Further Simulation Results for AVB Synchronization Transported using IEEE 1588 Peer-to-Peer Transparent Clocks

Geoffrey M. Garner
SAMSUNG Electronics (Consultant)

IEEE 802.1 AVB TG
IEEE 1588
2006.07.12

gmgarner@comcast.net

Outline

- Introduction
- Review of Simulation model
- Simulation cases
- Results
- Future work
- Appendix – Confidence Interval for a Quantile of a Distribution

Introduction

- A previous presentation [5] contained initial simulation results for the transport of synchronization over an AVB network for cases using
  - The IEEE 1588 Peer-to-Peer (P2P) Transparent Clock (TC)
  - The scheme described in the white paper [3]
    - This scheme is a boundary clock approach, in the sense that each clock is synchronized to its upstream neighbor rather than directly to the grandmaster (however, it differs in a number of ways from some schemes that use IEEE 1588 Boundary Clocks)
- The results in [5] showed
  - For a single hop, the scheme using P2P TCs gave somewhat better performance
    - MTIE for this scheme was lower by approximately an order of magnitude for observation intervals ranging from 10 μs to 20 s
  - For more than one hop, the schemes gave similar performance, with the scheme of [3] slightly better for shorter observation intervals and the P2P TC based scheme slightly better for longer observation intervals
- The results in [5] were based on single simulation runs
  - The results depend on the actual frequency offsets of the free-running node clocks, as well as the random noise generation
    - The noise generation model is described in [1] and [2] and references given there

Introduction (Cont.)

- The node clock frequency offsets were initialized randomly
- The random phase noise generation in each node depended on the initial state of the random number generator
- The resulting phase error accumulation time histories at each node are therefore samples of a random process
  - The MTIE results at each node are also samples of distributions of MTIE
  - The fact that a single MTIE sample is within a mask does not guarantee that any specified portion of the distribution of MTIE is within the mask
- It is therefore of interest to determine results for the distribution of MTIE
  - Specifically, it is of interest to determine results for a specified quantile (percentile) of MTIE
- The current presentation obtains simulation results for a 99% confidence interval for the 0.95 quantile of MTIE
  - For each observation interval, a range is given, such that we are 99% confident that the the 0.95 quantile of MTIE is within the range
    - I.e., the probability that the 0.95 quantile of MTIE is not in the range is 1%
  - Recall that the 0.95 quantile of MTIE, for a given observation interval, is a value that is exceeded by 5% of the MTIE samples (peak-to-peak values) of the process for that observation interval, measured over all samples of the process

Introduction (Cont.)

- Confidence intervals for the 0.95 quantile of MTIE are obtained by running multiple, independent replications of each simulation
  - Specifically, each simulation case is run 300 time
  - At the conclusion of each run, the state of the random number generator is saved and used to initialize the random number generator for the next run
  - The number of times the random number generator is invoked is counted, to verify that the supply of pseudo-random values is not exceeded
  - It is shown in the Appendix that, given 300 independent random samples of a population, a 99% confidence interval for the 0.95 quantile may be obtained by placing the samples in ascending order
    - A 99% confidence interval extends from the 275th to 294th smallest samples
    - The midpoint of the range (the 285th smallest sample) is taken as a point estimate
- The same simulation cases with network assumptions considered in [5] are considered here
  - I.e., 10 ms synch interval, 100 ms frequency update interval, other basic parameter values as in [5]
  - 3 dB filter bandwidths were chosen to be 1 Hz, 0.1 Hz, 0.01 Hz (3 cases)
- For each case, results are presented for both synchronization transport using IEEE 1588 P2P TCs and using the scheme of [3] (resulting in 2 sets of 300 runs each for each of the 3 cases)

Review of Simulation Model (from [5])

- The figure below (taken from [4]) shows synchronization transported from a Grandmaster (GM), through a chain of P2P TCs, to a slave clock (this figure shows the case of a GM and 3 additional nodes)

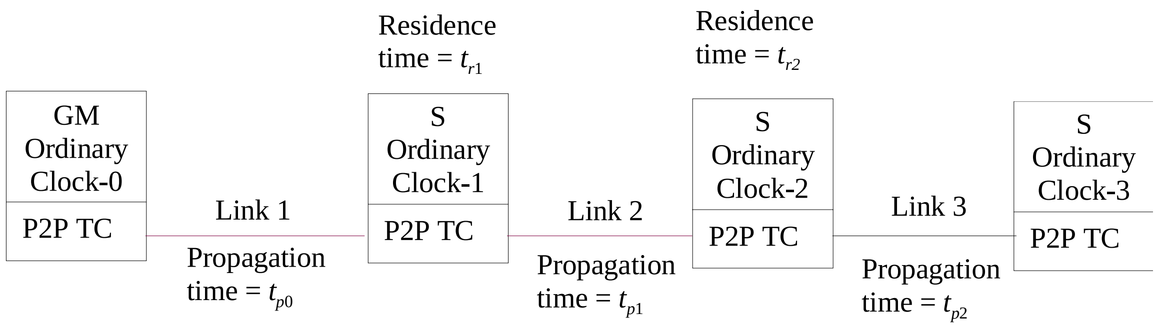

- The slave offset is computed when a Sync and corresponding Follow_Up message are received by the slave, using Eq. (2-4) of [4]

$$slave\_offset = t_2 - t_1 - total\_propagation\_plus\_residence\_time$$

- Where
  - t_2 = time that the Sync message is received by the slave
  - t_1 = time that the Sync message is sent by the GM

Review of Simulation Model (from [5])

- The total_propagation_plus_residence_time is given by Eqs. (2-2) and (2-3) in [4] (Eq. (2-2) shows how the needed sum is contained in the correction Sync and Follow_Up correction fields

$$total\_propagation\_plus\_residence\_time = \sum_{i=1}^{N-1} t_{ri} + \sum_{i=1}^{N} t_{pi}$$

$$Sync\_correction\_field + Follow\_Up\_correction\_field = \sum_{i=1}^{N-1} \left(t_{ri} + t_{pi}\right)$$

- The GM time t_1 is assumed to be perfect (and therefore has zero phase error)
- The times t_2, as well as the t_{ri} used to compute residence times, are taken from the syntonized P2P TC time (i.e., the flextimer). This is computed from the free-running P2P TC phase using Eqs. (2-5) – (2-9) of [4]
  - The method is summarized on the following 2 slides

Review of Simulation Model (from [5])

- GM time estimate, corresponding to time Sync is received at TC is

$$t_{GM} = t_1 + total\_propagation\_plus\_residence\_time$$

- Let
- $t_{GM,i}$ = estimated GM time at sync interval i based on the received Sync and Follow_Up messages
- $t_{b,i}$ = time indicated by free-running oscillator in P2P TC
- $y_{TC,i}$ = measured frequency offset of GM relative to free-running oscillator in P2P TC

$t_{f,i}$ = syntonized time, synthesized from the measured frequency offset and the time indicated by the free-running P2P TC oscillator

- Then

$$y_{TC,kM} = \frac{t_{GM,kM} - t_{GM,(k-1)M}}{t_{b,kM} - t_{b,(k-1)M}} - 1$$

$$y_{TC,kM+i} = y_{TC,kM} \quad \text{for } i = 1,2,\ldots,M$$

Review of Simulation Model (from [5])

- The syntonized time is synthesized by assuming that the frequency offset of the GM relative to the TC has been constant and equal to the current measurement since the last measurement was made. Then

$$1 + y_{TC,kM} = \frac{t_{f,kM+i} - t_{f,kM}}{t_{b,kM+i} - t_{b,kM}} \quad \text{for } i = 1,2,\ldots,M$$

- May solve for syntonized time

$$t_{f,kM+i} = t_{f,kM} + (1 + y_{TC,kM})(t_{b,kM+i} - t_{b,kM}) \quad \text{for } i = 1,2,\ldots,M$$

$$t_{f,kM+i} = t_{f,kM+i-1} + (1 + y_{TC,kM})(t_{b,kM+i} - t_{b,kM+i-1}) \quad \text{for } i = 1,2,\ldots,M$$

- The syntonized time is used to
  - Measure residence time
  - Measure propagation delay
  - Measure t_2

Review of Simulation Model (from [5])

- The free-running clock phase error is computed based on frequency offset (initialized randomly at initialization of the simulation run within the input frequency tolerance range), and the clock noise model described in [1] and [2] and the reference given there
- The simulation model computes only phase errors relative to the GM, which is assumed to be perfect
  - t_2 and the t_{ri} are syntonized clock phase offsets (final equation on slide 8)
  - Nominal residence time is assumed to be one sync interval
    - It is assumed that each Sync message is held at a TC until the corresponding Follow_Up message arrives
      - In worst case, the Follow_Up message takes on the order of a Sync interval to arrive (Follow_Up messages must be transported at least at the same average rate as Sync messages, otherwise there would be a monotonically increasing backlog of Sync and Follow_Up messages to process
  - Time at each node is shifted by the nominal propagation delay (this is equivalent to setting nominal propagation delay to zero) and only propagation delay measurement error is modeled
    - This is valid because, if propagation delay could be measured perfectly, it could be corrected for perfectly. Any remaining error is due to the propagation delay measurement error
  - Using the figure on the next slide (taken from [4]) to illustrate the exchange of Pdelay messages, the propagation delay measurement error can be modeled

Review of Simulation Model (from [5])

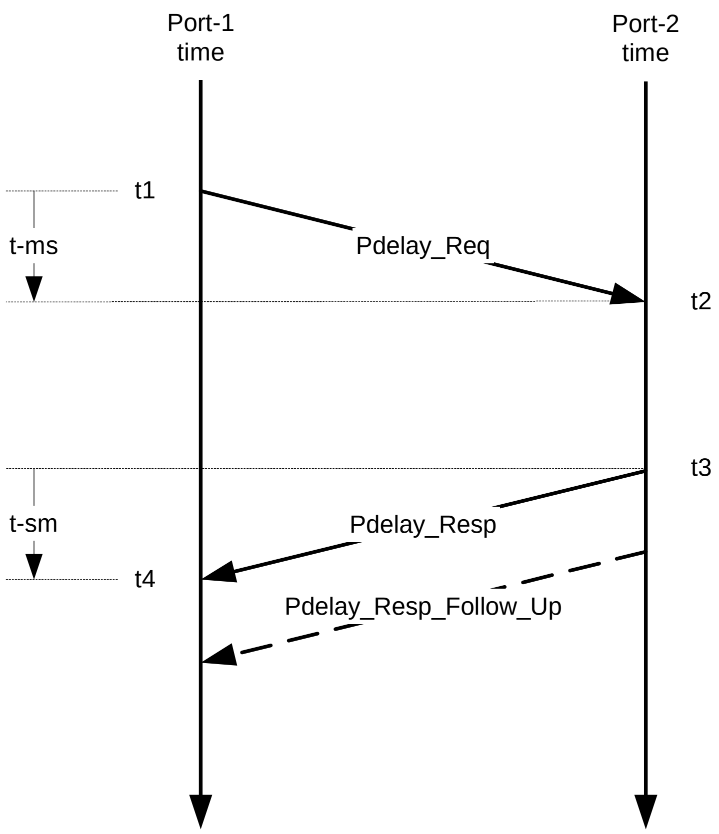

- Propagation delay error = 0.5*(t3 – t2) * (frequency offset between the nodes)
- Frequency offset is computed over one sync interval, using the syntonized timer values
- The Pdelay turnaround time, t3 – t2, is specified as an input parameter by the user
  - 1 ms is used in the cases here

Simulation Cases

- Parameters common to all simulation cases
  - Free-run clock accuracy = ± 100 ppm
    - Initialize frequency offset of each clock randomly within this range
  - Phase measurement granularity = 40 ns (25 MHz free-running clocks)
  - Frequency measurement granularity = 2.3283 x 10^{-10} (32 bit accuracy)
  - Consider up to 10 hops (AVB will require 7 hops; we exceed that to see how much MTIE increases beyond 7 hops)
  - Sync interval = 10 ms
  - Frequency update interval = 100 ms
  - Pdelay turnaround time (t3 – t2 on slide 10) = 1 ms
  - Free running clock has phase noise based on model used in [1] and [2]
  - Endpoint filter is 2^{nd} order, with 0.1 dB gain peaking
  - Asymmetry in PHY latency and cable delay not modeled
- We simulate both the use of P2P TCs and the scheme of [3]
  - The scheme of [3] does not need Pdelay turnaround time
  - The model for [3] does assume that message exchanges between successive nodes are synchronized (if this is not assumed, results will be worse)

Simulation Cases

- Case 1
  - 3 dB bandwidth = 1 Hz
- Case 2
  - 3 dB bandwidth = 0.1 Hz
- Case 3
  - 3 dB bandwidth = 0.01 Hz
- In all cases, MTIE is computed after the initial phase transient has decayed

Case 1 Results from [5] Based on Single Run - P2P TC

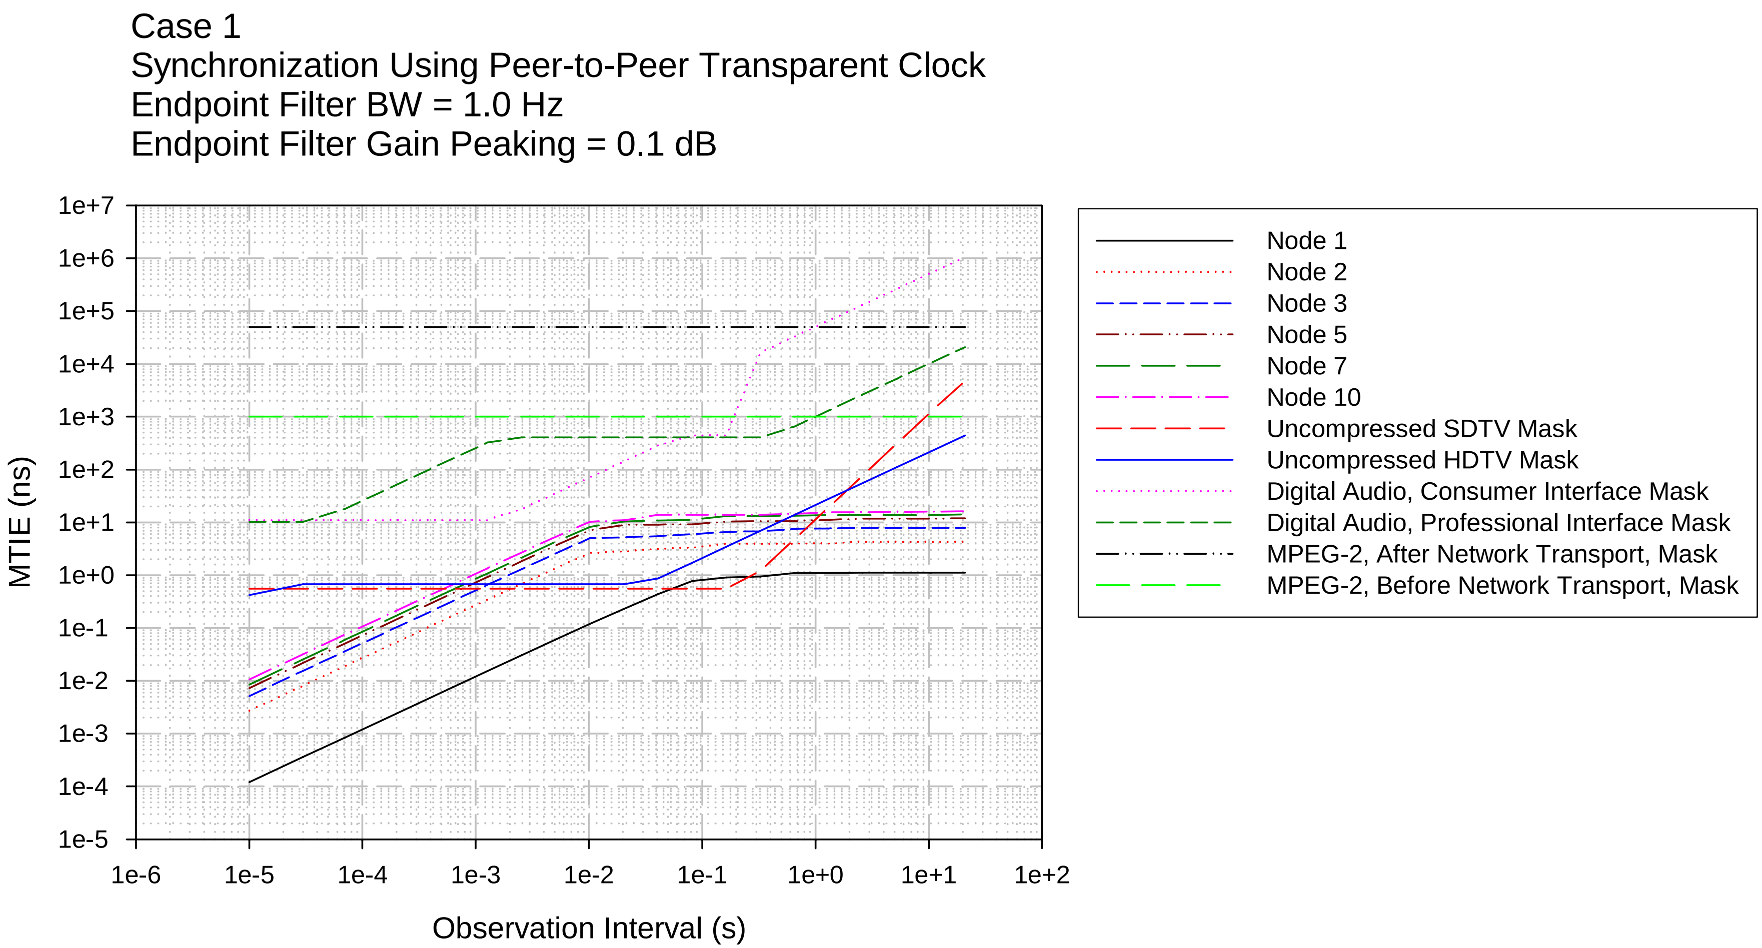

Case 1 Results for 300 Replications - P2P TC

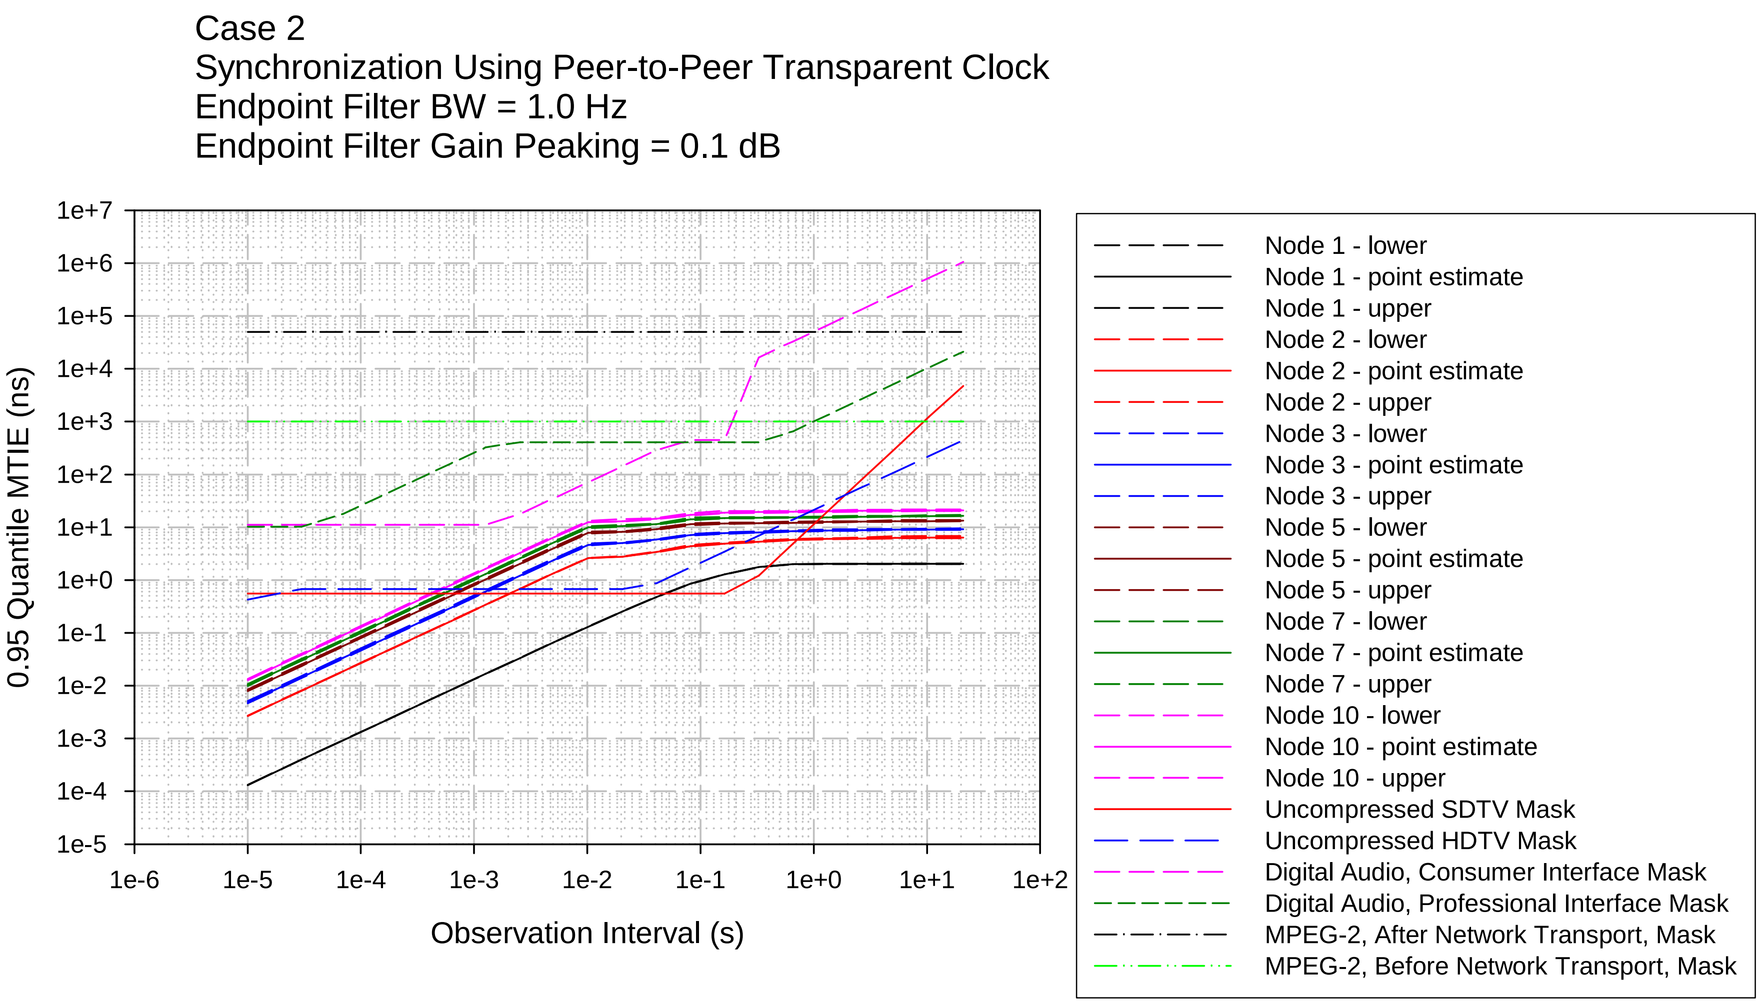

Case 1 Results from [5] Based on Single Run - Scheme of [3]

Case 1
Synchronization Using AVB White Paper Scheme
Message exchanges between successive nodes synchronized
Endpoint Filter BW = 1.0 Hz
Endpoint Filter Gain Peaking = 0.1 dB

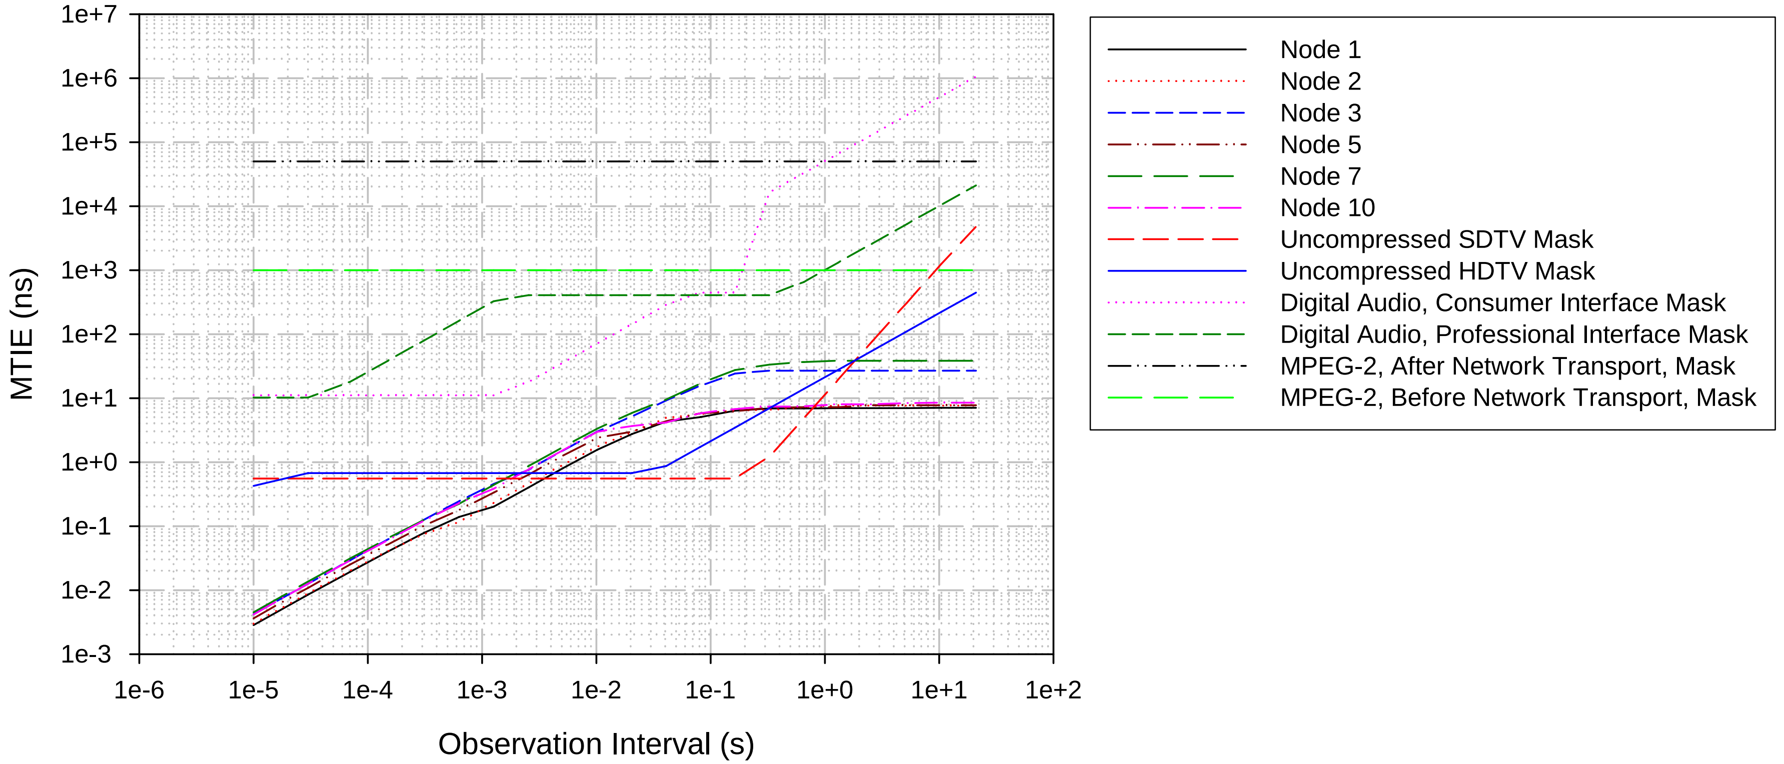

Case 1 Results for 300 Replications - Scheme of [3]

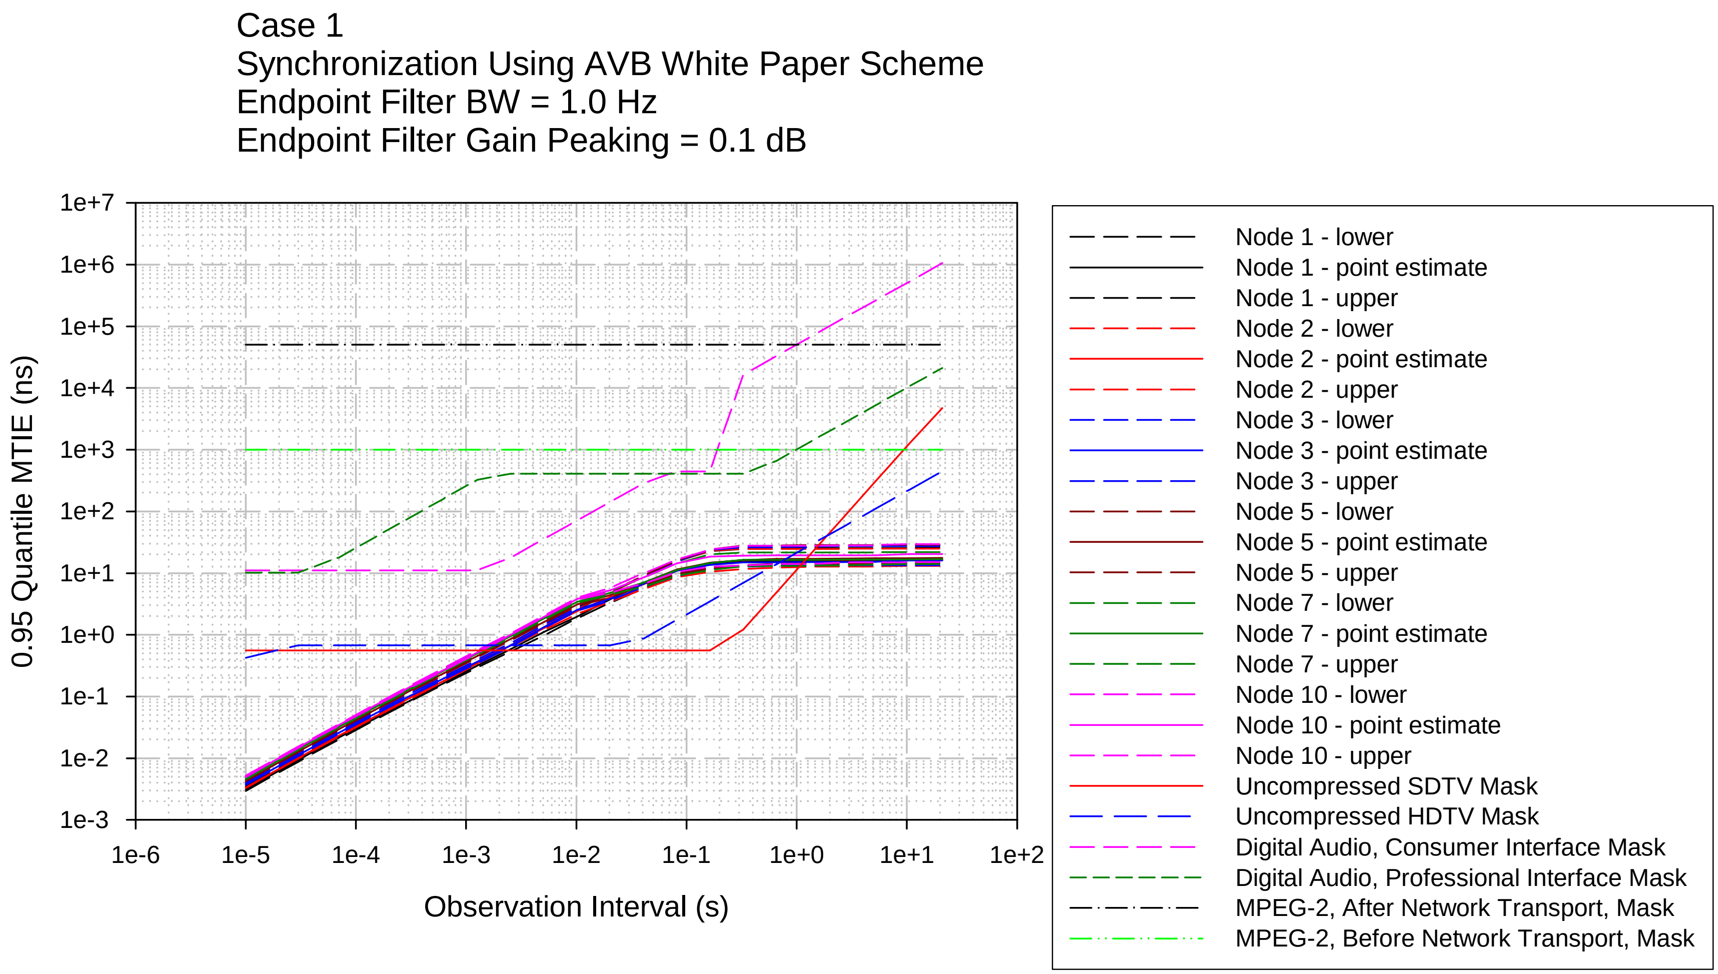

Comparison of single run and 300-run results for Case 1

- Results for P2P TC transport
  - MTIE results based on 300 independent runs are above those from [5] for single run, as expected
    - Often the degree to which the single run results are exceeded is not huge
- Results for transport using scheme of [3]
  - MTIE results based on 300 independent runs are above those from [5] for single run, for nodes 1, 2, 4-6, and 8-10, as expected
  - MTIE results based on MTIE results based on 300 independent runs are below those from [5] for single run, for nodes 3 and 7
  - Phase error plots on the following 2 slides show low frequency, transient-like behavior for the single run results from [5]
    - This type of behavior can be produced by specific combinations of frequency offsets at successive nodes
- Note that the MTIE results based on 300 runs are 0.95 quantiles
  - These MTIE values can be exceeded in 5% of cases
  - This can occur in both the above cases (it happened to occur in the case based on the scheme of [3] here)

Case 1, Node 3 Results from [5] Based on Single Run - Scheme of [3]

Case 1, Node 3
Synchronization Using AVB White Paper Scheme
Endpoint Filter BW = 1.0 Hz
Endpoint Filter Gain Peaking = 0.1 dB

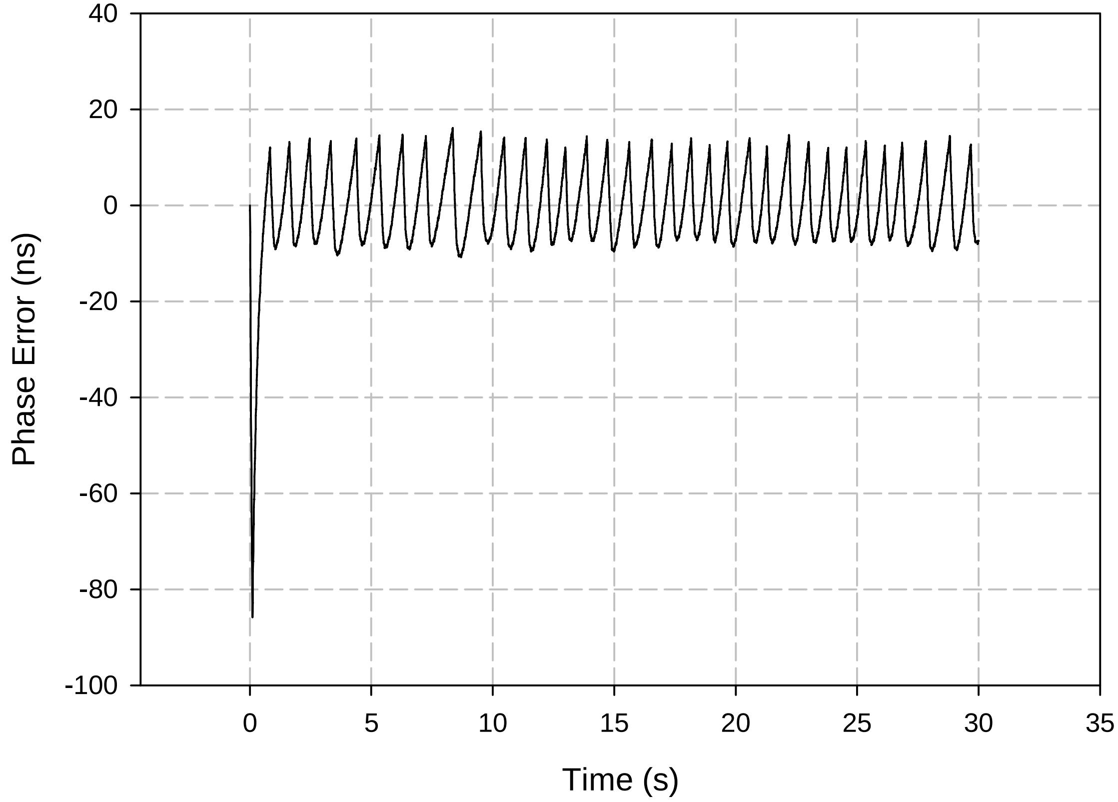

Case 1, Node 7 Results from [5] Based on Single Run - Scheme of [3]

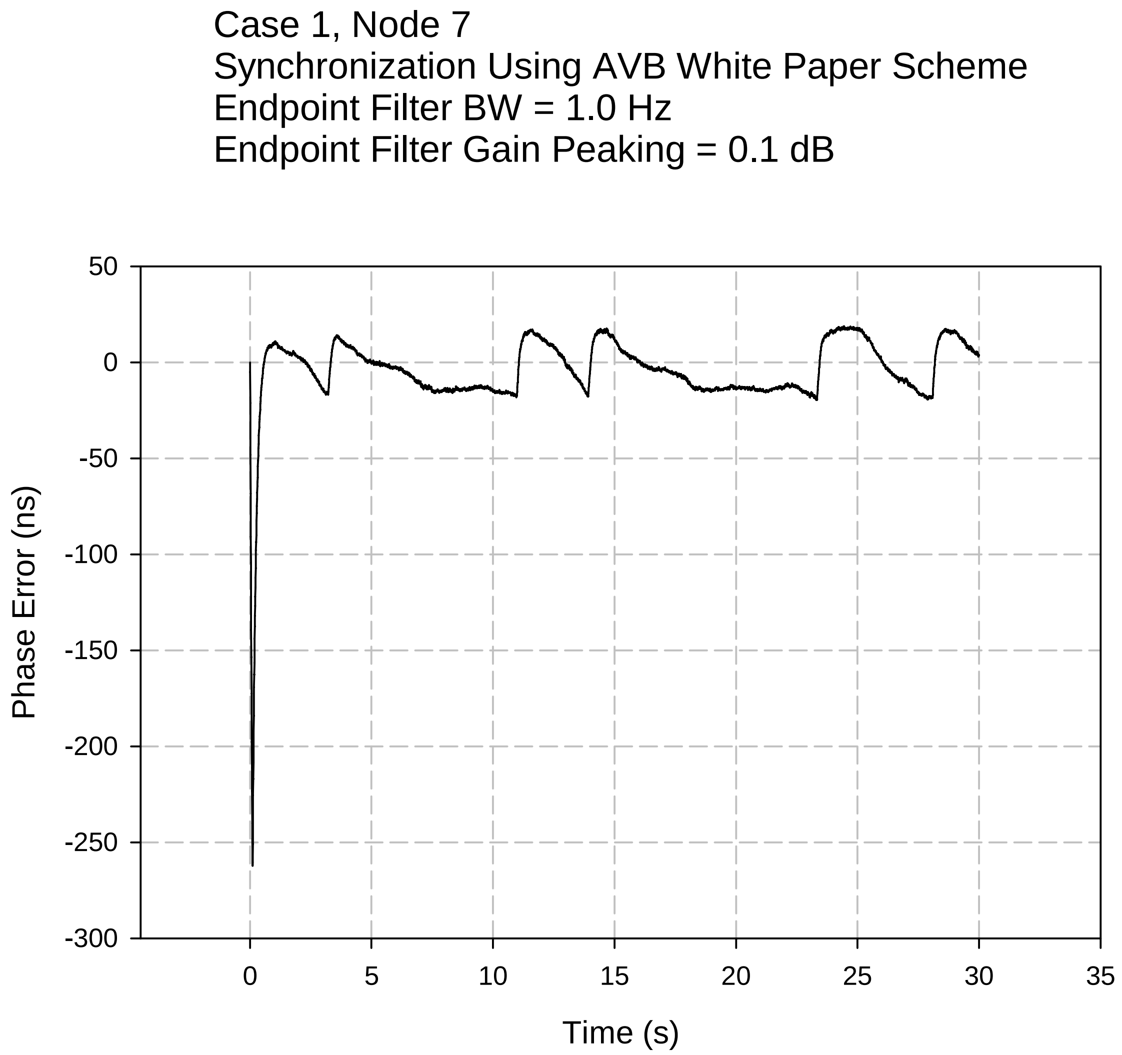

Case 2 Results from [5] Based on Single Run – P2P TC

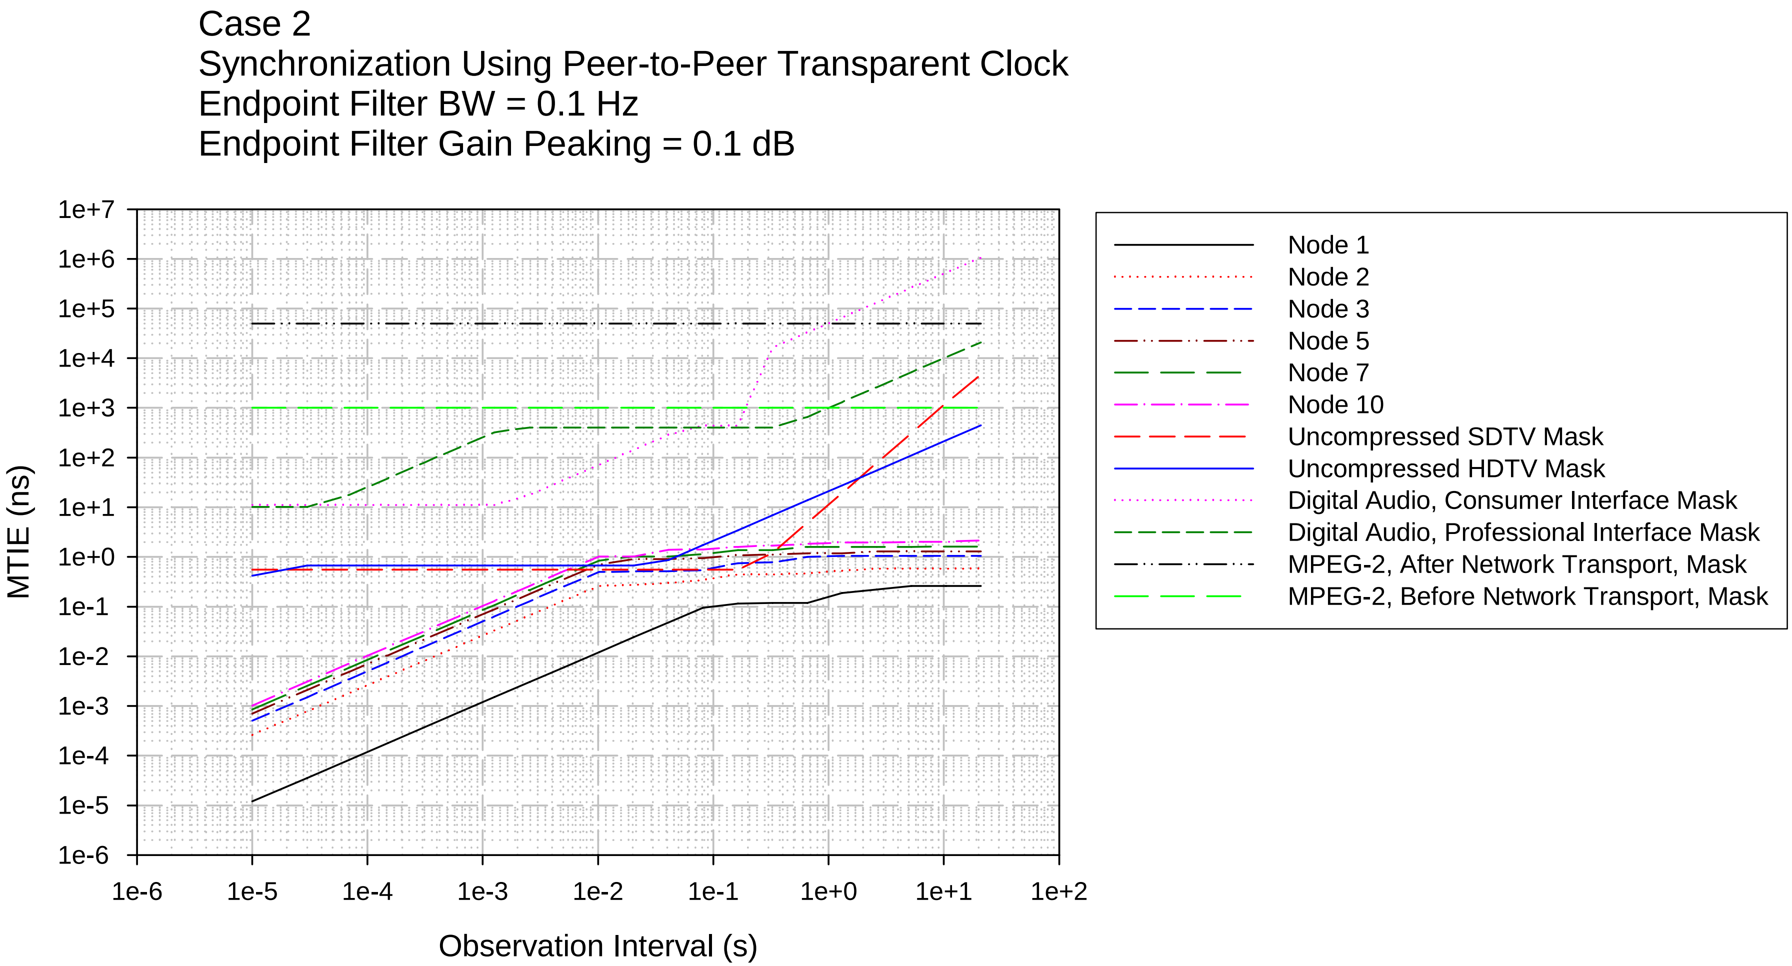

Case 2 Results for 300 Replications - P2P TC

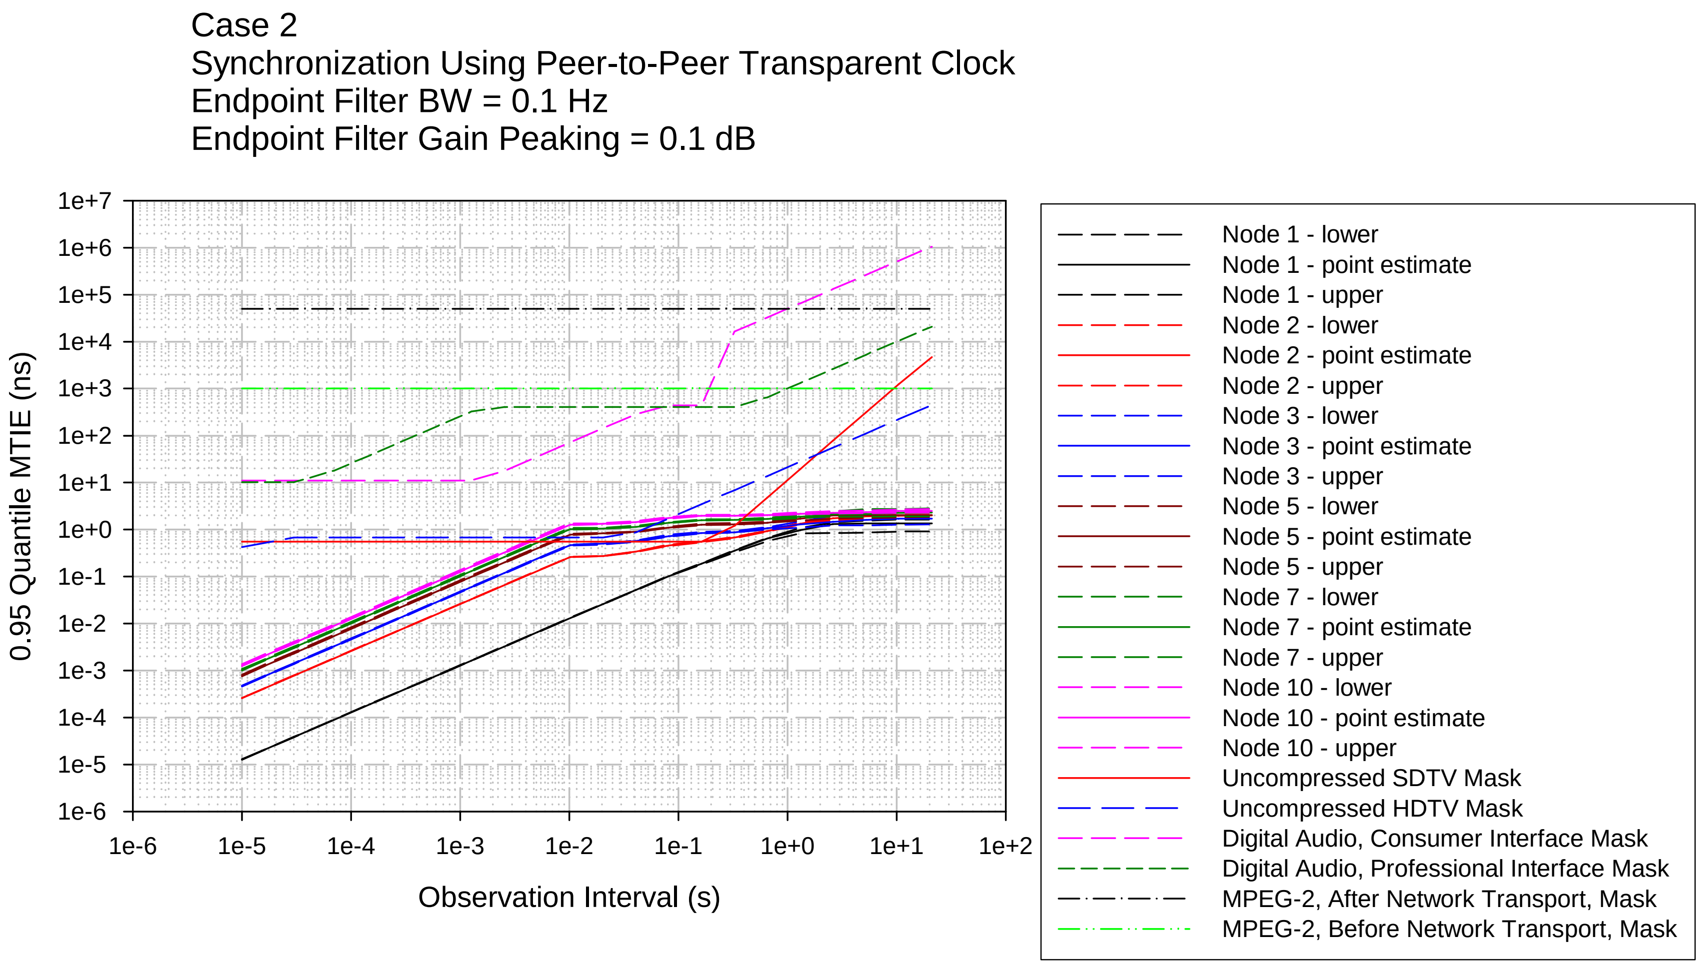

Case 2 Results from [5] Based on Single Run - Scheme of [3]

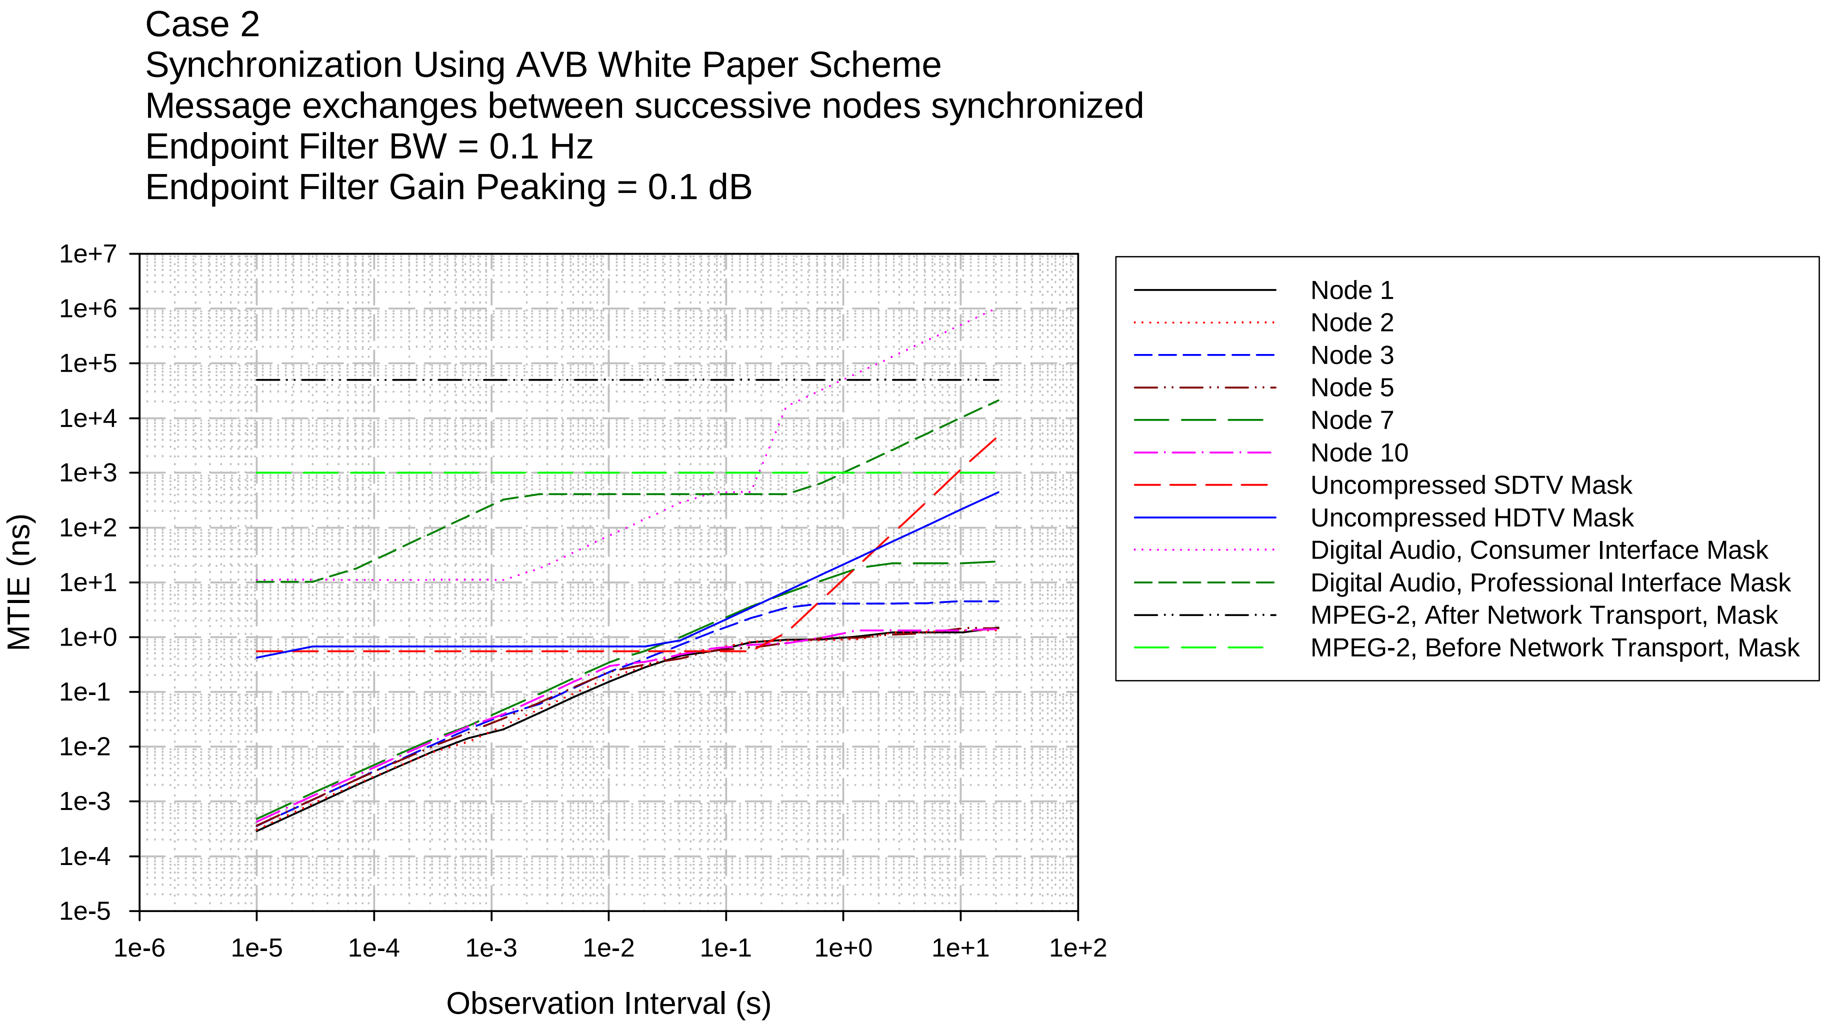

Case 2 Results for 300 Replications - Scheme of [3]

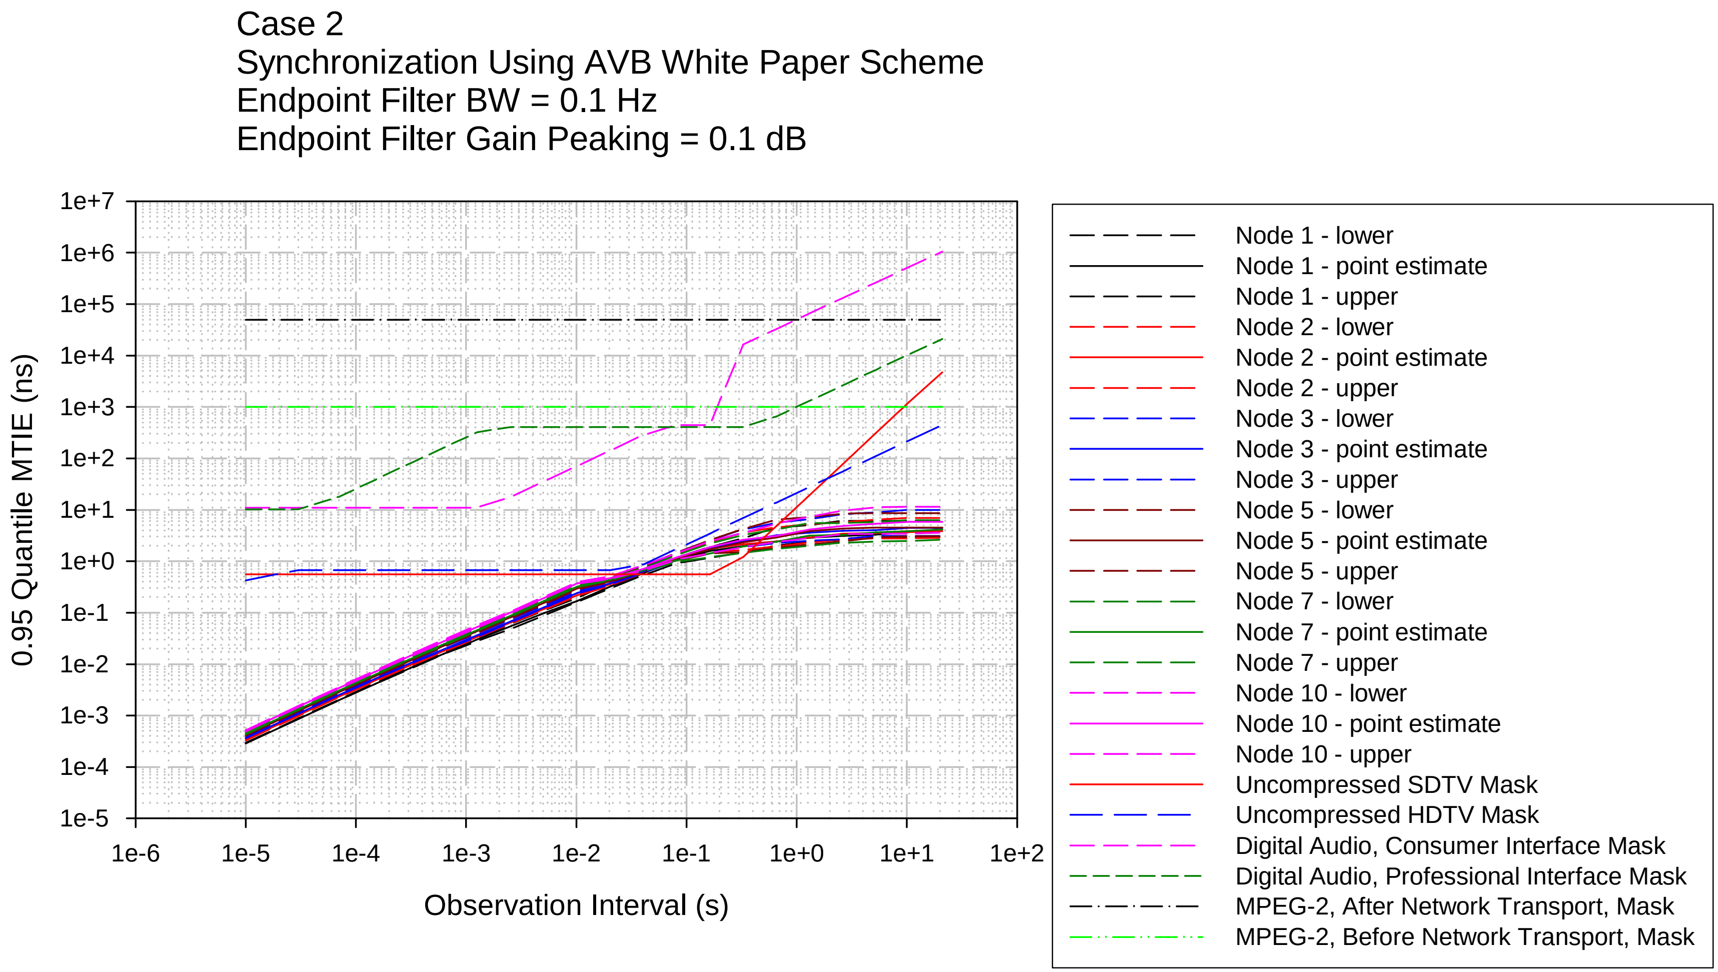

Case 3 Results from [5] Based on Single Run - P2P TC

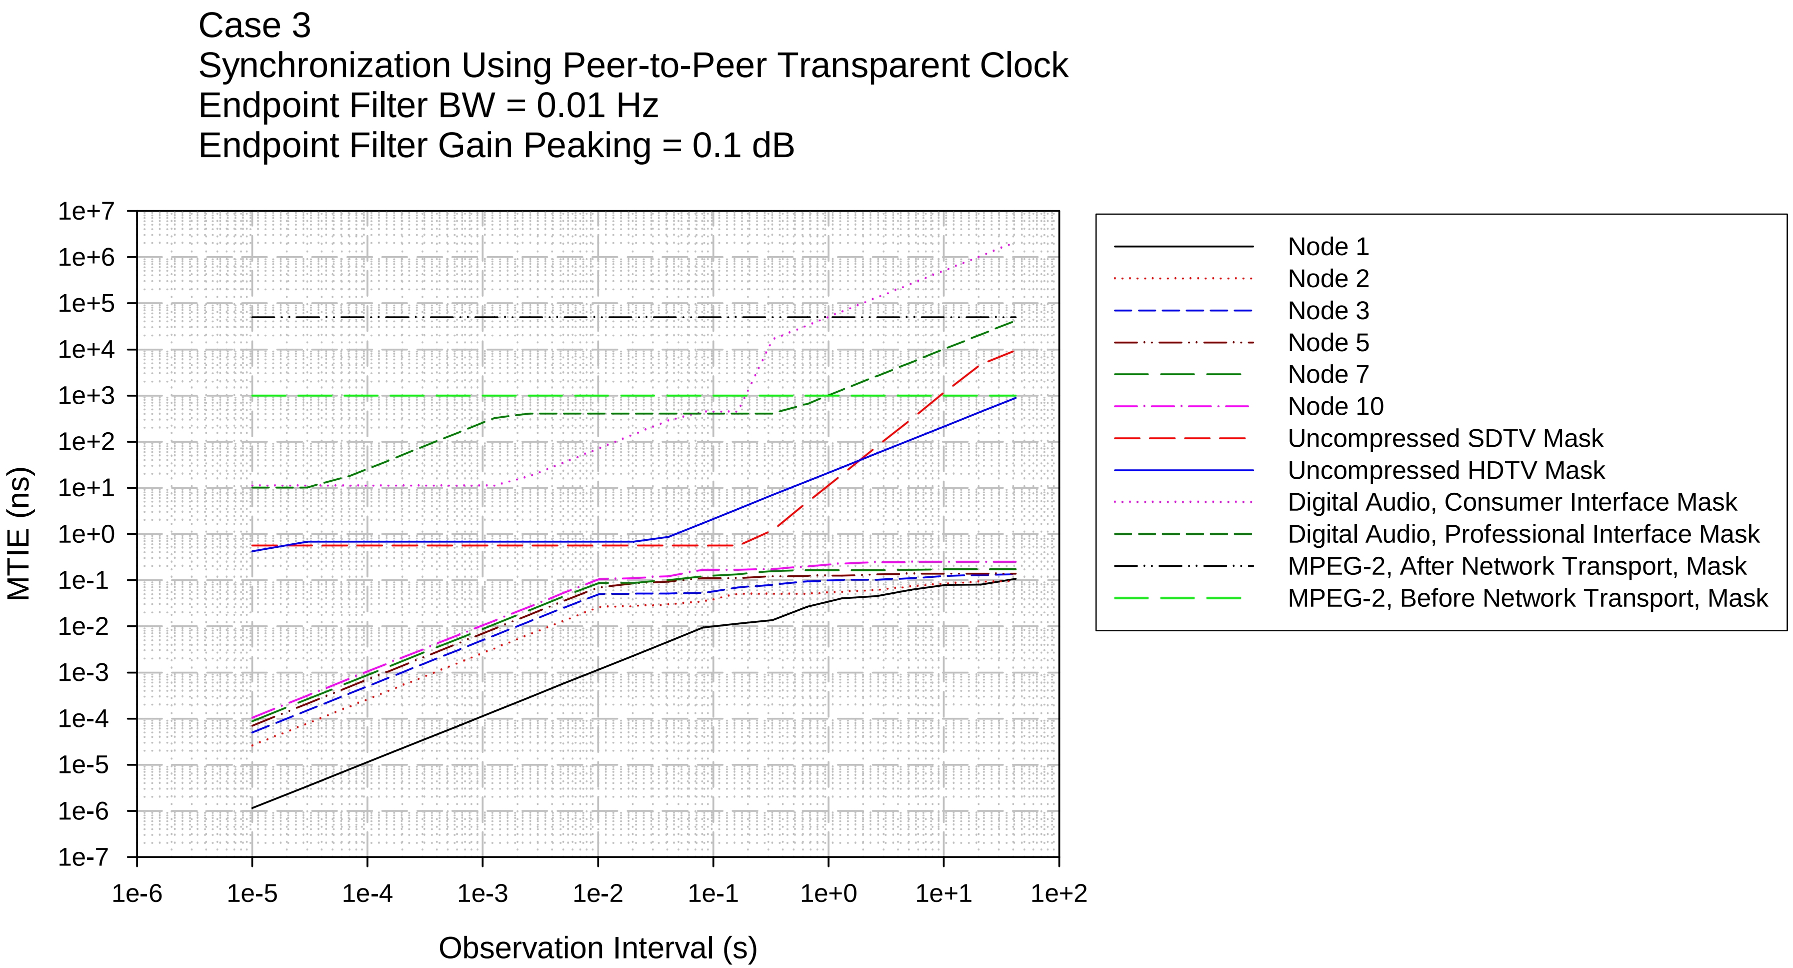

Case 3 Results for 300 Replications – P2P TC

Case 3
Synchronization Using Peer-to-Peer Transparent Clock
Endpoint Filter BW = 0.01 Hz
Endpoint Filter Gain Peaking = 0.1 dB

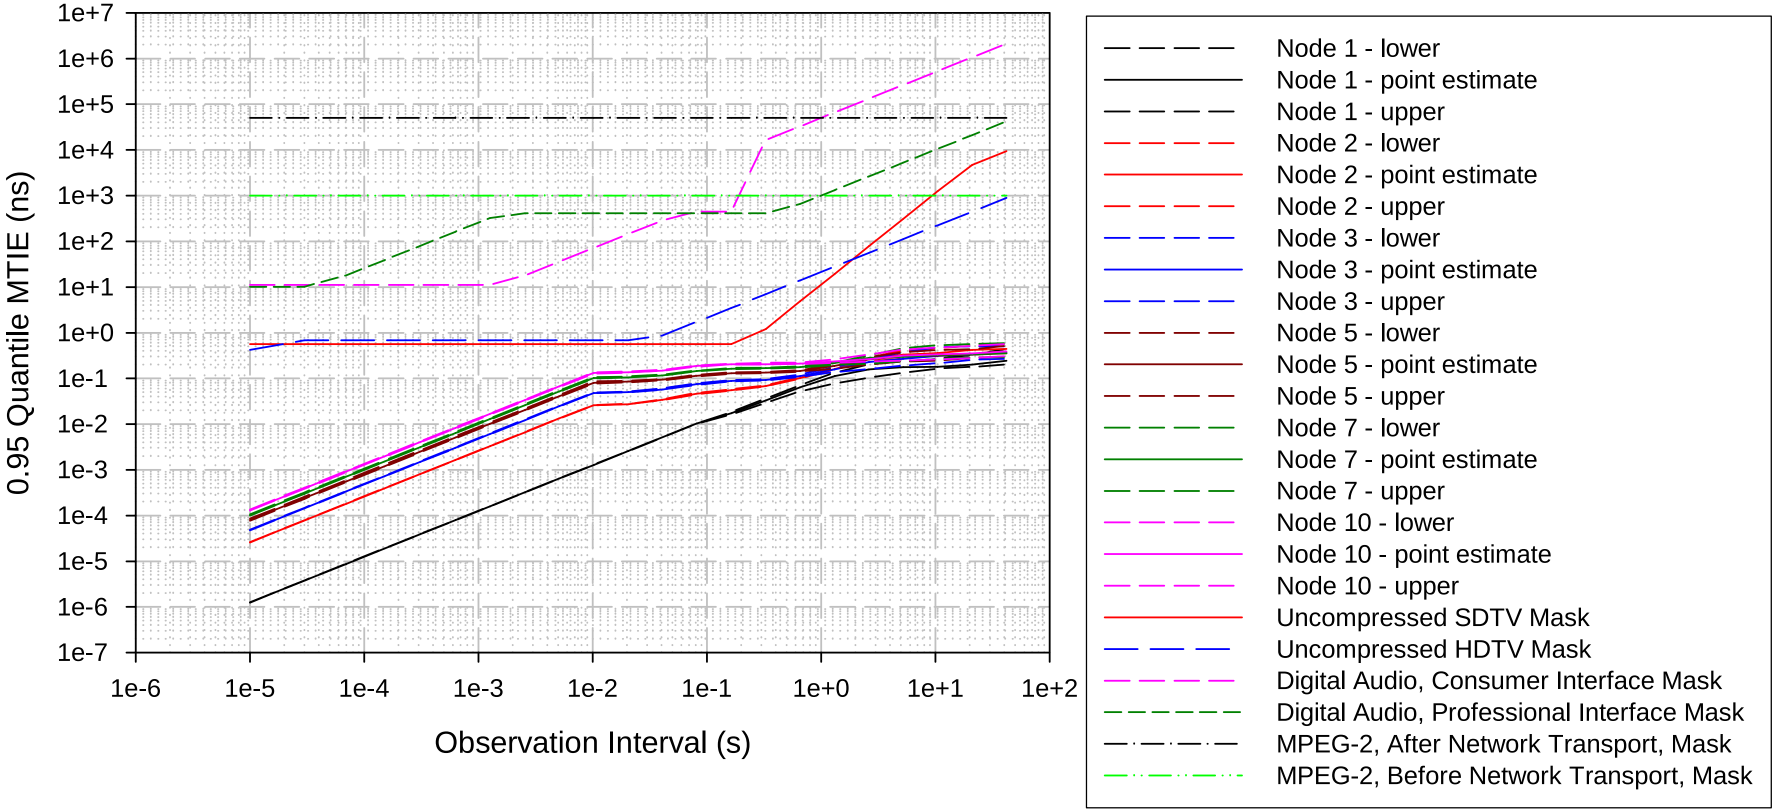

Case 3 Results from [5] Based on Single Run - Scheme of [3]

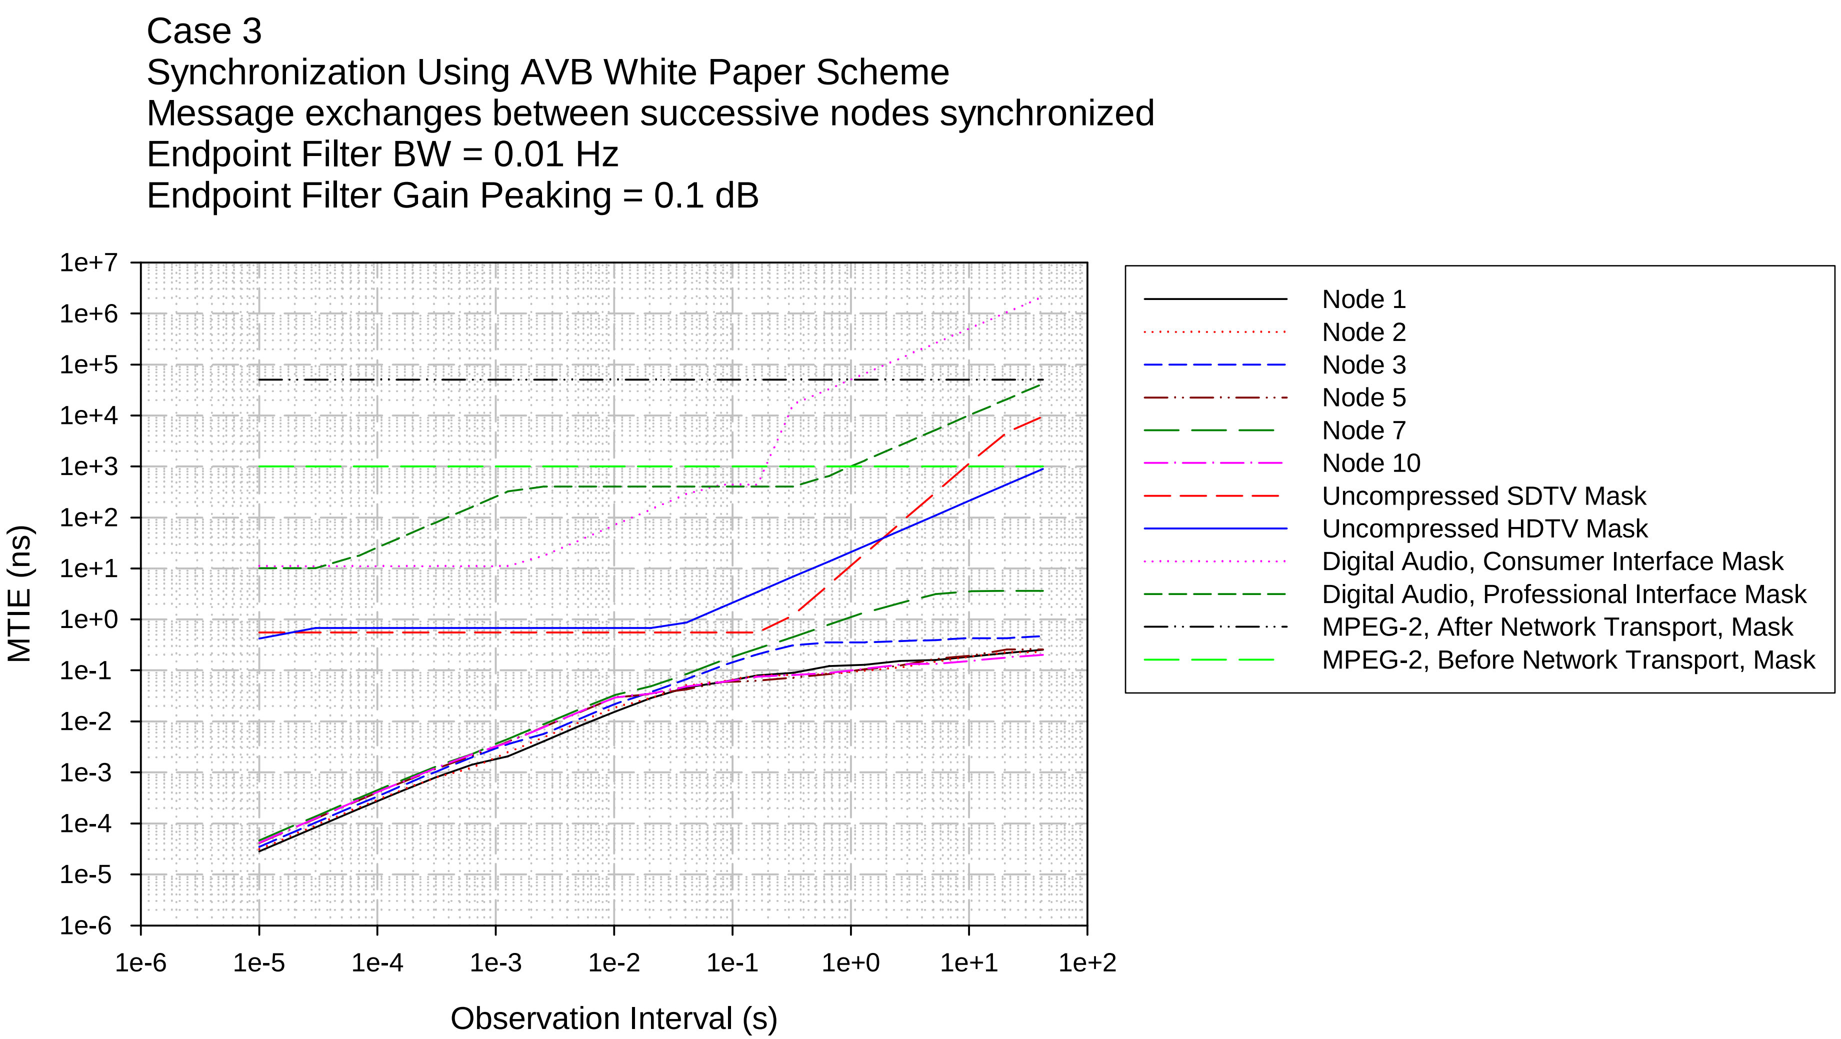

Case 3 Results for 300 Replications - Scheme of [3]

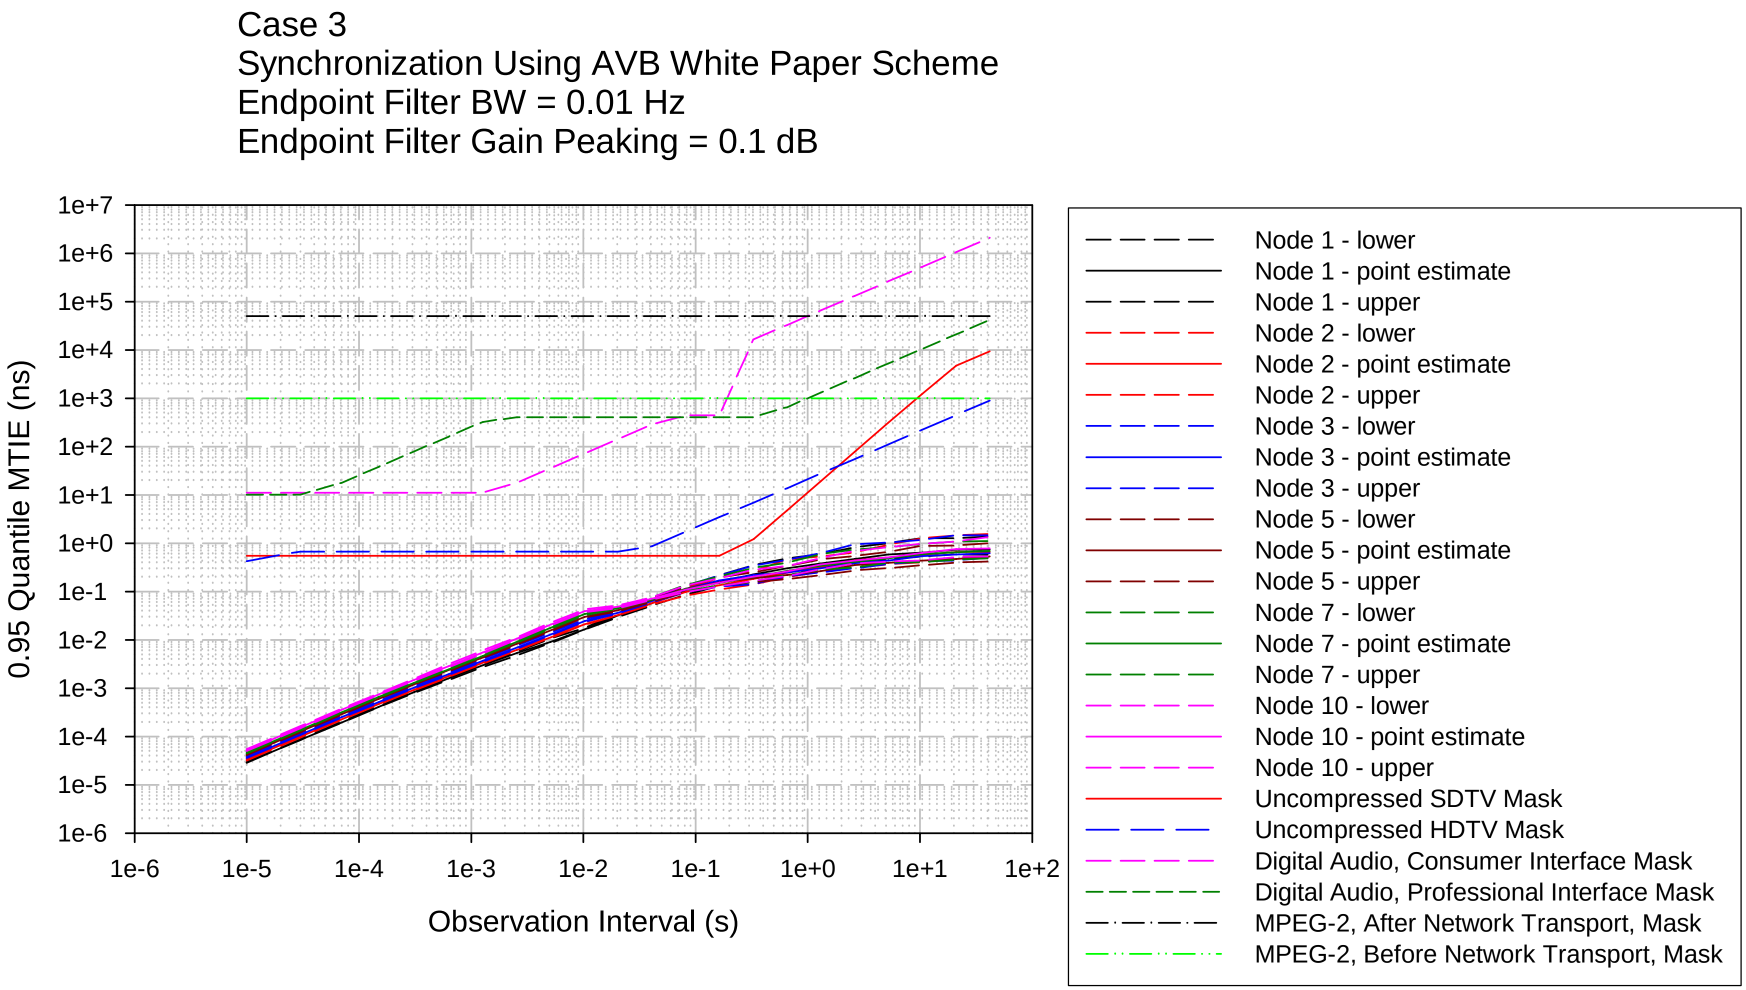

Comparison of Node 1 Results for all Cases - 300 Replications

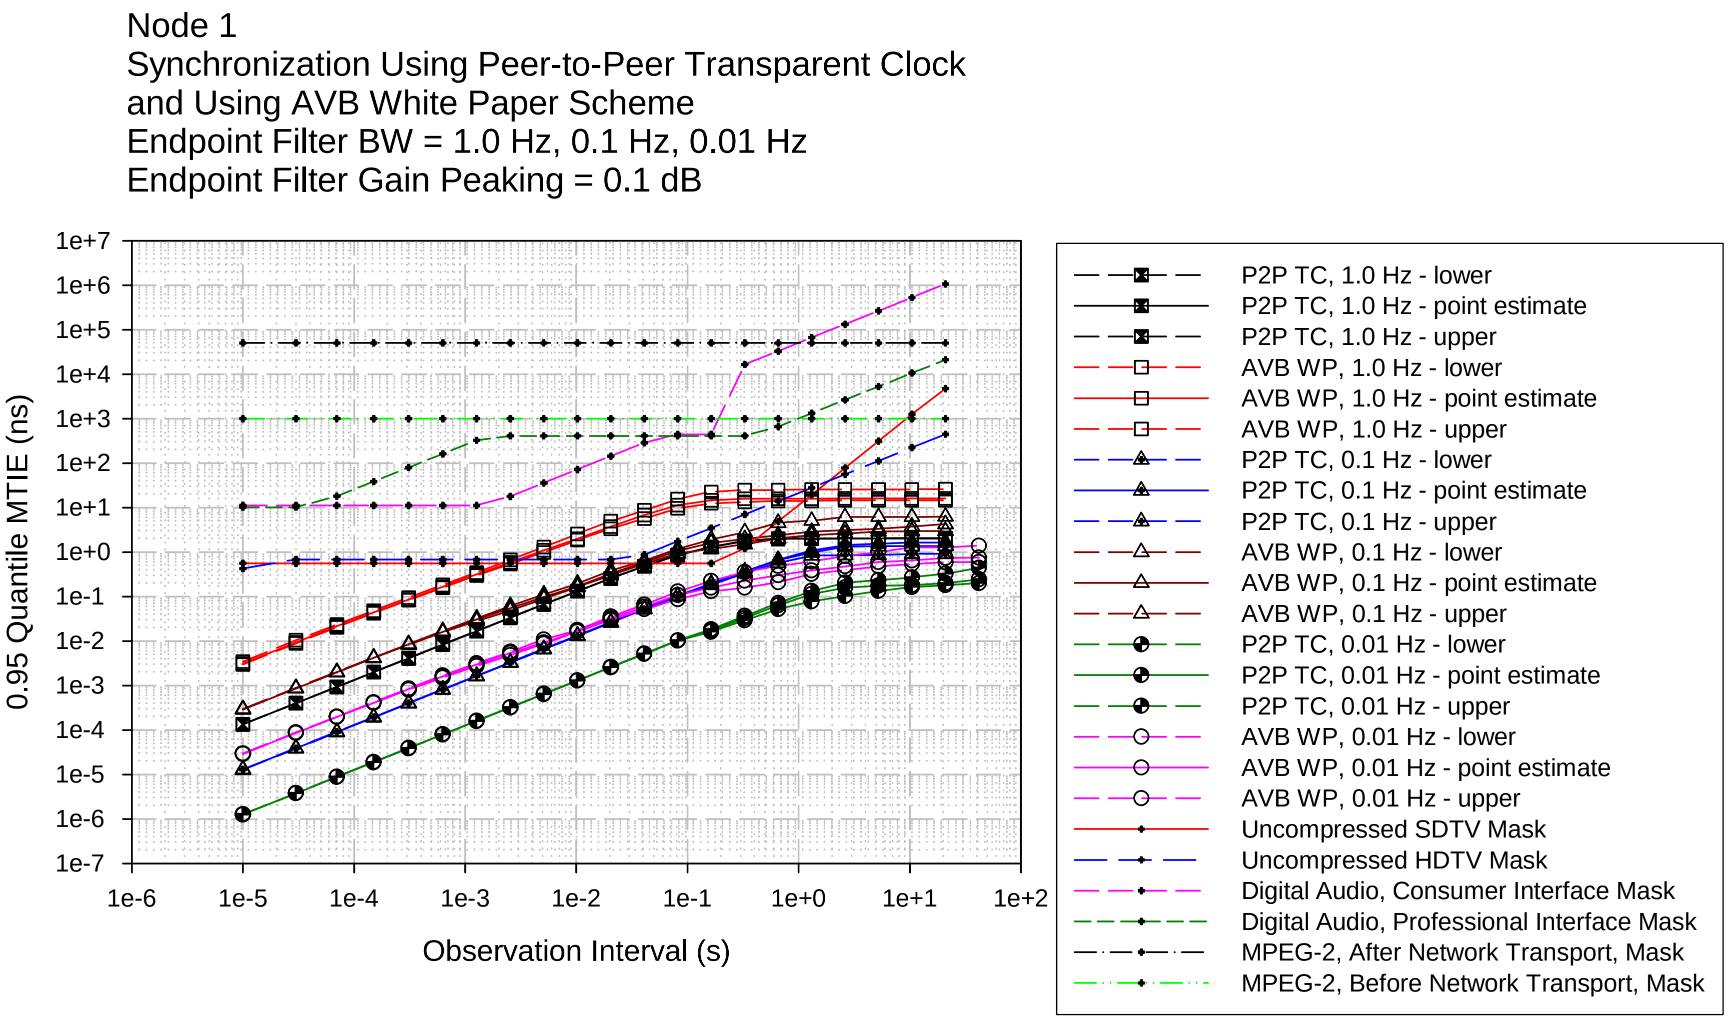

Comparison of Node 7 Results for all Cases – 300 Replications

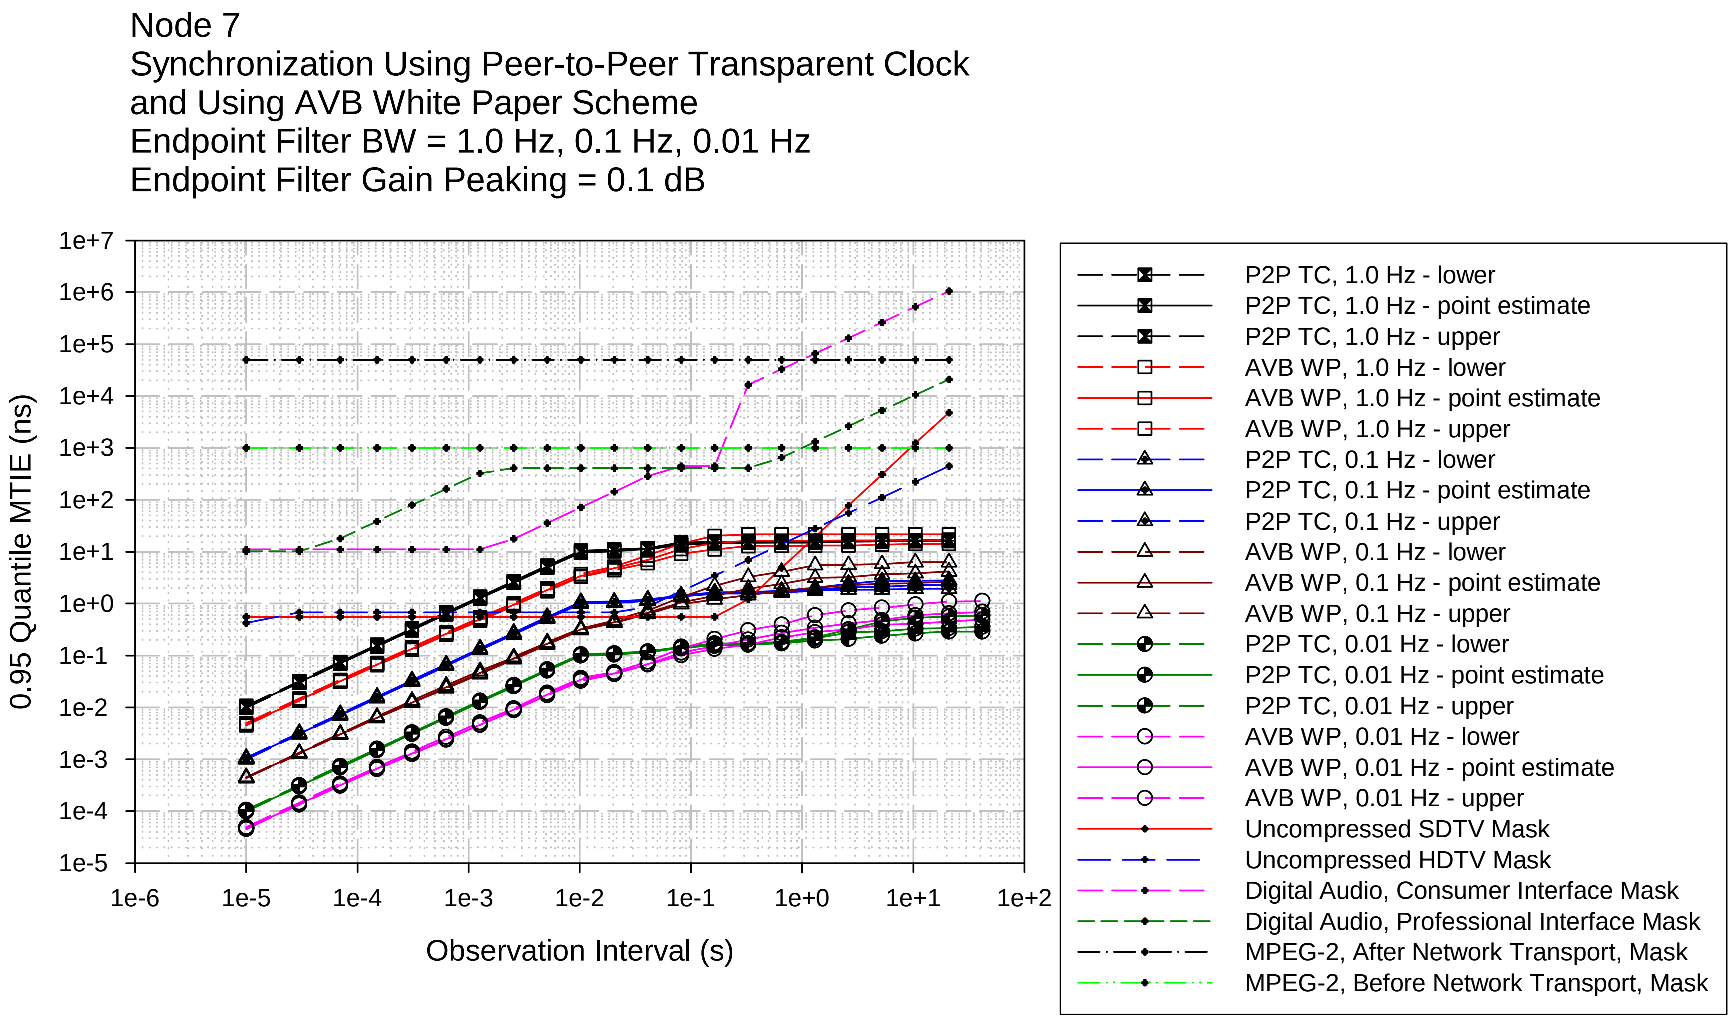

Discussion of Results

- Results obtained here are similar to those of [5] based on single runs
  - For a single hop, the scheme using P2P TCs gives somewhat better performance
    - MTIE for this scheme was lower by approximately an order of magnitude for observation intervals ranging from 10 μs to 20 s (see slide 29)
  - For more than one hop, the schemes give similar performance, with the scheme of [3] slightly better for shorter observation intervals and the P2P TC based scheme slightly better for longer observation intervals
  - As in [5], results for scheme of [3] assume that message exchanges between successive nodes are synchronized (i.e., node *n* sends Sync message to node *n*+1 immediately after it has computed slave offset for message exchange with node *n*-1
    - If this is not done, results for this scheme will be worse
- Both methods meet MTIE masks for uncompressed digital video with 0.01 Hz filter
- Both methods meet MTIE masks for digital audio with 1 Hz filter
  - Results of [2] suggest that 10 Hz filter may be ok for this case (these results are for the scheme of [3], but the results in the current presentation for narrower bandwidths indicate that the results of the two schemes are similar)

Discussion of Results

- For cases that use P2P TCs, examination of the computed residence times and propagation delay errors indicates that the main contribution to phase error is the effect of the 40 ns phase measurement granularity
  - The results for a single hop are much better than those for multiple hops for this scheme because, for a single hop, residence time is not used
- Each residence time measurement is truncated to the next lower multiple of 40 ns
  - The truncated errors accumulate; when they reach 40 ns, the residence time measurement jumps by 40 ns
  - On the next residence time measurement, i.e., the next sync interval, the residence time jumps back to its previous value, which is 40 ns smaller
  - The resulting residence time measurement history is approximately constant, with 40 ns pulses that last for 1 sync interval occurring at some frequency
  - In the worst case, we may assume that the time between successive pulses is much longer than the endpoint filter time constant
  - Since the pulse time is much shorter than the filter time constant for the bandwidths considered here (e.g., the shortest filter time constant is $1/[(2\pi)(1 \text{ Hz})] = 0.16$ s, while the pulse width is equal to the sync interval, or 0.01 s), the pulses may be modeled as impulses input to the filter

Discussion of Results

- Consider a linear, 2nd order filter with undamped natural frequency ω_n, damping ratio ζ, and 20 dB/decade roll-off
  - The transfer function is given by

$$H(s) = \frac{2\varsigma\omega_n s + \omega_n^2}{s^2 + 2\varsigma\omega_n s + \omega_n^2}$$

- The impulse response may be obtained by taking the inverse Laplace Transform of the transfer function; the result is

$$h(t) = \omega_n e^{-\varsigma\omega_n t}\left[2\varsigma \cosh\sqrt{\varsigma^2 - 1}\,\omega_n t + \frac{1 - 2\varsigma^2}{\sqrt{\varsigma^2 - 1}}\sinh\sqrt{\varsigma^2 - 1}\,\omega_n t\right]$$

- The impulse response of this filter takes on its maximum value at time zero
  - The maximum value is equal to $2\zeta\omega_n$
- Since the phase measurement is always truncated to the next lower multiple of 40 ns, the filtered phase error contribution of one node is equal, in worst case, to $2\zeta\omega_n(40\ \mathrm{ns})$

Discussion of Results

- Can use this as a rule of thumb to conservatively estimate the MTIE contribution from phase measurement granularity at one node
- For example, for a 0.1 Hz filter with 0.1 dB gain peaking, the undamped natural frequency and damping ratio are 0.071781 rad/s and 4.3188, respectively
  - The filtered phase error due to a 40 ns impulse is 0.62 ns
  - Long-term MTIE for 10 nodes for this case is approximately 2.1 ns for the single run of [5] (slide 21) and 2.8 ns for the 0.95 quantile estimate based on 300 runs (slide 22)
    - For the single run of [5], there were no instances where pulses from all 10 nodes lined up in time
    - The 300 runs also did not produce any instances where the pulses from all 10 nodes lined up in time
    - If the pulses from all 10 nodes lined up in time, the resulting long-term MTIE for node 10 would be on the order of 6.2 ns
- Note that while the rule of thumb gives an estimate of long-term MTIE, it does not give MTIE as a function of observation interval and, in particular, short observation intervals
  - Simulation is necessary to obtain the actual MTIE performance for various observation intervals, and also to include other effects (e.g., errors in propagation delay measurement, errors due to asymmetries in cable delay and PHY latency)

Future Work

- Modeling of asymmetry in cable delay
  - If cable delays in the 2 direction on a link are different and unknown, but fixed, the asymmetry will contribute to an error in time synchronization but not jitter/wander
- Modeling of asymmetry and uncertainty in PHY latency
  - Likely assume that the time stamp measurement is made at the (G)MII
  - Results will depend on magnitude, asymmetry, and time variation in unknown PHY delay
    - Fast variation in unknown PHY delay may be easy to filter and result in lower jitter/wander
    - Slow variation in unknown PHY delay may be difficult to filter
    - Asymmetry in PHY delay in the two directions that is fixed has the same effect of asymmetry in cable delay that is fixed
- Additional cases with different parameters

References

1. Geoffrey M. Garner and Kees den Hollander, *Analysis of Clock Synchronization Approaches for Residential Ethernet*, Samsung presentation to IEEE 802.3 ResE SG, San Jose, CA, September, 2005.
2. Geoffrey M. Garner, *Additional Simulation Results for ResE Synchronization Using Filtered Phase and Instantaneous Frequency Adjustments*, Samsung presentation to IEEE 802.3 ResB SG, Vancouver, BC, November, 2005.
3. *Residential Ethernet(RE) (a working paper)*, Draft 0.142, maintained by David V. James and based on work by him and other contributors, November 16, 2005.
4. Geoffrey M. Garner, *Description of Use of IEEE 1588 Peer-to-Peer Transparent Clocks in A/V Bridging Networks*, Revision 2, working paper, May 12, 2006 (working paper presented at May, 2006 IEEE 802.1 Interim Meeting).
5. Geoffrey M. Garner, *Initial Simulation Results for AVB Synchronization Transport using IEEE 1588 Peer-to-Peer Transparent Clocks*, Samsung presentation to IEEE 802.1, Beijing, May, 2006.
6. Athanasiois Papoulis, *Probability, Random Variables, and Stochastic Processes*, Third Edition, McGraw-Hill, 1991, pp. 254 – 255.

Appendix – Confidence Interval for a Quantile of a Distribution Obtained from Samples of the Distribution

- Note: The material in this Appendix is contained in standard references on Probability and Statistics; e.g., see [6]
- Let X_i, $i = 1, 2, \ldots, N$ be samples of a population. In the case here , X_i is a sample of MTIE for a particular observation interval S. For convenience, we omit writing the S dependence explicitly (i.e., we should really write $X_i(S)$)
- Assume the samples are independent of each other. Since the samples are of the same population, they are identically distributed
- Assume the samples have been placed in ascending order, i.e.,

$$X_1 \le X_2 \le X_3 \le \cdots \le X_N$$

- Let x_β be the β^{th} quantile of the population distribution. The event

$$X_r \le x_\beta \le X_s$$

occurs if at least r and at most s-1 of the samples are less than x_β. The probability that any sample is less than x_β is β.

Appendix (Cont.)

- Then the probability that $X_r \leq x_\beta \leq X_s$ is

$$P\{X_r \leq x_\beta \leq X_s\} = \sum_{k=r}^{s-1} \frac{N!}{k!(N-k)!} \beta^k (1-\beta)^{N-k}$$

- $\beta = 0.95$ for the 0.95 quantile. For $N = 300$, $r = 275$, and $s = 294$, computation of the above binomial sum gives

$$P\{X_r \leq x_\beta \leq X_s\} = 0.991$$

- I.e., this is a 99% confidence interval for the 0.95 quantile