and audio masks



Simulation Case 1, Scenario 1 Results (Cont.)

- ❑ As expected, delay variation occurs as faster streams overtake slower streams (i.e., streams beat against each other)
- ❑ Peak-to-peak unfiltered phase variation is approximately 47 μs (2 packets)
- ❑ Filtering reduces this to approximately 20 μs
- ❑ Unfiltered phase variation plots show evidence of additional lower frequency envelope
- ❑ MTIE is within MPEG-2 mask for case of transport to residence via a service provider (but not clear what budget component ResE gets)
- ❑ MTIE exceeds masks for digital audio and uncompressed digital video

Simulation Case 1, Scenario 2

❑ Stream 1

- talker_1 to listener_1, nominal rate

❑ Stream 2

- talker_2 to listener_2, rate offset by -1 ppm

❑ Stream 3


- talker_3 to listener_3, rate offset by +1 ppm

❑ Simulated for 405 s, with traffic turned on at 5 s


- Plots measure time from when traffic is turned on

Simulation Case 1, Scenario 2 Results

Case 1, Scenario 2, talker_1 -> listener_1
Unfiltered delay variation




Case 1, Scenario 2, talker_1 -> listener_1
Filtered delay variation




Simulation Case 1, Scenario 2 Results (Cont.)

Case 1, Scenario 2, talker_2->listener_2
Unfiltered delay variation




Case 1, Scenario 2, talker_2->listener_2
Filtered delay variation




Simulation Case 1, Scenario 2 Results (Cont.)

Case 1, Scenario 2, talker_3->listener_3
Unfiltered delay variation



Case 1, Scenario 2, talker_3->listener_3
Filtered delay variation




Simulation Case 1, Scenario 2 Results (Cont.)

- ❑ Small peaks for streams 2 and 3 are due to these streams beating against each other (their relative frequency offset is 2 ppm, or twice their offset relative to stream 1)
- ❑ Large peaks for streams 2 and 3 occur when stream 3 overtakes streams 1 and 2 at the same time
 - Stream 1 result indicates that stream 1 overtakes stream 2 and then is immediately overtaken by stream 3
- ❑ Unfiltered phase variation is now of lower frequency compared to Scenario 1, due to smaller frequency offsets

Simulation Case 1, Scenario 2 Results (Cont.)

Case 1, Scenario 2 stream MTIE curves
and comparison with video and audio masks



Simulation Case 1, Scenario 2 Results (Cont.)

- ❑ 1 Hz filter has little impact, as period of phase variation is considerably longer than filter time constant ($1/2\pi \text{ s} = 0.159 \text{ s}$)
- ❑ MTIE reaches MPEG-2 mask for case of transport to residence via a service provider (but not clear what budget component ResE gets)
- ❑ MTIE exceeds masks for digital audio and uncompressed digital video

Simulation Case 2, Scenario 1

- ❑ Similar to Case 1, Scenario 1, except now have 6 traffic streams instead of 3, each with half the traffic volume as in Case 1
- ❑ 6 talker nodes connected to one Ethernet switch
- ❑ 6 listener nodes connected to a second Ethernet switch
- ❑ The two Ethernet switches are connected together
- ❑ Simulated for 105 s, with traffic turned on at 5 s
 - Unfiltered phase plots measure time from $t = 0$
 - Filtered phase plots measure time from when traffic is turned on

Simulation Case 2, Scenario 1 (Cont.)

❑ Stream 1

- talker_1 to listener_1, nominal rate

❑ Stream 2

- talker_2 to listener_2, rate offset by -100 ppm

❑ Stream 3

- talker_3 to listener_3, rate offset by +100 ppm

❑ Stream 4

- talker_4 to listener_4, nominal rate

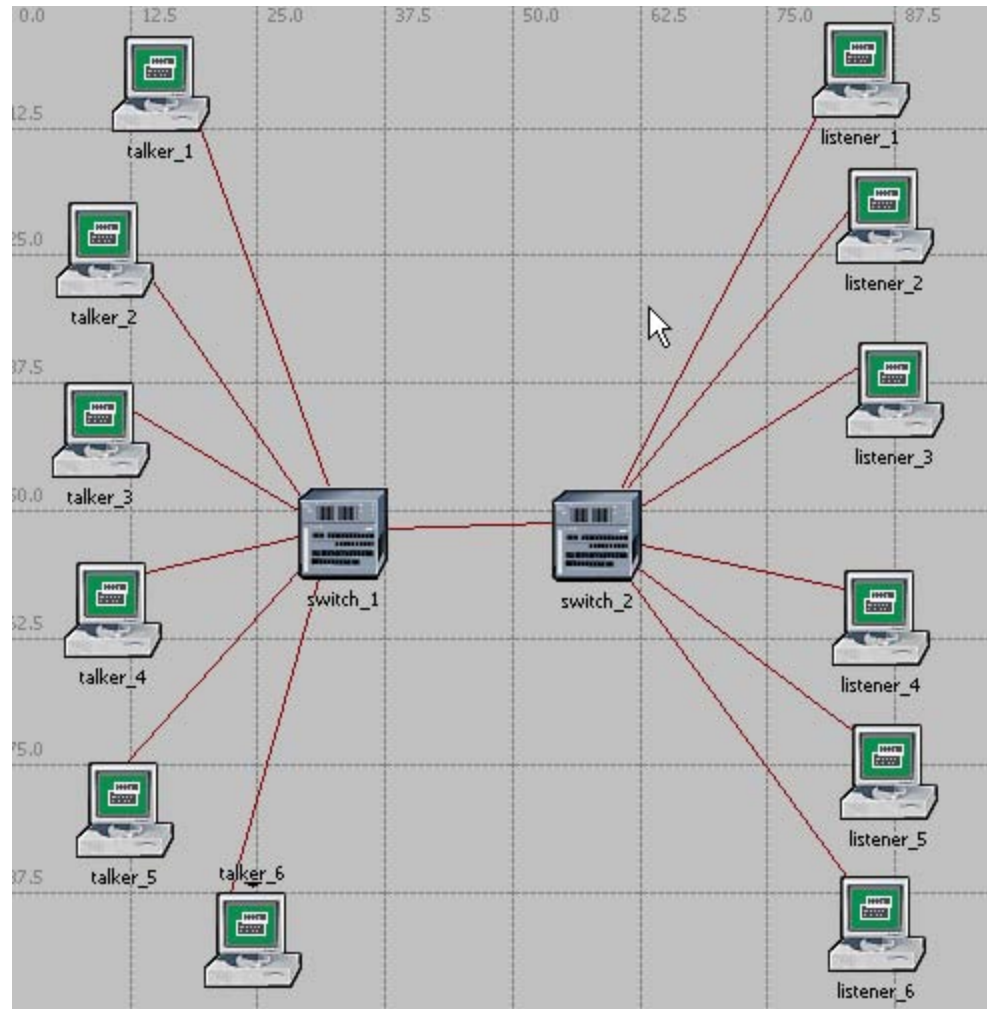
❑ Stream 5

- talker_5 to listener_5, rate offset by -50 ppm

❑ Stream 6

- talker_6 to listener_6, rate offset by +50 ppm


Simulation Case 2, Scenario 1 (Cont.)




Switch_1 to Switch_2 link
utilization = 54%

Simulation Case 2, Scenario 1 Results

Case 2, Scenario 1, talker_1 -> listener_1
Unfiltered delay variation




Case 2, Scenario 1, talker_1 -> listener_1
Filtered delay variation




Simulation Case 2, Scenario 1 Results (Cont.)

Case 2, Scenario 1, talker_3->listener_3
Unfiltered delay variation

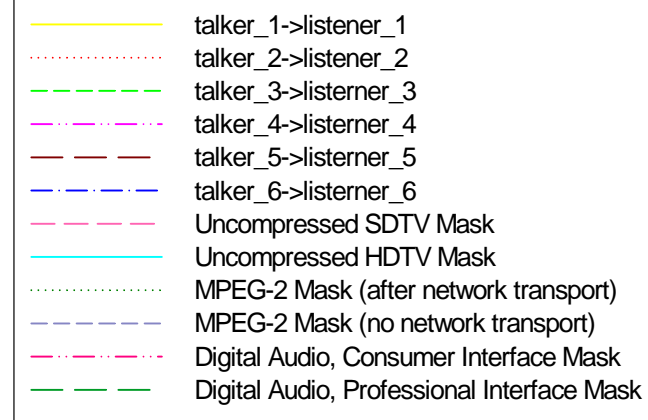



Case 2, Scenario 1, talker_3->listener_3
Filtered delay variation



Simulation Case 2, Scenario 1 Results (Cont.)

Case 2, Scenario 1 stream MTIE curves
and comparison with video and audio masks



Simulation Case 2, Scenario 1 Results (Cont.)

- ❑ Phase variation patterns are more complicated compared to Case 1, due to larger number of streams beating against each other
- ❑ MTIE for filtered phase is larger than for Case 1, Scenario 1, and now exceeds MPEG-2 mask for case of transport to residence via a service provider
- ❑ MTIE exceeds masks for digital audio and uncompressed digital video

Simulation Case 2, Scenario 2

- ❑ Similar to Case 2, Scenario 1, except now have added a best effort traffic stream as the 7th stream
- ❑ 7 talker nodes connected to one Ethernet switch
- ❑ 7 listener nodes connected to a second Ethernet switch
- ❑ The two Ethernet switches are connected together
- ❑ Simulated for 105 s, with traffic turned on at 5 s
 - Unfiltered phase plots measure time from $t = 0$
 - Filtered phase plots measure time from when traffic is turned on
 - For this scenario, only selected traffic stream results are shown as the results for all the time-sensitive traffic streams were similar

❑Stream 1

- talker_1 to listener_1, nominal rate

❑Stream 2

- talker_2 to listener_2, rate offset by -100 ppm

❑Stream 3

- talker_3 to listener_3, rate offset by +100 ppm

❑Stream 4

- talker_4 to listener_4, nominal rate

❑Stream 5

- talker_5 to listener_5, rate offset by -50 ppm

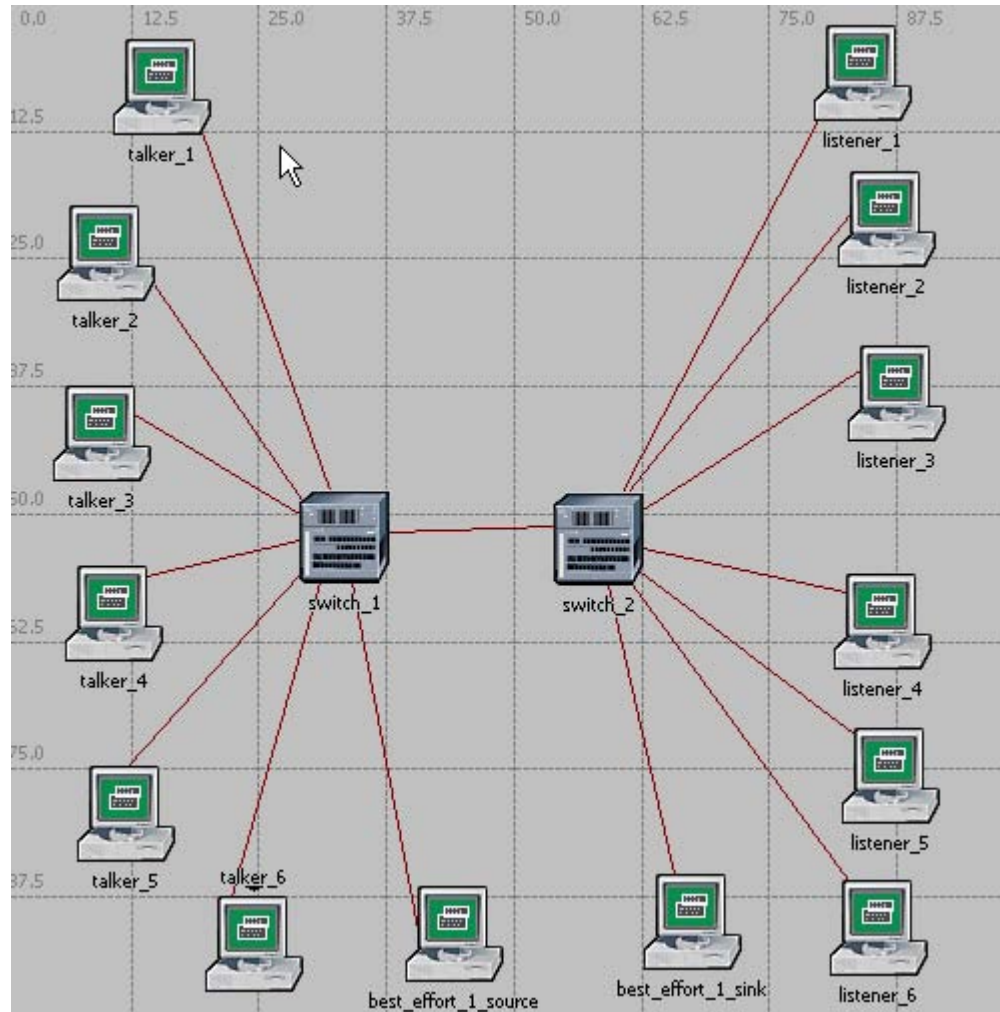
❑Stream 6

- talker_6 to listener_6, rate offset by +50 ppm

❑Stream 7

- Best_effort_1_source to best_effort_1 sink
- Packet size = 1538 bytes (12304 bits)
- Poisson packet arrivals, mean inter-arrival time = 0.49216 ms (chosen to make total switch_1 to switch_2 link utilization 80%)


Simulation Case 2, Scenario 2 (Cont.)




Switch_1 to Switch_2 link
utilization = 80%

Simulation Case 2, Scenario 2 Results (Cont.)

Case 2, Scenario 2, talker_2->listener_2
Unfiltered delay variation




Case 2, Scenario 2a, talker_2->listener_2
Filtered delay variation



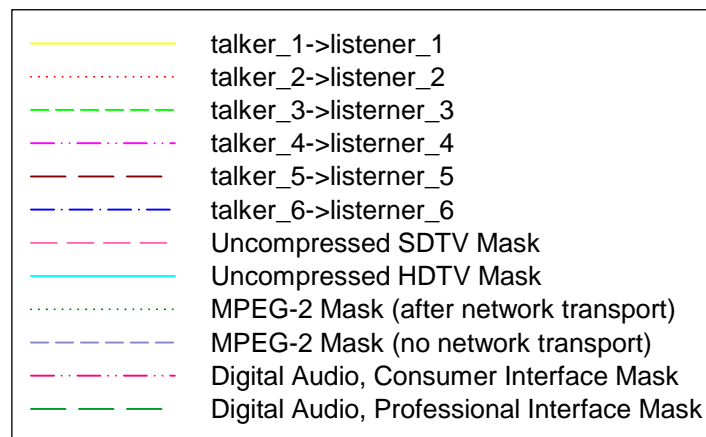

Simulation Case 2, Scenario 2 Results (Cont.)

Case 2, Scenario 2, best_effort_1_source->best_effort_1_sink
Unfiltered delay variation



Simulation Case 2, Scenario 2 Results (Cont.)

Case 2, Scenario 1 stream MTIE curves
and comparison with video and audio masks



Simulation Case 2, Scenario 2 Results (Cont.)

- ❑ Phase variation patterns for unfiltered phase are much less regular compared to cases with only time-sensitive traffic
 - Regular patterns are destroyed by the random best-effort traffic
- ❑ Nonetheless, phase variation patterns for filtered phase bear a greater resemblance to corresponding patterns of Case 2, Scenario 1 (no best-effort traffic)
- ❑ MTIE for filtered phase is similar to that for Case 2, Scenario 1 (though slightly larger for shorter observation intervals)
 - The larger unfiltered phase variation is of frequency greater than 1 Hz and is filtered
- ❑ MTIE exceeds masks for MPEG-2, digital audio, and uncompressed digital video

Conclusions

- ❑ For 50% link utilization, three time-sensitive traffic streams, 256 byte packets, and 100 Mbit/s links, the MTIE masks for digital audio and uncompressed digital video are exceeded
 - The MTIE mask for MPEG-2 is not exceeded for larger frequency offsets and just reached for smaller frequency offsets, but note that ResE gets only a budget component of this mask
 - The small frequency offset case shows evidence of an additional low-frequency envelope, that would result in larger MTIE
- ❑ For 50% link utilization due to six time-sensitive traffic streams, 256 byte packets, and 100 Mbit/s links, the MTIE masks for MPEG-2, digital audio, and uncompressed digital video are exceeded. Adding best-effort traffic to increase the link utilization to 75% does not change MTIE appreciably.
- ❑ MTIE will be larger for larger time-sensitive traffic stream packet size
- ❑ MTIE will be larger for a larger number of time-sensitive traffic streams
- ❑ MTIE will be smaller for larger link speed (e.g., 1 Gb Ethernet)

Conclusions (Cont.)

- ❑ A filter bandwidth that is considerably less than 1 Hz is required to effectively filter phase variation to levels within the digital audio and video MTIE masks for cases where the different time sensitive traffic streams have small frequency offsets relative to each other
 - A filter with such a bandwidth that also had acceptable noise generation would be impractical (i.e., expensive)
- ❑ The results indicate that timing recovery for the time sensitive traffic streams by filtering the streams at the egress (e.g., with a PLL function) will not enable the respective jitter and wander requirements to be met (the requirements are embodied in the MTIE masks)

Conclusions (Cont.)

□ Note that the cases here are not worst-case

- Time-sensitive streams with larger packet sizes would give worse performance (larger delay variation)
- Networks with more time-sensitive streams and/or higher link utilization due to the time-sensitive streams (e.g., 70%) will give worse performance (larger delay variation)

□ Conclusion – Timing recovery using PLL filtering of the time-sensitive data packet arrivals will not provide a good enough clock even in realistic scenarios, let alone worst-case scenarios

- An alternative scheme that transports timing through some other means is needed


References

1. 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.
2. ITU-T Recommendation G.8251, *The Control of Jitter and Wander within the Optical Transport Network (OTN)*, ITU-T, Geneva, November, 2001, Amendment 1, June, 2002, Corrigendum 1, June, 2002.
3. Geoffrey Garner, *Jitter Analysis for Asynchronous Mapping of a Client Signal into an Och*, Lucent Contribution to ITU-T Q 11/15 Interim Meeting, Ottawa, ON, July, 2000.

Additional Results


Simulation Case 1, Scenario 1 Results (Cont.)

Case 1, Scenario 1, talker_2->listener_2
Unfiltered delay variation
Detail of first 10 s




Simulation Case 1, Scenario 1 Results (Cont.)

Case 1, Scenario 1, talker_3->listener_3
Unfiltered delay variation




Case 1, Scenario 1, talker_3->listener_3
Filtered delay variation




Simulation Case 1, Scenario 1 Results (Cont.)

Case 1, Scenario 1, talker_3->listener_3
Unfiltered delay variation
Detail of first 10 s




Simulation Case 2, Scenario 1 Results (Cont.)

Case 2, Scenario 1, talker_2->listener_2
Unfiltered delay variation




Case 2, Scenario 1, talker_2->listener_2
Filtered delay variation




Simulation Case 2, Scenario 1 Results (Cont.)

Case 2, Scenario 1, talker_4 -> listener_4
Unfiltered delay variation




Case 2, Scenario 1, talker_4 -> listener_4
Filtered delay variation




Simulation Case 2, Scenario 1 Results (Cont.)

Case 2, Scenario 1, talker_5->listener_5
Unfiltered delay variation




Case 2, Scenario 1, talker_5->listener_5
Filtered delay variation



Simulation Case 2, Scenario 1 Results (Cont.)

Case 2, Scenario 1, talker_6->listener_6
Unfiltered delay variation



Case 2, Scenario 1, talker_6->listener_6
Filtered delay variation

