Initial 60802 Error Generation Time Series Simulation Results Version 1

Geoffrey M. Garner
Analog Devices (Consultant)

gmgarner@alum.mit.edu

IEEE 802.1 TSN TG
2024.01.22

Introduction - 1

- Tables 12 and 13 of IEC/IEEE 60802/D2.1 [1] contain error generation limits for a PTP Relay Instance and a PTP End Instance, respectively
- Annex D, Subclause D.4, of 60802 describes an approach to testing these requirements
  - While D.4 is informative and is not a test specification, it provides sufficient information to enable simulations of possible test setups to be performed, to see whether meeting the error generation limits is reasonable
  - A more recent version of D.4 is contained in [2]
- In [3], Monte Carlo simulation results are given for the error generation tests of D.4 for a PTP Relay Instance, and compare results for cases with and without the use of drift tracking and compensation algorithms
  - The results in [3] (see slide 16) meet the requirements of Table 12 of [1]
- One of the next steps described in [3] is to run time series simulations for the test cases of Annex D of [1], for both PTP Relay Instances and PTP End Instances
  - Time series simulations are needed for PTP End instances in particular because dTE results End Instances are after any endpoint filtering

Introduction - 2

- The current presentation describes initial time series (i.e., time domain) simulation results
- These initial results are based on single replications of each simulation case
  - If necessary, multiple replications of the simulation cases can be run subsequently

Description of Test Setup and Model for Simulator - 1

- A possible setup for testing a PTP Relay Instance is shown in Figures D.2, D.3, and D.4 of [1] and [2]
- A possible setup for testing a PTP End Instance is shown in Figure D.5 of [1] and [2]
- For convenience, Figures D.2 and D.5 of [1] and [2] are reproduced on the next slide (Figures D.3 and D.4 differ from D.2 only on that the three figures label different outputs for the different tests)
- In the time series simulator, the ClockSource and LocalClock at the Grandmaster (GM) PTP Instance are the same clock, while in Annex D.4 they are different clocks
- Therefore, the Emulated ClockSource and Emulated LocalClock of Figures D.2 and D.5 of [1] and [2] must be modeled as two separate nodes in the time series simulator
- In addition, since the test cases require the values of fields sent in the Sync/Follow_Up messages of the device under test (DUT), a node that follows the DUT is needed in the time series simulation model

Description of Test Setup and Model for Simulator - 2

- Figure D.2 of [1] and [2] – possible test setup for PTP Relay Instance

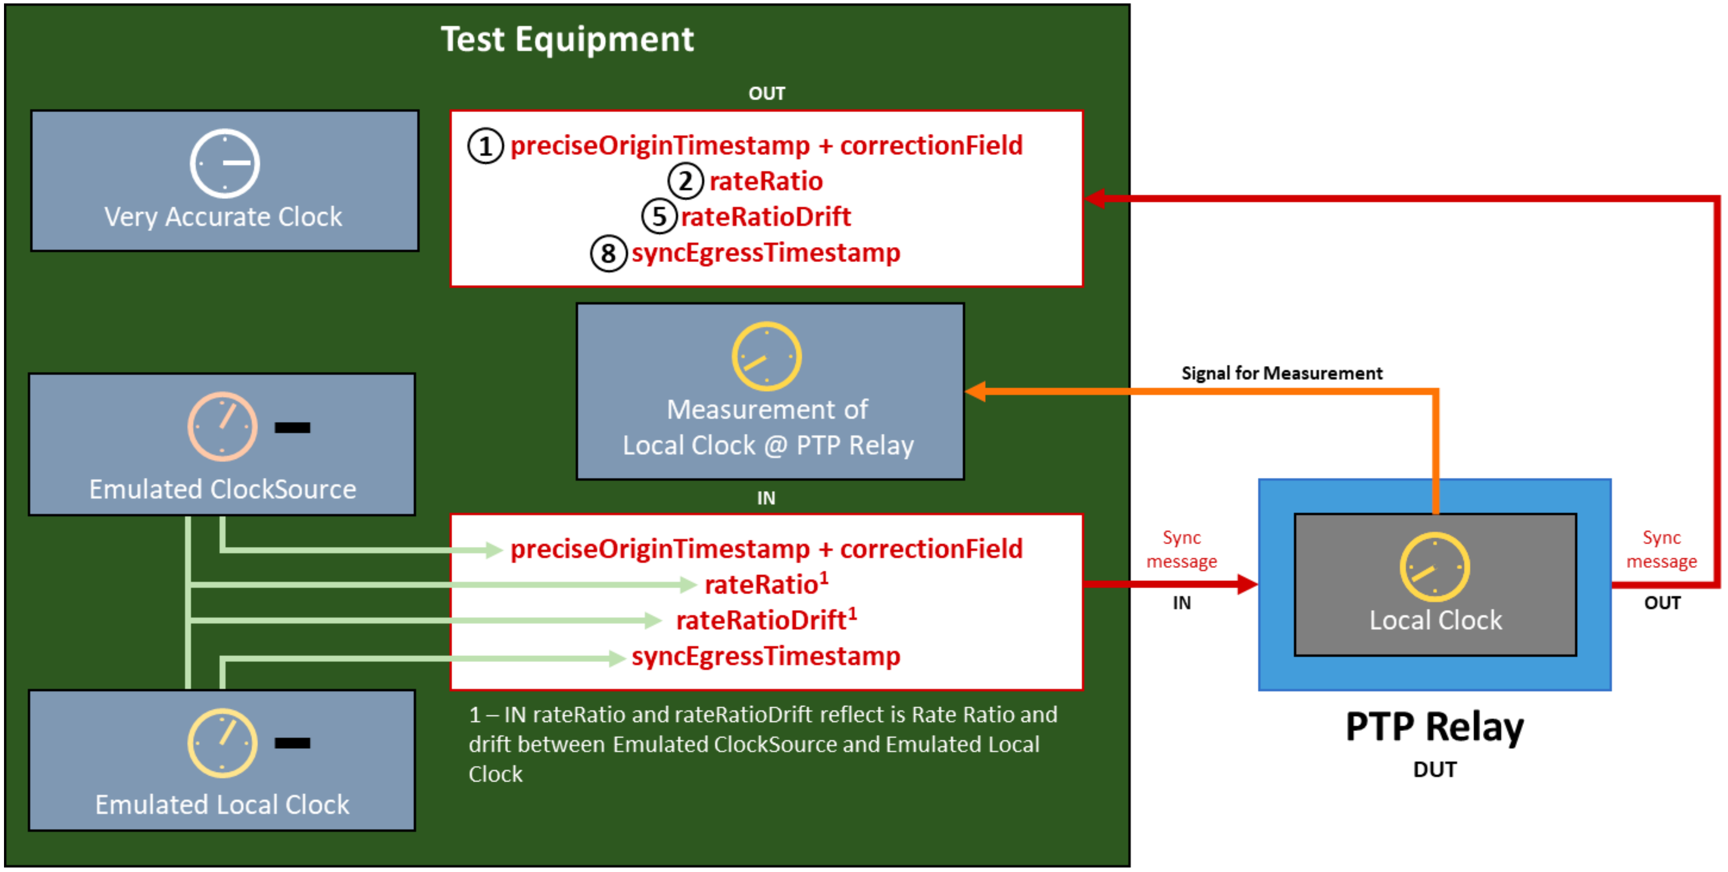

Description of Test Setup and Model for Simulator - 3

- Figure D.5 of [1] and [2] – possible test setup for PTP End Instance

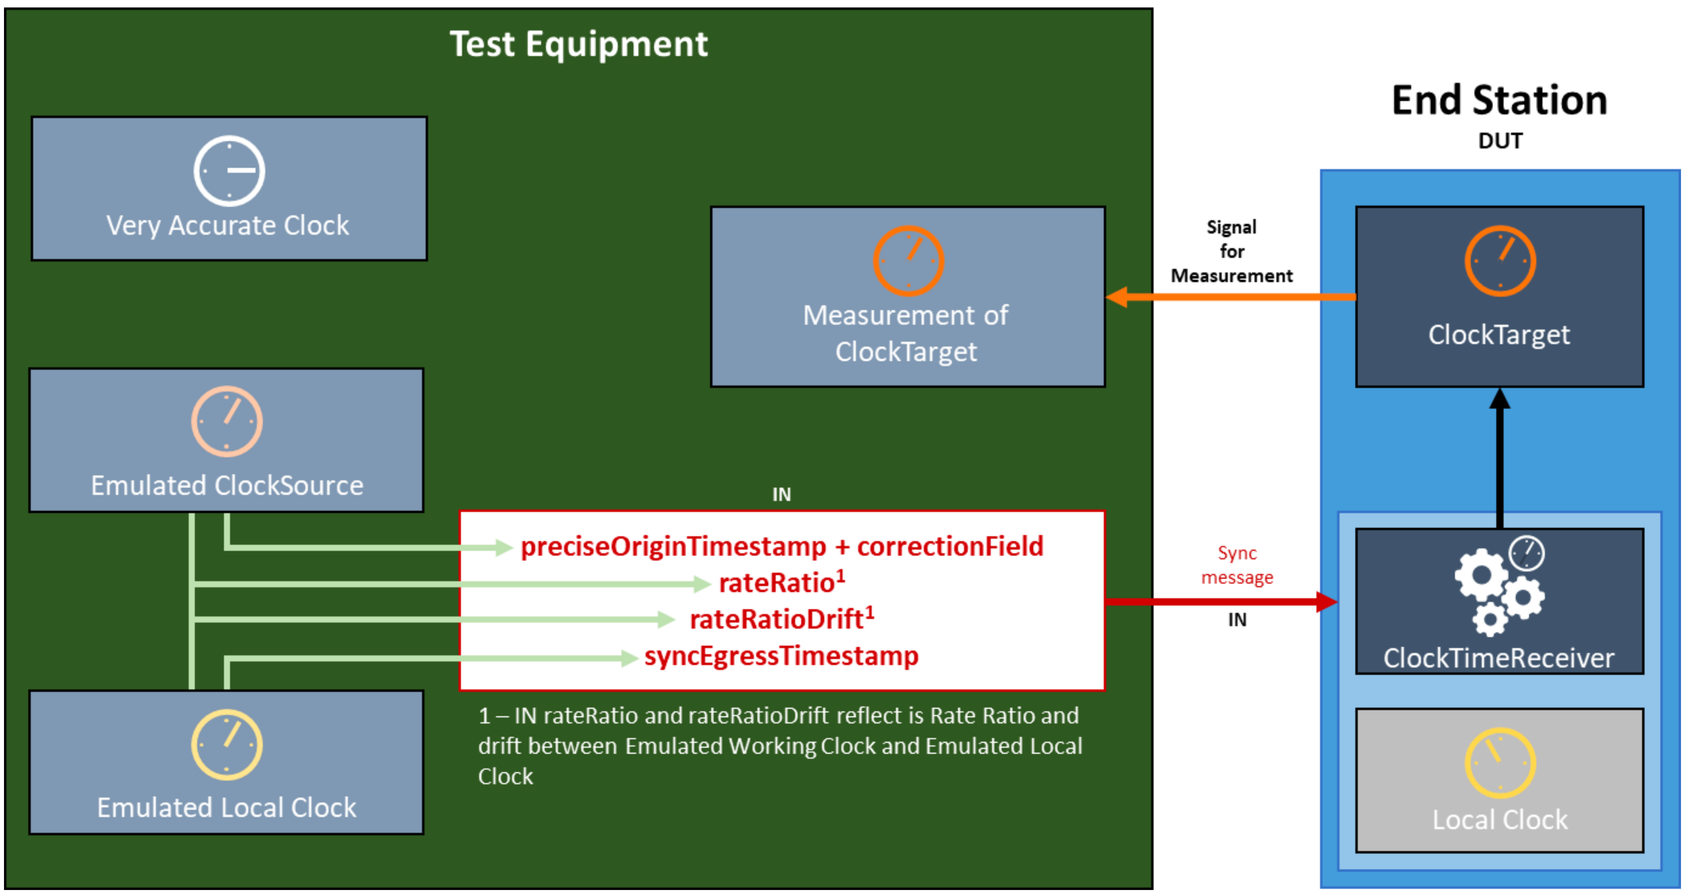

Description of Test Setup and Model for Simulator - 4

- The above means that the time series simulation model has four nodes:
  - Node 1 – Emulated ClockSource
  - Node 2 – Emulated LocalClock
  - Node 3 – DUT
  - Node 4 – Sink node that receives messages sent by the DUT
- A schematic of the simulator model is shown on the next slide
  - For convenience, these nodes are referred to by node number in the remainder of this presentation
- For the PTP Relay Instance tests, node 4 is not used; the tests only need values of fields of the Sync message (the time series simulator does not model two-step behavior explicitly) sent by the DUT
- For the PTP End Instance only filtered dTE is needed (see Table 13 of [1]), and this is the filtered dTE of the DUT (node 3)
  - Since both filtered and unfiltered dTE results are produced at each node, the same simulation runs can cover the PTP Relay Instance and PTP End Instance cases

Description of Test Setup and Model for Simulator - 5

Schematic of Time Domain Simulation Model

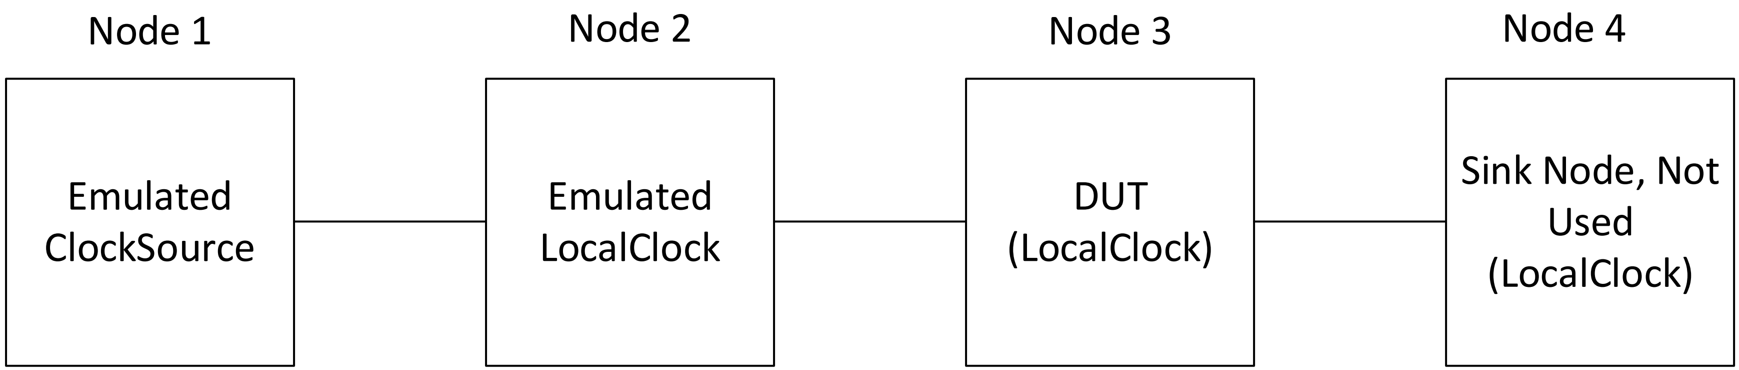

Assumptions Common to All Simulation Cases - 1

- The timestamp granularity is assumed to be 8 ns, based on a 125 MHz clock
  - The timestamp is truncated to the next lower multiple of 8 ns
  - This error is present only at node 3 (DUT); it is zero at other nodes
- The dynamic timestamp error is assumed to be uniformly distributed over [-6 ns, +6 ns]
  - This error is present only at node 3; it is zero at other nodes
- Pdelay Interval
  - Pdelay is used only to compute meanLinkDelay, and not neighborRateRatio (NRR)
  - NRR is computed using successive Sync message (using the syncEgressTimestamp)
  - The nominal Pdelay interval is 125 ms
  - The actual Pdelay interval is assumed to be uniformly distributed in the range [(0.9)(125 ms), (1.3)(125 ms)] = [112.5 ms, 162.5 ms]

Assumptions Common to All Simulation Cases - 2

Sync Interval

- The Sync interval is assumed to be uniformly distributed in the range [119 ms, 131 ms]

Residence time

- The residence time is assumed to be a truncated normal distribution with mean of 5 ms and standard deviation of 1.8 ms, truncated at 1 ms and 15 ms
- Probability mass greater than 15 ms and less than 1 ms is assumed to be concentrated at 15 ms and 1 ms, respectively (i.e., truncated values are converted to 15 ms or 1 ms, respectively)
- Residence time is present at node 3; cases are run with and without residence time present at node 2
- Residence time at node 4 is irrelevant, and is set to zero

Assumptions Common to All Simulation Cases - 3

❑Pdelay Turnaround Time

- The Pdelay turnaround time is assumed to be a truncated normal distribution with mean of 10 ms and standard deviation of 1.8 ms, truncated at 1 ms and 15 ms
- Probability mass greater than 15 ms and less than 1 ms is assumed to be concentrated at 15 ms and 1 ms, respectively (i.e., truncated values are converted to 15 ms or 1 ms, respectively)
- Pdelay turnaround time is present at node 3; cases are run with and without Pdelay turnaround time present at nodes 1 and 2
- Pdelay turnaround time at node 4 is irrelevant, and is set to zero

Assumptions Common to All Simulation Cases - 4

- Link Delay
  - Link delay is assumed to be uniformly distributed between 5 ns and 500 ns
  - Link delays are generated randomly at initialization and kept at those values for the entire simulation
  - Link asymmetry is not modeled
  - For the single replication simulation cases here, link delay is
    - 36.865 ns for the link between nodes 1 and 2
    - 454.21 ns for the link between nodes 2 and 3
    - Irrelevant and set to zero for the link between nodes 3 and 4

Assumptions Common to All Simulation Cases - 5

❑Mean Link Delay Averaging

- The averaging function is assumed to be an IIR filter that uses 0.999 of the previously computed value and 0.001 of the most recent measurement
- This is equivalent to the filter of the NOTE of B.4 of 802.1AS-2020, taken as a first-order filter, i.e.,

$$y_k = a_1 y_{k-1} + b_0 x_k$$

- where y_k is the k^{th} filter output, x_k is the k^{th} measurement, $a_1 = 0.999$, and $b_0 = 0.001$

Assumptions Common to All Simulation Cases - 6

❑Clock drift

- In all cases, node 3 (DUT) is assumed to have a stable LocalClock, i.e., its clock drift is zero
- Clock drift is present in nodes 1 and 2 in some of the cases (the specific cases are described in detail later)
- Clock drift, when present, is as follows:
  - The clock frequency is stable at -100 ppm (relative to nominal) from initialization to 1000 s
  - The clock frequency then drifts from -100 ppm to +100ppm at a rate of 1 ppm/s (i.e., over 200 s), to time 1200 s
  - The clock frequency is constant at +100 ppm from 1200 s until the end of the simulation time
    - Note: In initial discussions, it was indicated that the clock drift of 1 ppm/s would begin after 100 s; however, it was found that more time was needed for the initial transient, due to starting the simulation with non-zero frequency offset, to decay away. An initialization time of 1000 s was chosen as a conservative value after which the transient has decayed.

Assumptions Common to All Simulation Cases - 7

- Drift tracking and compensation algorithms are used in all cases at nodes 1, 2, and 3
  - The algorithms are irrelevant at node 4
- All the simulation cases here use (see [1] – [4], and references cited in those presentations, for details):
  - mNRRcompNAP = 8; mNRRsmoothingNA = 4

Endpoint filter (PLL) Parameters - 1

- Simulation results presented in [4] indicated that the endpoint filter needs to have a 3dB bandwidth in the range 0.7 Hz to 1 Hz and a maximum gain peaking of 2.1985 dB
  - The corresponding undamped natural frequency is 3.1011 rad/s
- The PLL model used in the simulator is second-order and linear, with 20 dB/decade roll-off
  - It is based on a discretization that uses an analytically exact integrating factor to integrate the second-order system
  - As a result, the PLL model in the simulator is stable regardless of the time step, i.e., sampling rate (though aliasing of the input or noise is possible)
  - Details are given in Appendix VIII.2.2 of [5] (except that the relation between gain peaking and damping ratio is based on the exact result in 8.2.3 of [6] (see Eqs. (8-13 – 8-15 there)

Endpoint filter (PLL) Parameters - 2

PLL noise generation (Cont.)

- PLL noise generation is modeled as described in [4], i.e., using the same local oscillator phase variation model used for the LocalClock
- The noise is computed by passing the XO phase noise through a high-pass filter with the same 3dB bandwidth and damping ratio as the low-pass PLL filter, and adding the result to the PLL output that was computed from the input
- However, the PLL is associated with the PTP End Instance, i.e., the DUT, and it was indicated in the assumptions related to clock drift that the DUT is assumed to have a stable LocalClock, i.e., its clock drift is assumed to be zero
- **This means that the PTP End Instance endpoint filter is assumed to have zero noise generation**

Endpoint filter (PLL) Parameters - 3

❑PLL noise generation (Cont.)

- Since the DUT drift is emulated by applying drift to the Emulated LocalClock, this is equivalent to assuming that the actual clock drift of the DUT is zero
- **At the very least, this means the actual tests should be done at constant temperature**
- This means that the actual DUT noise generation when the test is performed is the noise generation at constant temperature
- **The underlying assumption is that the DUT noise generation at constant temperature is negligible compared to the PTP End Instance requirement on filtered dTE**
- Note that in the time series simulations for a chain of PTP Instances [4], temperature was assumed to be varying according to the given temperature profile

Summary of Simulation Cases - 1

Case	Mean Link Delay Averaging Used	Residence time set to zero at node 2	Pdelay turnaround time set to zero at nodes 1 and 2	Clock drift present at node 1	Clock drift present at node 2	Simulation time (s)
1	yes	no	no	no	no	100000
1N	yes	yes	yes	no	no	100000
1A	no	no	no	no	no	100000
1AN	no	yes	yes	no	no	100000
2	yes	no	no	yes	no	2000
2N	yes	yes	yes	yes	no	2000
3	yes	no	no	yes	yes	2000
3N	yes	yes	yes	yes	yes	2000

- In the above case numbering convention:
  - N means that residence time and Pdelay turnaround time are set to zero
  - A means that link delay averaging is not used (these cases were run so that the behavior of link delay averaging could be seen more easily)

Summary of Simulation Cases – 2

- Note that cases 1, 1N, 1A, and 1AN have no clock drift
- These cases had a much longer simulation time than the cases with clock drift (100000 s vs 2000 s) to help understand the relatively long-term variations in measured mean link delay (to be described shortly)

Results for Measured Mean Link Delay - 1

- Results are presented for measured mean link delay because the variation in measured mean link delay shows impacts subsequent results for filtered and unfiltered dTE and PreciseOriginTimestamp+CorrectionField-(Working Clock at GM)
- The impact is not large, and is visible mainly in the results for cases without clock drift (because the effect of clock drift is much larger)
  - The mean link delay results are shown mainly for explanation
- The following two slides show results for measured mean link delay for the link between node 2 (emulated LocalClock) and node 3 (DUT), for cases with and without link delay averaging (cases 1 and 1A, respectively), for simulation time of 100,000 s
- The two slides after that show results for cases 1 and 1A, but with only dynamic timestamp error, i.e., timestamp granularity is set to zero
- The two slides after that show the detail of the first 5000 s for the case 1 and 1A results, with both dynamic timestamp error and timestamp granularity

Measured Mean Link Delay – Case 1

Case 1
Measured mean link delay by node 3, for link between nodes 2 and 3

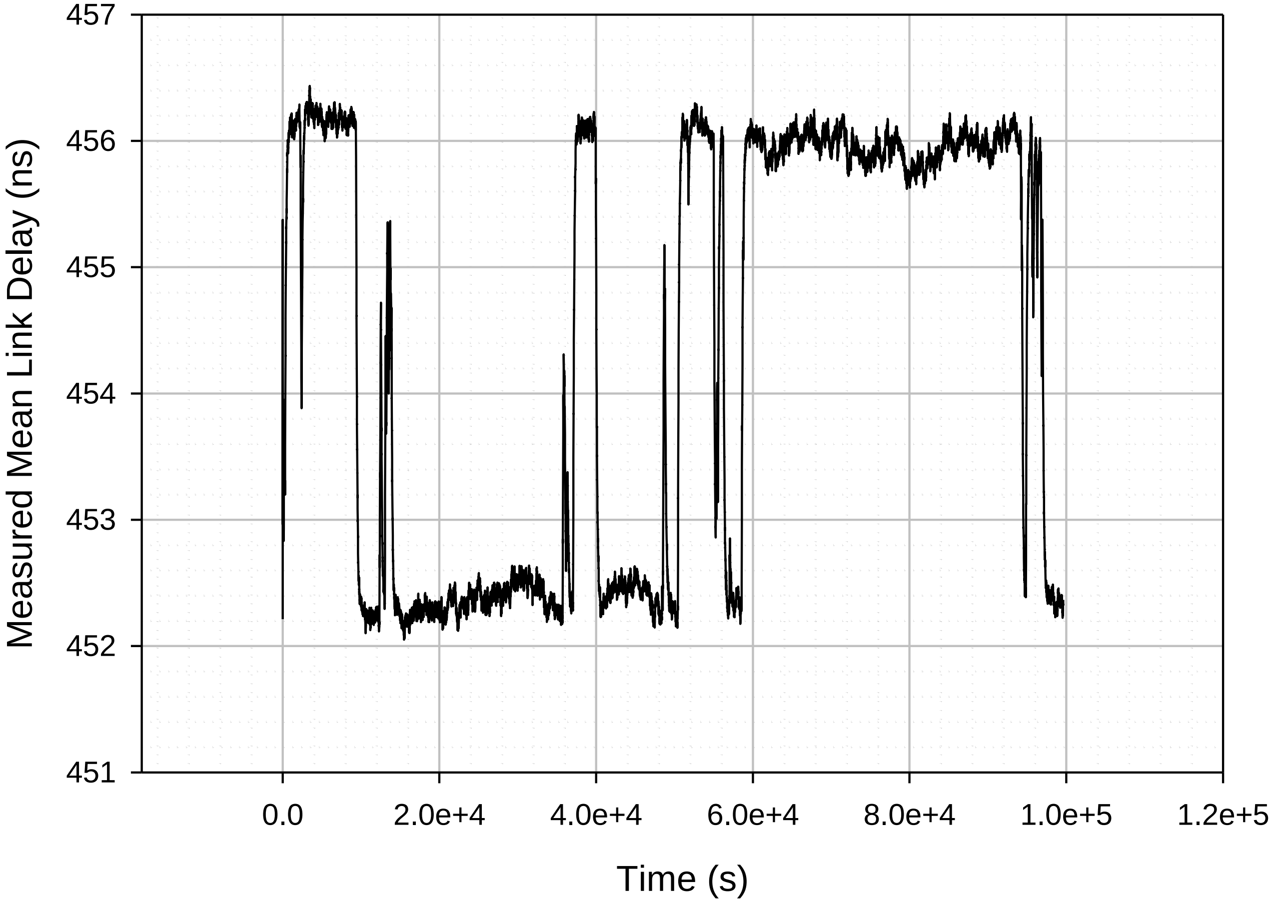

Measured Mean Link Delay – Case 1A

Case 1A
Measured mean link delay by node 3, for link between nodes 2 and 3

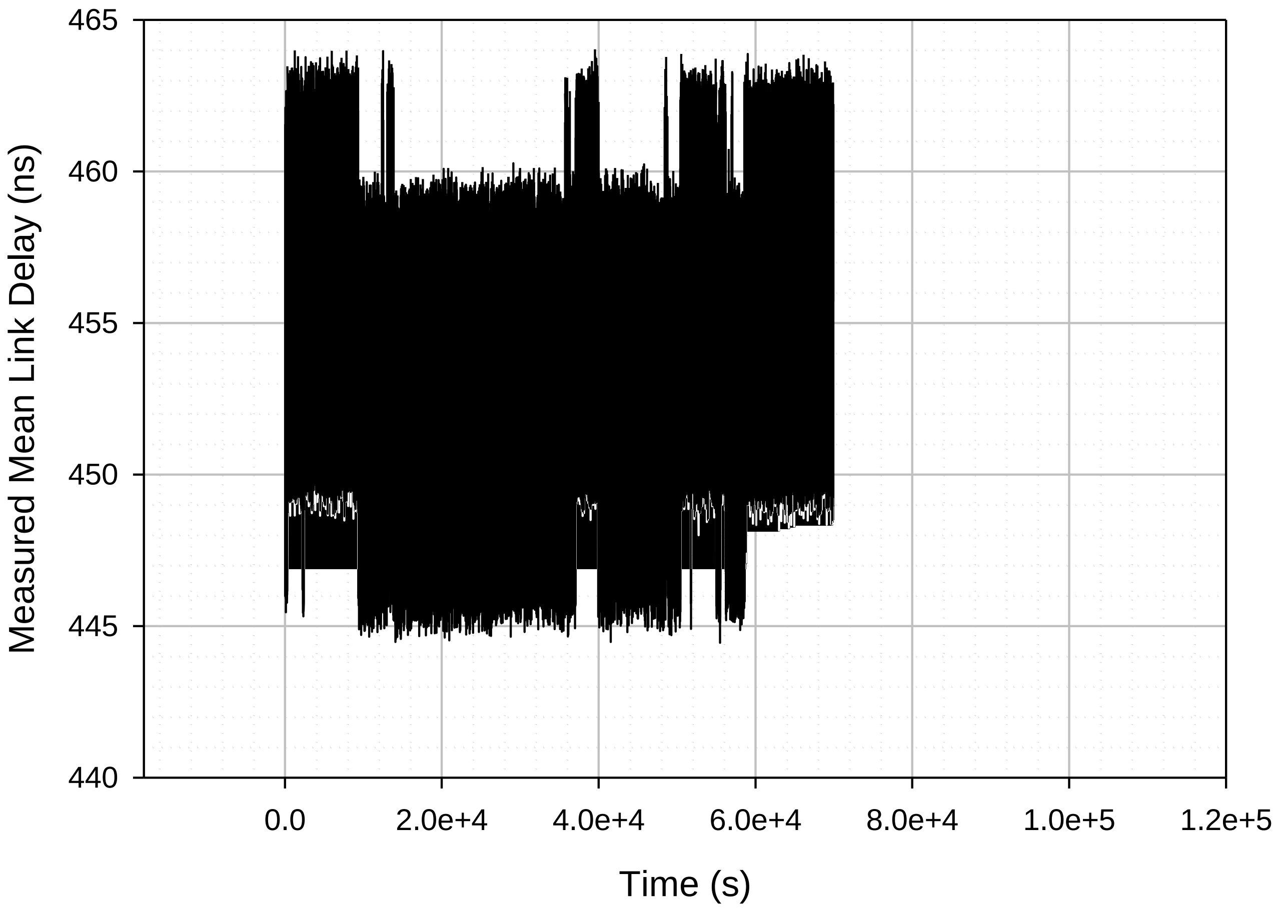

Measured Mean Link Delay – Case 1, No TS Gran

Case 1
Measured mean link delay by node 3, for link between nodes 2 and 3
Timestamp granularity at node 3 set to zero

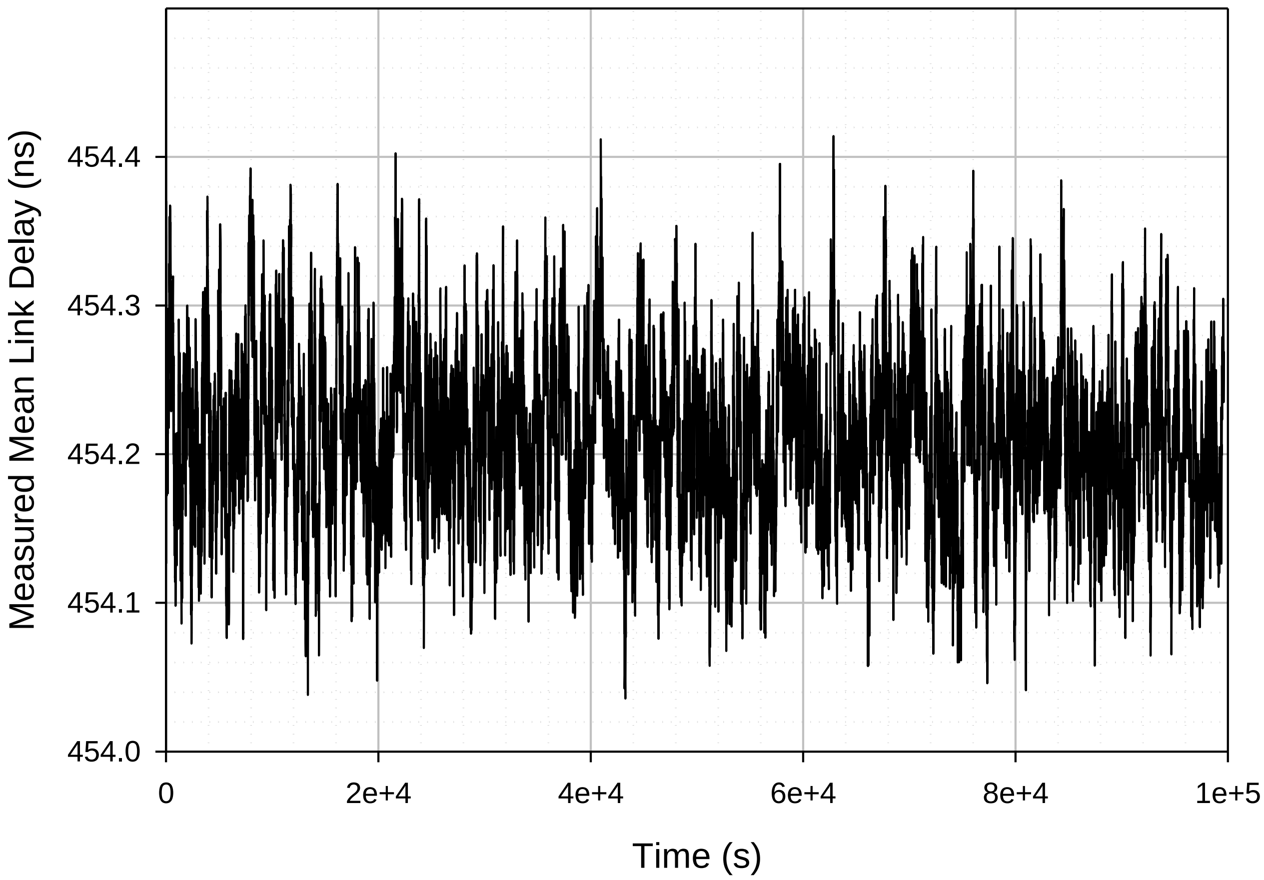

Measured Mean Link Delay – Case 1A, No TS Gran

Case 1A
Measured mean link delay by node 3, for link between nodes 2 and 3
Timestamp granularity at node 3 set to zero

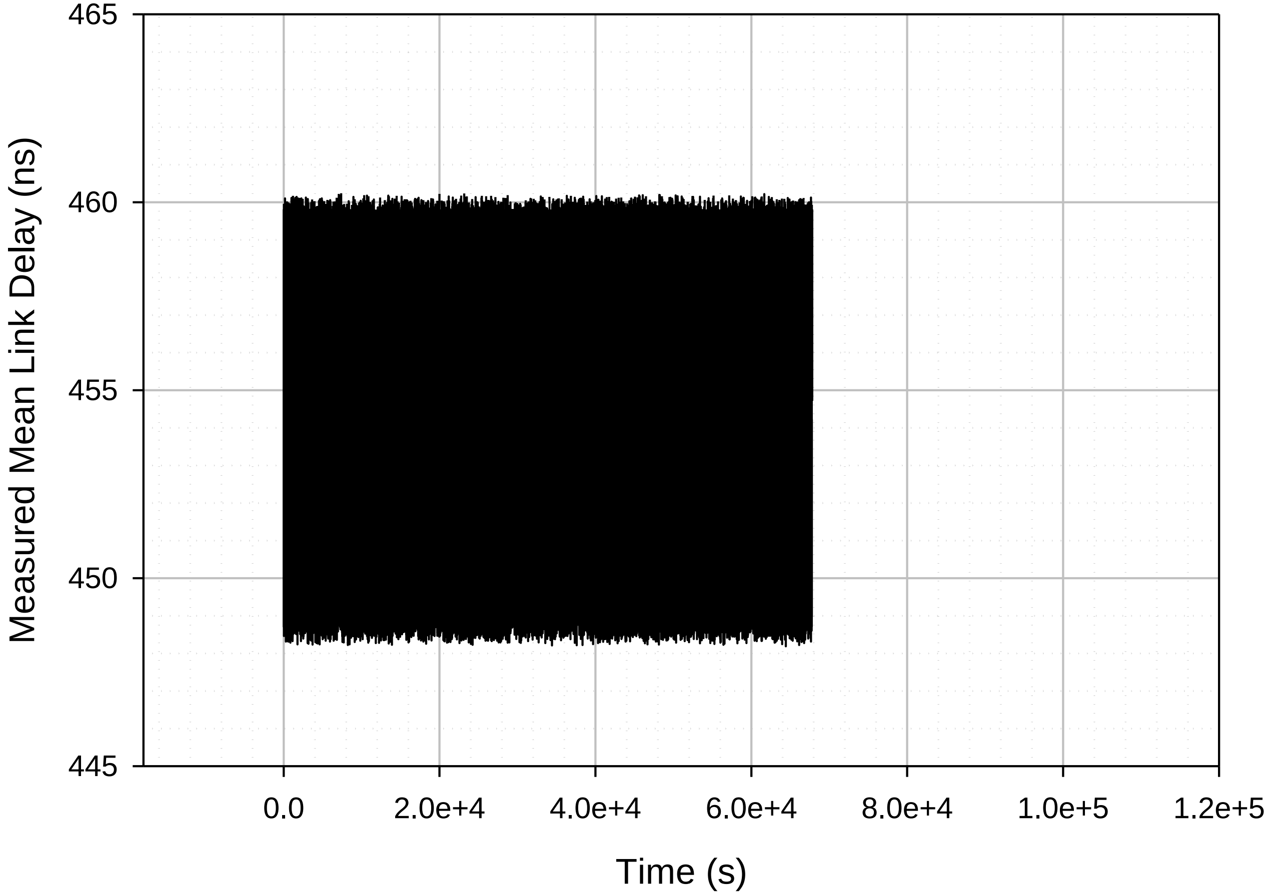

Measured Mean Link Delay - Case 1 (Detail of 0 - 5000 s)

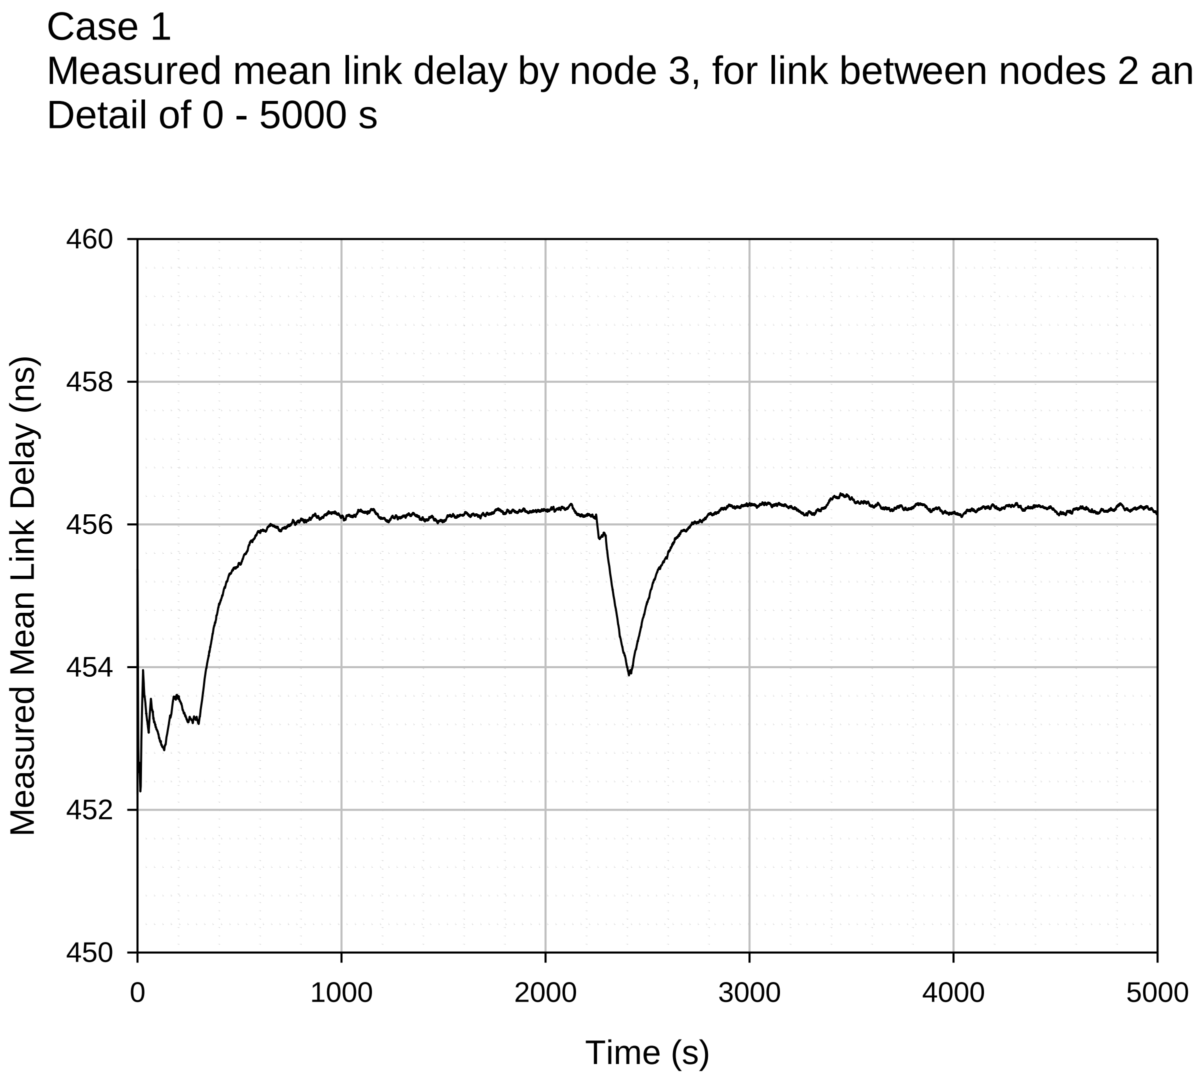

Measured Mean Link Delay - Case 1A (Detail of 0 - 5000 s)

Case 1A
Measured mean link delay by node 3, for link between nodes 2 and 3
Detail of 0 - 5000 s

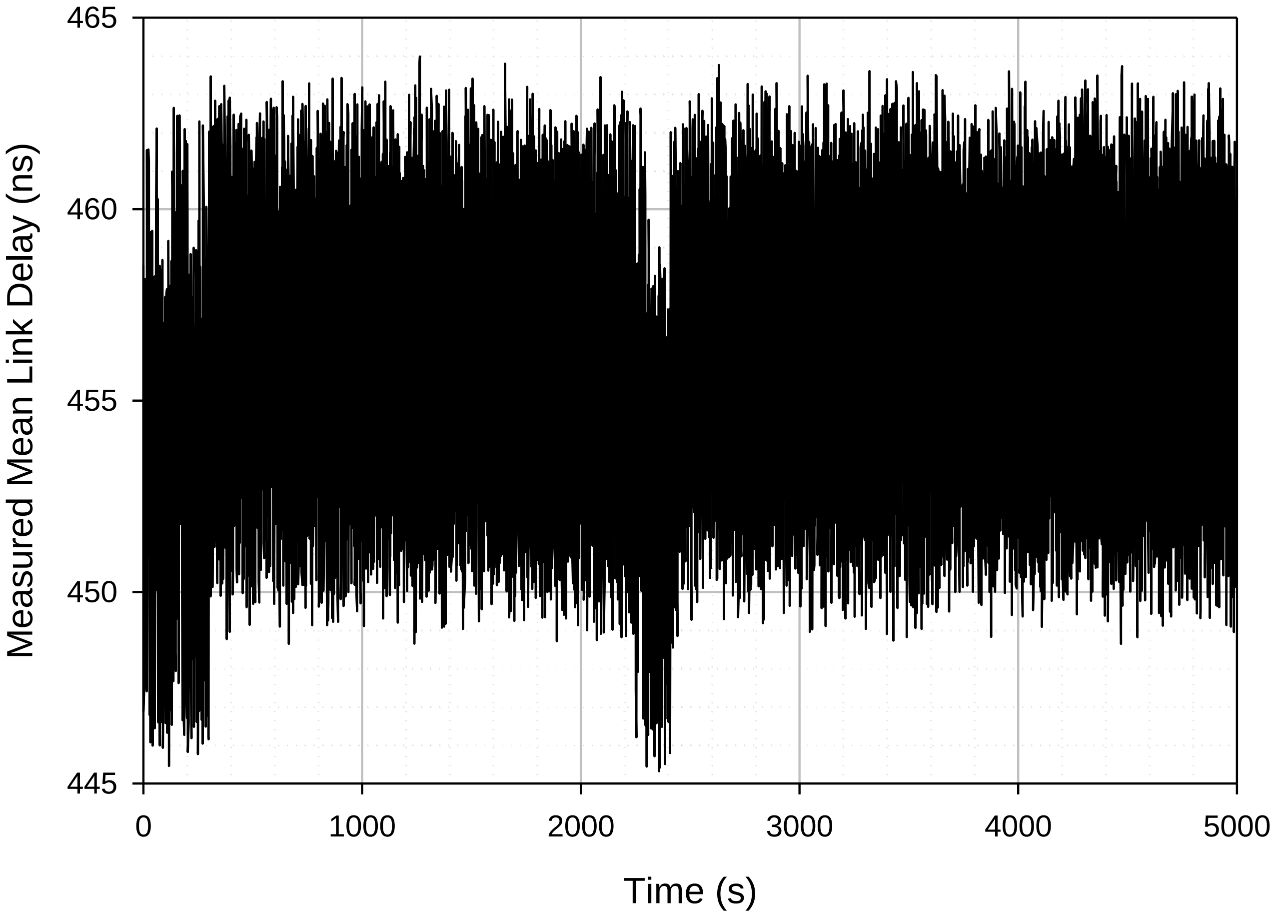

Results for Measured Mean Link Delay - 2

- The results without link delay averaging and with timestamp granularity set to zero show a fast variation whose peak-to-peak magnitude is approximately 12 ns, due to the [-6 ns, +6 ns] dynamic timestamp error
- The results with link delay averaging, but with timestamp granularity set to zero show that the peak-to-peak of 12 ns is reduced to approximately 0.37 ns by the averaging filter
- The addition of 8 ns timestamp granularity, but without link delay averaging, results in relatively infrequent jumps (i.e., over periods of thousands of seconds) of approximately 4 ns, plus some jumps that are less than 4 ns
- The inclusion of the averaging filter (and with both timestamp granularity and dynamic timestamp error) reduces the $\pm$6 ns fast variation; however, the infrequent jumps remain

Results for Measured Mean Link Delay - 2

- The infrequent jumps in measured link delay is caused by the truncation due to timestamp granularity
  - Note that there is timestamp granularity error and dynamic timestamp error only at the DUT, which means that there are sources of these errors in the peer delay computation on the DUT's upstream link, one truncation when timestamping the transmitted Pdelay_Req, and one truncation when timestamping the received Pdelay_Resp
  - Note also that, if 4 ns were added to these timestamps, the effect of that would cancel; the variation in these results is ***not*** due to not adding the 4 ns
  - If one Pdelay_Req message has a timestamp truncated by almost 8 ns, and the associated Pdelay_Resp message has truncation close to zero, the result is a 4 ns error in measured link delay (without averaging), due to the division by 2 in the computation
  - Similarly, if a Pdelay_Resp timestamp is truncated and the corresponding Pdelay_Req timestamp is not truncated, the result is a -4ns error in measured mean link delay
  - However, these truncations are relatively infrequent, and most of the time both truncations do not occur in the same peer delay exchange; this results in a shift of the pattern by approximately 4 ns

Results for Measured Mean Link Delay - 3

- The actual link delay is 454.21 ns; much of the variation is centered on this value, but sometimes it is shifted by -4 ns (and it is centered on this value for the cases where timestamp granularity is set to zero)
- Note that there is not a regular, low-frequency, beating pattern because the send and receive timestamps are not being taken at fixed frequencies that differ by a small amount; rather, Pdelay_Req is being sent at intervals that have some random variation (which results in truncation being different on successive peer delay message exchanges)
  - Also, there is additional randomness in case 1 due to the nonzero pdelay turnaround time at node 2

Results for Measured Mean Link Delay - 4

- The mean link delay averaging uses a window of 1000 peer delay exchanges
  - The Pdelay interval varies uniformly from 112.5 ms to 162.5 ms; the mean of this distribution is137.5 ms, which means that 1000 exchanges would occur over an interval of 137.5 s on average
  - This means that the time constant of the averaging filter is on the order of hundreds of seconds
  - The averaging filter removes the fast variation in the measurement, but requires a few hundred seconds to respond to an isolated jump in the pattern
  - In addition, if the pattern jumps and then jumps back before the filter has fully responded, the result is a smaller amplitude change in the pattern (as seen in the result with the averaging filter around 2400 s (in the 0 – 5000 s detail)

Results for Cases 1, 1N, 2, 2N, 3, 3N -- 1

- In what follows, results are shown for each of cases 1, 1N, 2, 2N, 3, and 3N (in that order)
- In each set, the following results (plots) are shown:
  1) Unfiltered dTE
  2) PreciseOriginTimeStamp+CorrectionField-(Working Clock at GM), when Sync is transmitted
  3) rateRatio field of Sync minus actual rate ratio, when Sync is sent
  4) rateRatioDrift field of Sync minus actual rate ratio drift, when Sync is sent
  5) Filtered dTE
- Items 2, 3, and 4 pertain to a PTP Relay Instance
- Item 5 pertains to a PTP End Instance
- Item 1 pertains to both a PTP Relay Instance and PTP End Instance
- Items 2, 3, 4, and 5 are the main interest relative to the requirements of Tables 12 and 13 of [1]

Results for Cases 1, 1N, 2, 2N, 3, 3N -- 2

- In some cases, plots of the full time history are dominated by a large startup transient
  - In these cases, an additional plot is shown without the startup transient (either starting after 50 s or showing a smaller range for the vertical axis)
  - In some later cases, only the details without the startup transient are shown
- In addition to the plots, mean, standard deviation, minimum, and maximum are computed for items 1 – 5 above and shown in tables
  - The computation is done after removing the first 500 s of each time history, to eliminate the effect of any startup transient
  - In the tables, items 2, 3, and 4 above are abbreviated M, N, and P, respectively

Case 1 - Unfiltered dTE

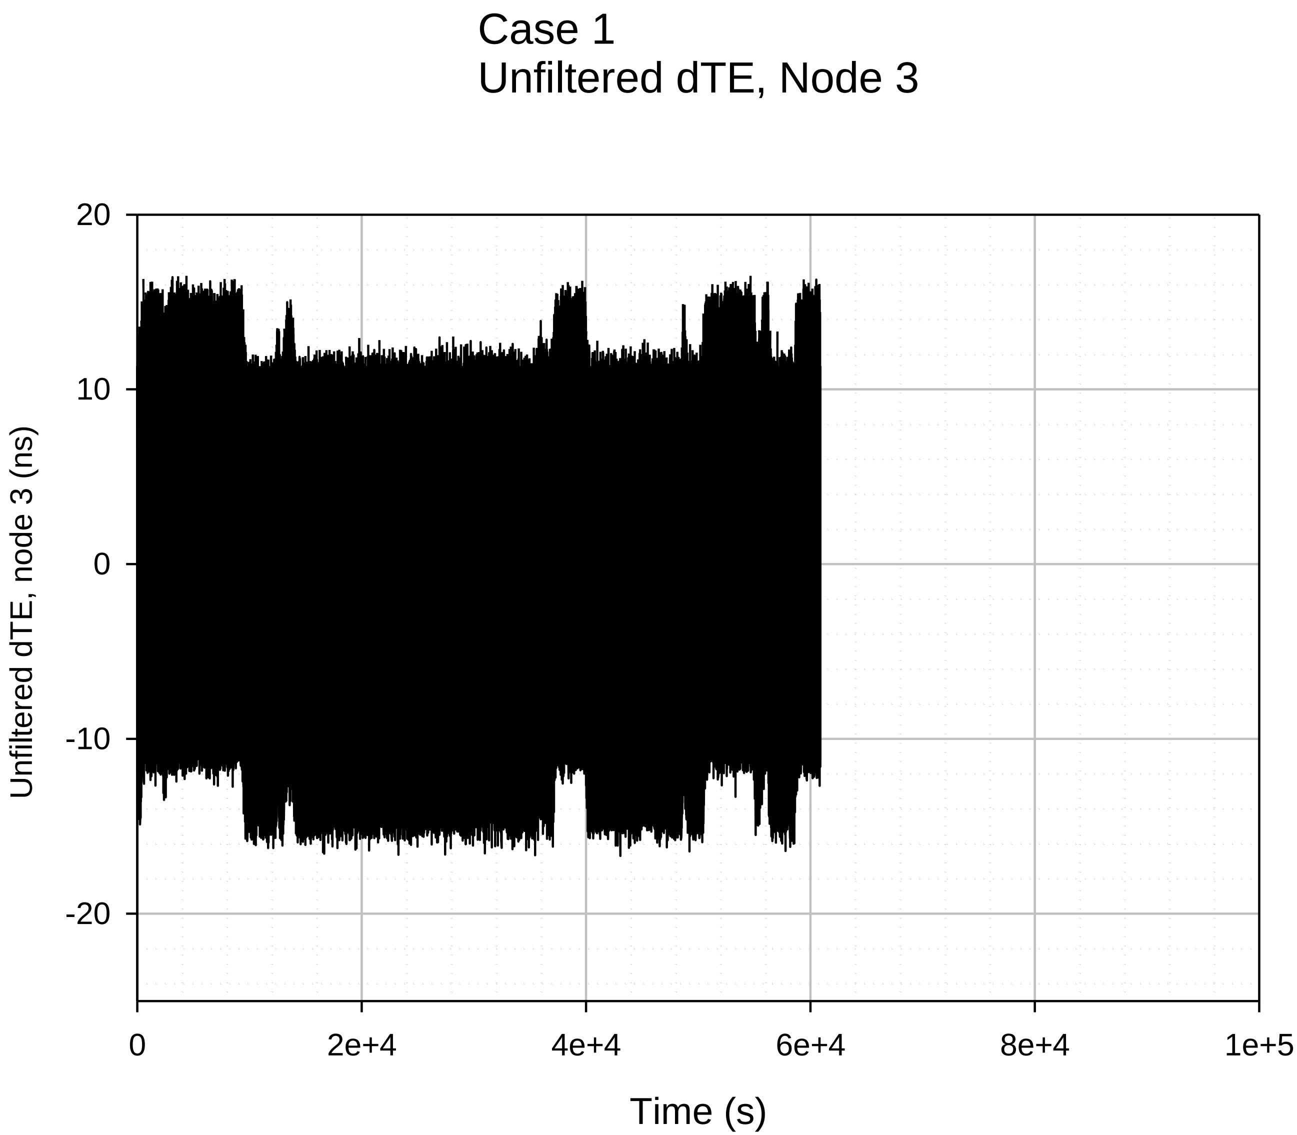

Case 1 - PreciseOriginTS+corrF-WorkingClock at GM

Case 1
preciseOriginTimestamp + correctionField - Working Clock at GM at transmission of Sync message
Node 3

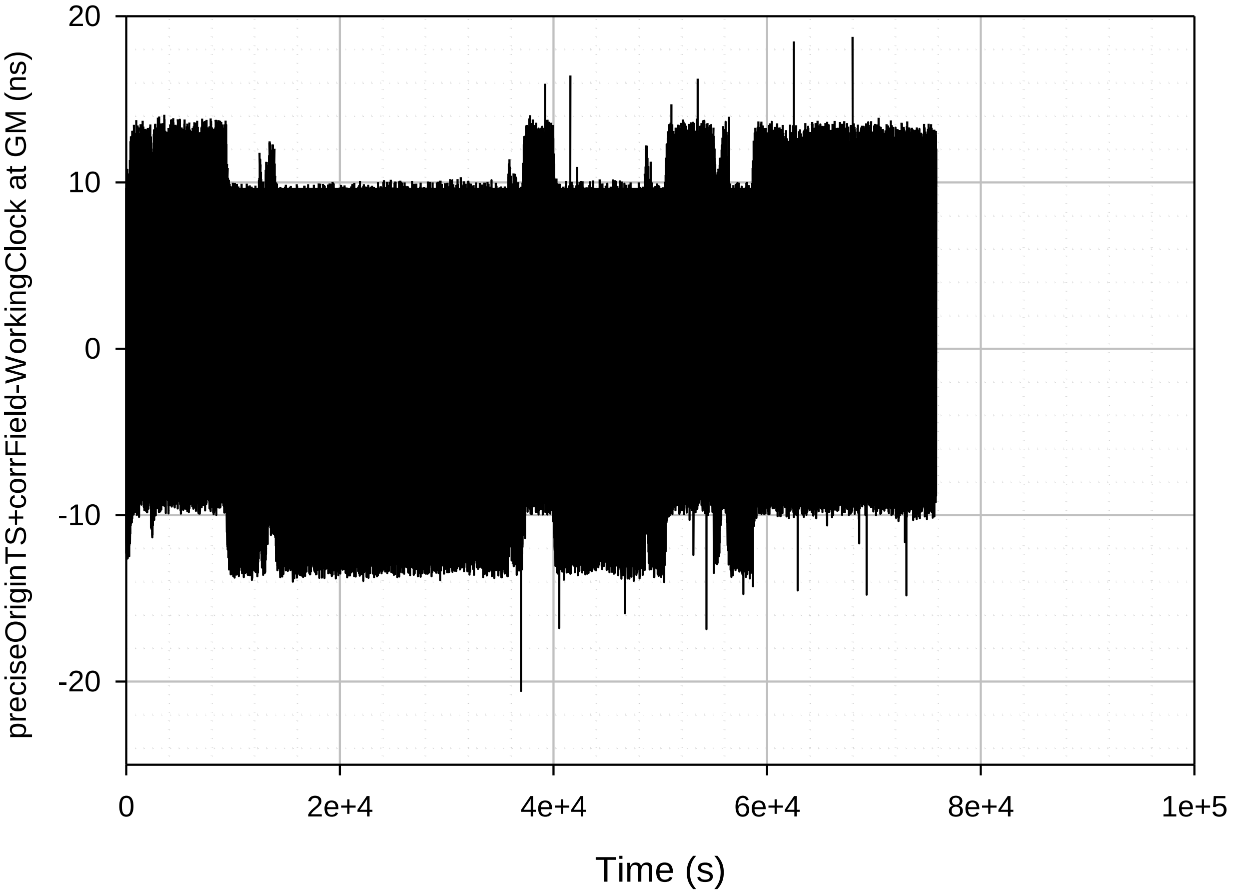

Case 1 - RR field of Sync Minus Actual RR

Case 1
rateRatio field of transmitted Sync message by DUT minus actual rate ratio of DUT working clock relative to Working Clock at GM at transmission of Sync message
Node 3

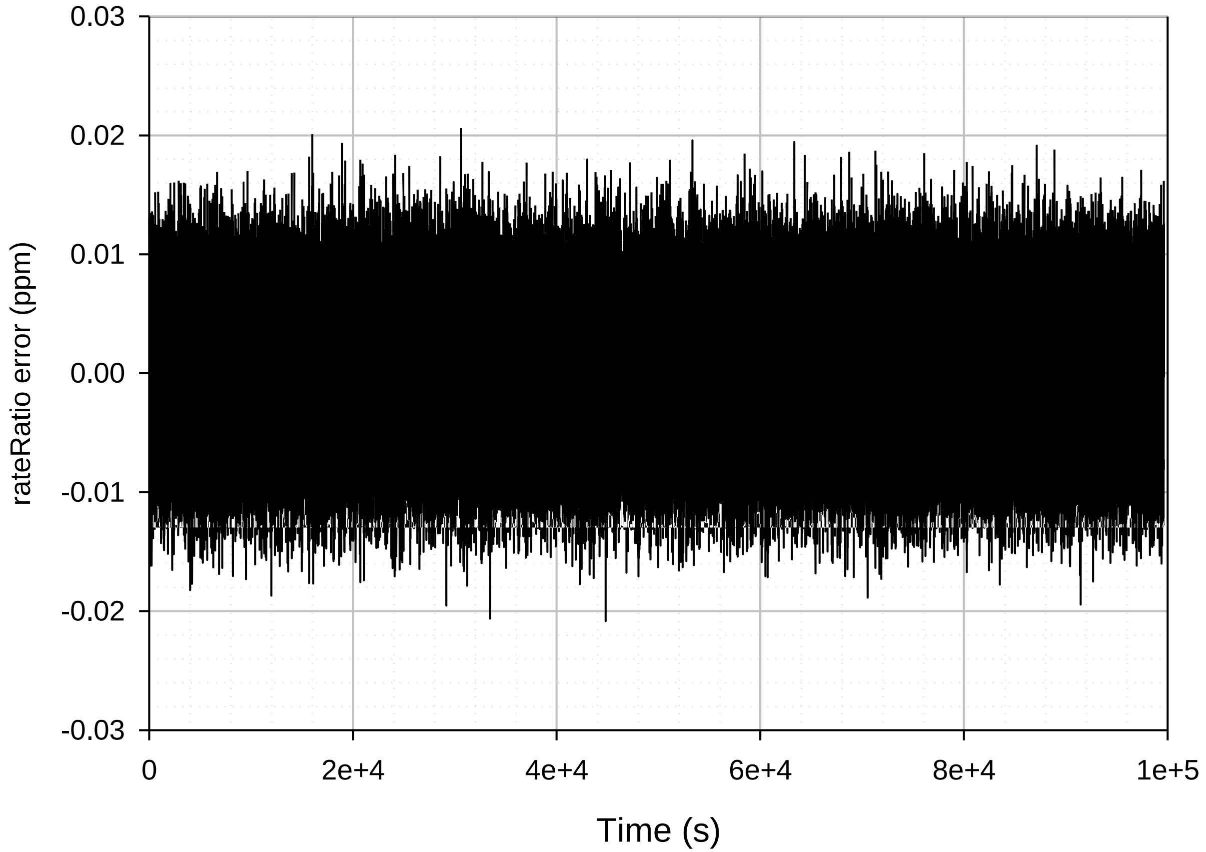

Case 1 – RR drift field of Sync Minus Actual RR Drift

Case 1
rateRatioDrift field of transmitted Sync message by DUT minus actual rate ratio drift rate of DUT working clock relative to Working Clock at GM at transmission of Sync message
Node 3

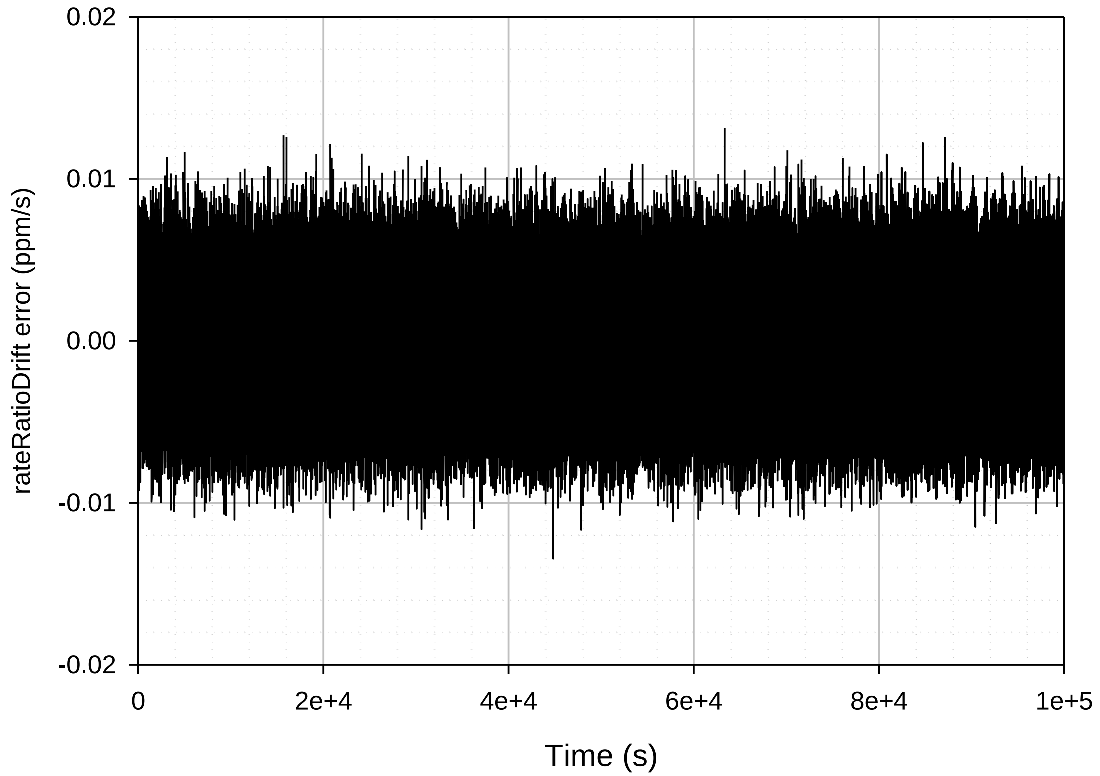

Case 1 - Filtered dTE (PTP End Instance)

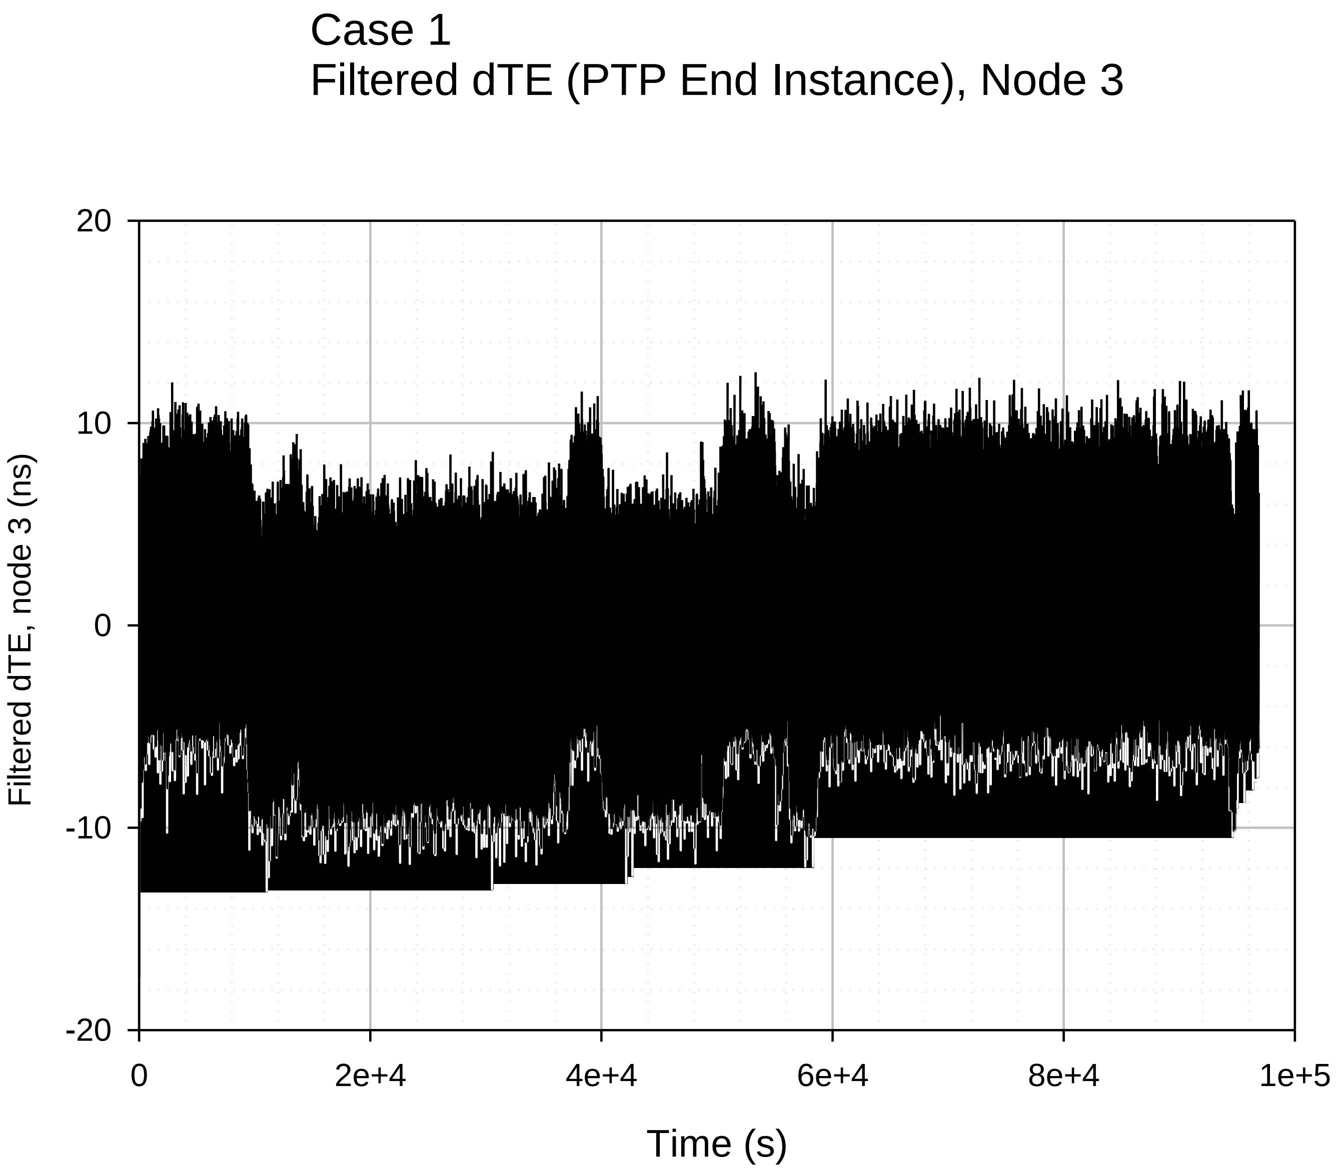

Case 1N - Unfiltered dTE

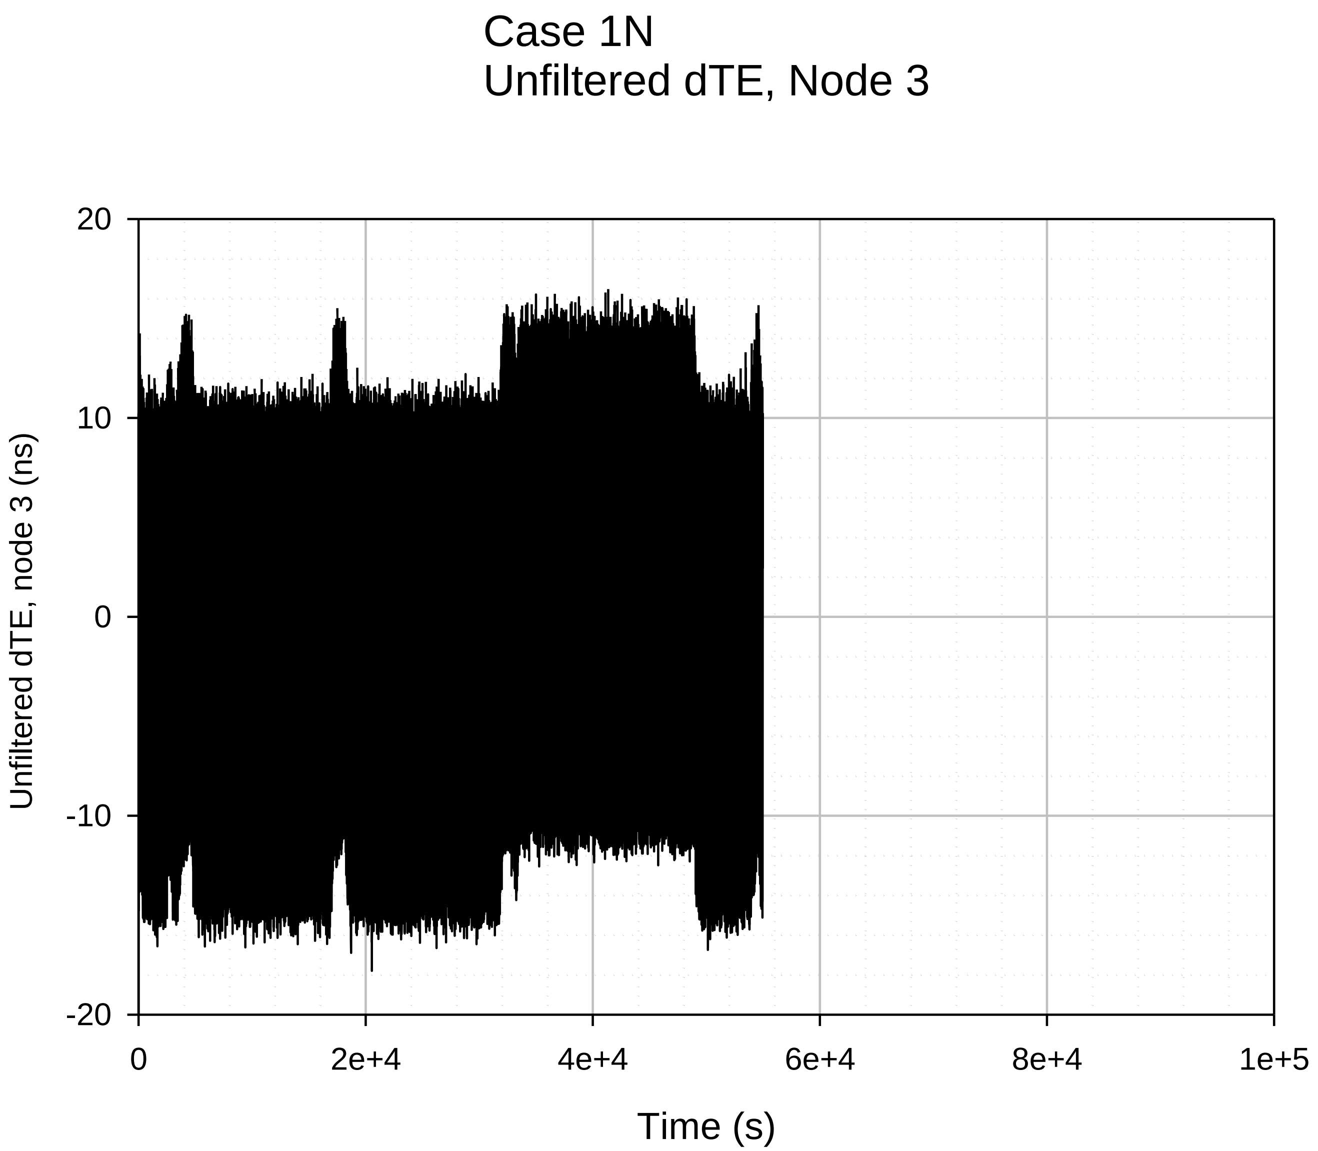

Case 1N - PreciseOriginTS+corrF-WorkingClock at GM

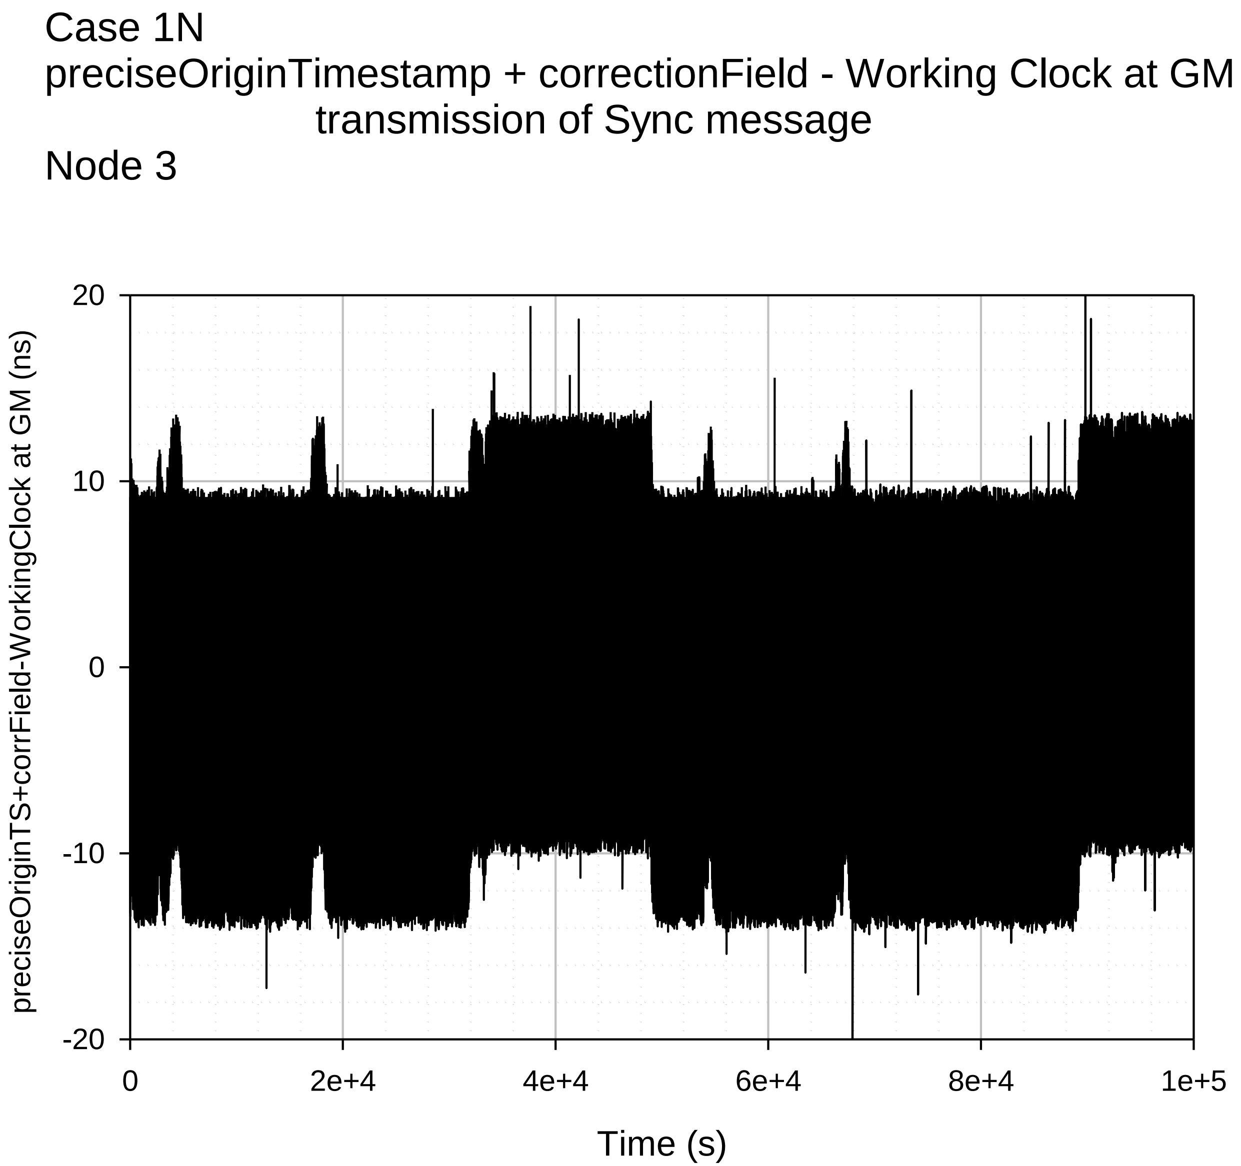

Case 1N - RR field of Sync Minus Actual RR

Case 1N
rateRatio field of transmitted Sync message by DUT minus actual rate ratio of DUT working clock relative to Working Clock at GM at transmission of Sync message
Node 3

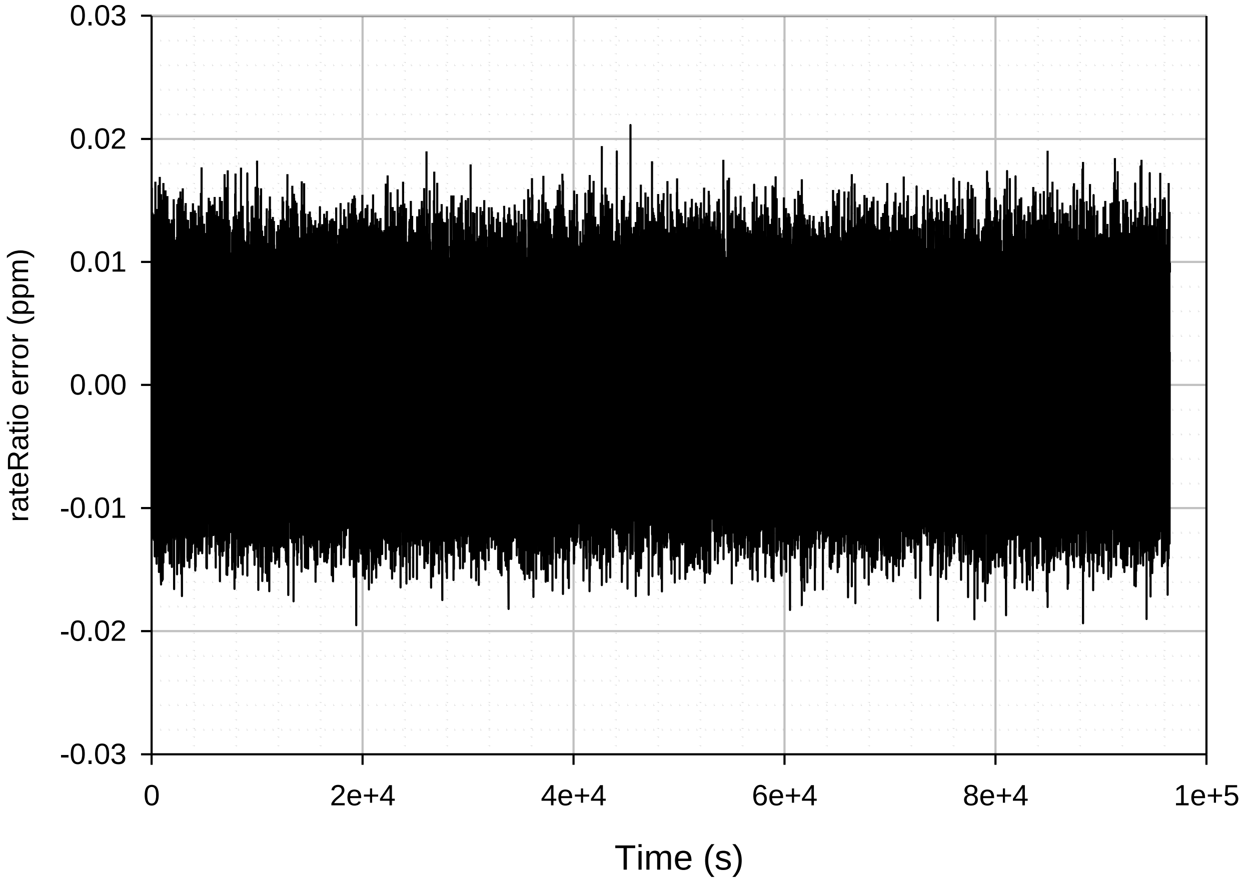

Case 1N – RR drift field of Sync Minus Actual RR Drift

Case 1N

rateRatioDrift field of transmitted Sync message by DUT minus actual rate ratio drift rate of DUT working clock relative to Working Clock at GM at transmission of Sync message

Node 3

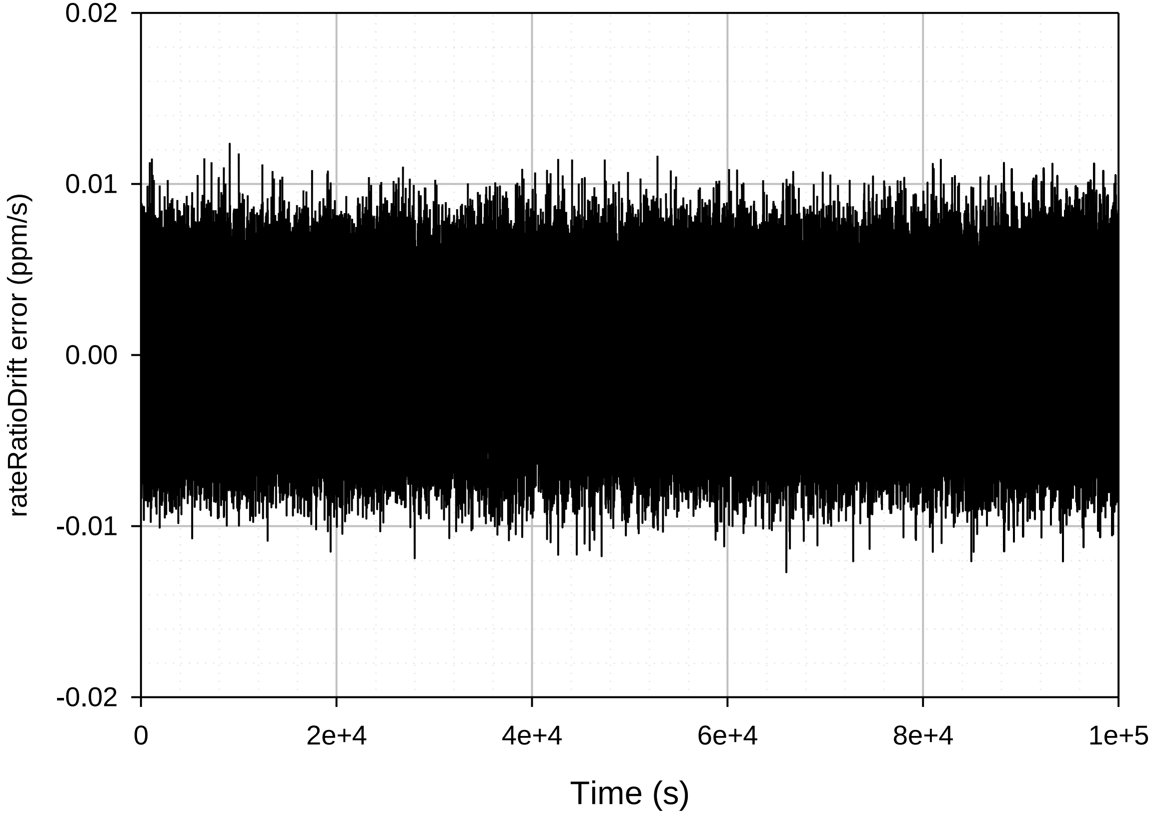

Case 1N - Filtered dTE (PTP End Instance)

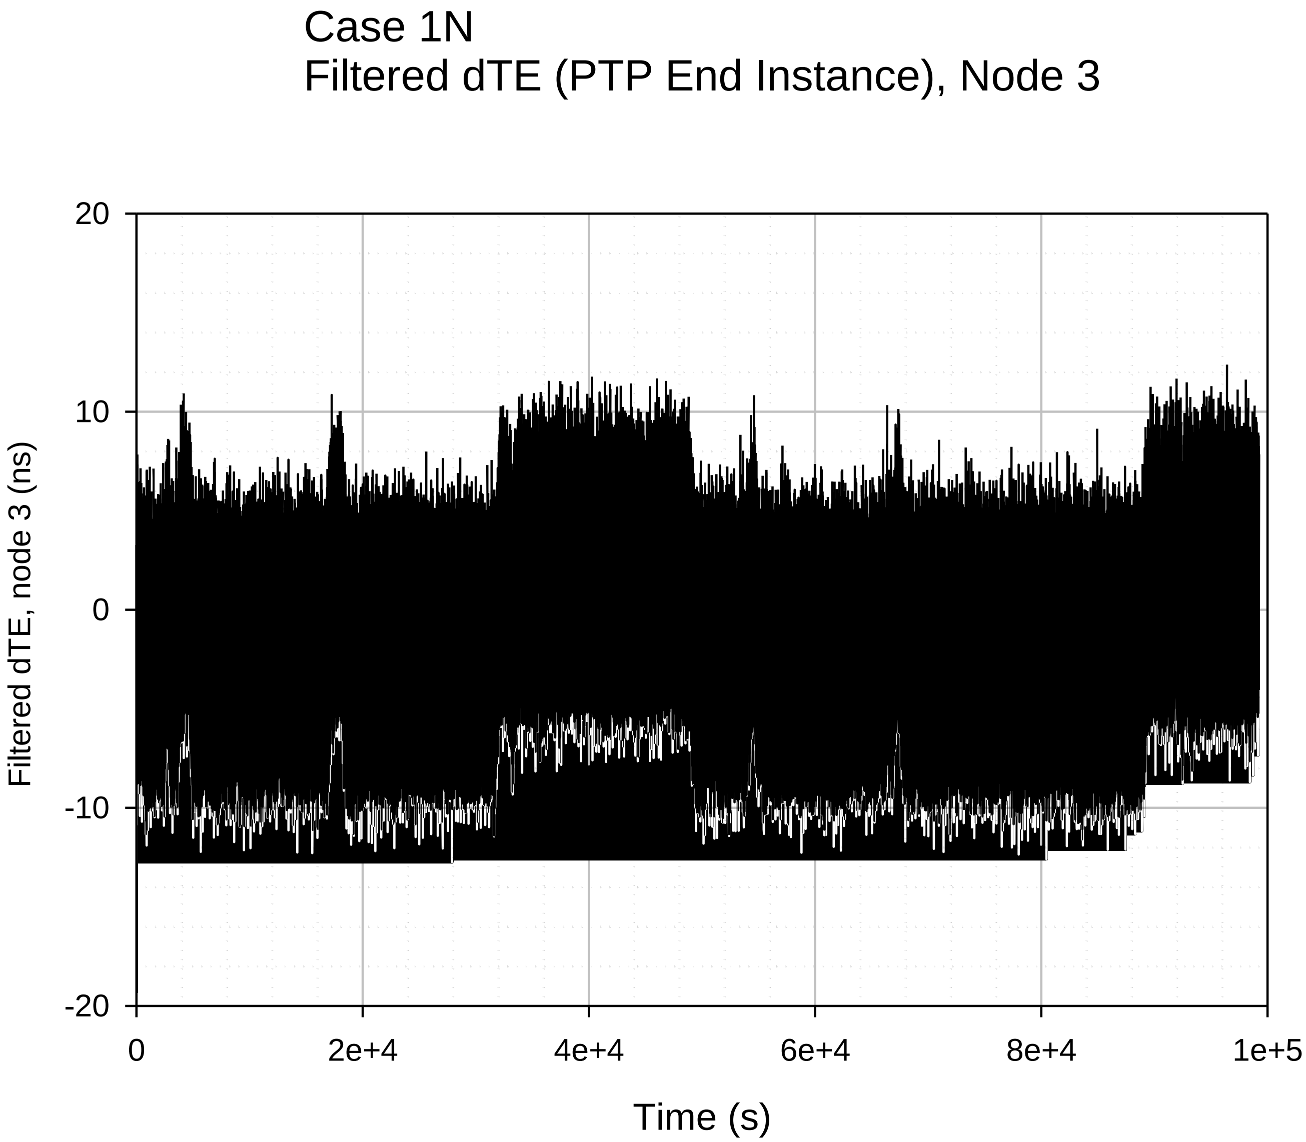

Numerical Results for Cases 1 and 1N

Case 1

Quantity	Mean	Standard Deviation	Minimum	Maximum
dTE_{unfilt} (ns)	0.195	5.142	-16.78	17.45
M (ns)	0.219	5.212	-20.57	18.75
N (ppm)	7.554×10^{-8}	4.501×10^{-3}	-0.0283	0.0262
P (ppm/s)	5.521×10^{-8}	2.777×10^{-3}	-0.0135	0.0131
dTE_{filt} (ns)	0.203	3.173	-13.14	12.56

Case 1N

Quantity	Mean	Standard Deviation	Minimum	Maximum
dTE_{unfilt} (ns)	-0.946	5.154	-17.84	17.27
M (ns)	-0.936	5.221	-20.43	20.11
N (ppm)	-1.436×10^{-8}	4.509×10^{-3}	-0.0195	0.0211
P (ppm/s)	-1.643×10^{-8}	2.781×10^{-3}	-0.0127	0.0124
dTE_{filt} (ns)	-0.931	3.202	-12.74	12.42

Cases 1 and 1N Results – Discussion - 1

- The case 1 and 1N (i.e., with non-zero residence time and Pdelay turnaround time at all nodes versus only at the DUT, respectively) are qualitatively very similar
  - Quantitatively, the minimum and maximum values and standard deviations for corresponding parameters are similar for both cases
  - The mean values for corresponding parameters differ mainly due to statistical variation (mean values are much smaller than minimum and maximum values in absolute value)
- Requirements of Tables 12 and 13 of [1] are met:
  - Mean of M is within range of -2 ns to +2 ns, and minimum and maximum values are within $\pm$20 ns range (row 1 of Table 12 of [1])
  - Mean of N is within range of -0.1 ppm to +0.1 ppm, and standard deviation is within $\pm$0.02 ppm (row 2 of Table 12 of [1])
  - Mean of P is within range of -0.1 ppm/s to +0.1 ppm/s, and standard deviation is within $\pm$0.02 ppm/s (row 5 of Table 12 of [1])
  - dTE_{filt} is within range of -15 ns to +15 ns (row 1 of Table 13 of [1])

Cases 1 and 1N Results – Discussion - 2

- The effect of the variation in mean link delay (described earlier) is evident in both filtered and unfiltered dTE, and in M
  - Note that results for unfiltered dTE and M are similar but not identical; M values are approximately equal to dTE values computed at times that Sync messages are transmitted by the DUT, while dTE values are also computed at times in between these transmit times
  - The results for minimum and maximum values, and standard deviation, for M, are in reasonable agreement with results for dTE for Test Type 1 on p.16 of [3]; however, the mean values are larger by several orders of magnitude. This is likely due to the variation in measured mean link delay error obtained here.
    - Related to this, note that means and standard deviations are computed here as time averages, but it has not been established the these have converged, or even that the processes are ergodic (i.e., roughly that time averages converge to ensemble averages). It would be desirable to at least run multiple replications of the simulations to consider this further.
  - Results for minimum and maximum values, and standard deviation for N and P are larger than corresponding results for rateRatio and rateRatioDrift given for Test Type 1 on p.16 of [3]; this also is likely due to variation in measured mean link delay error obtained here.

Case 2 - Unfiltered dTE

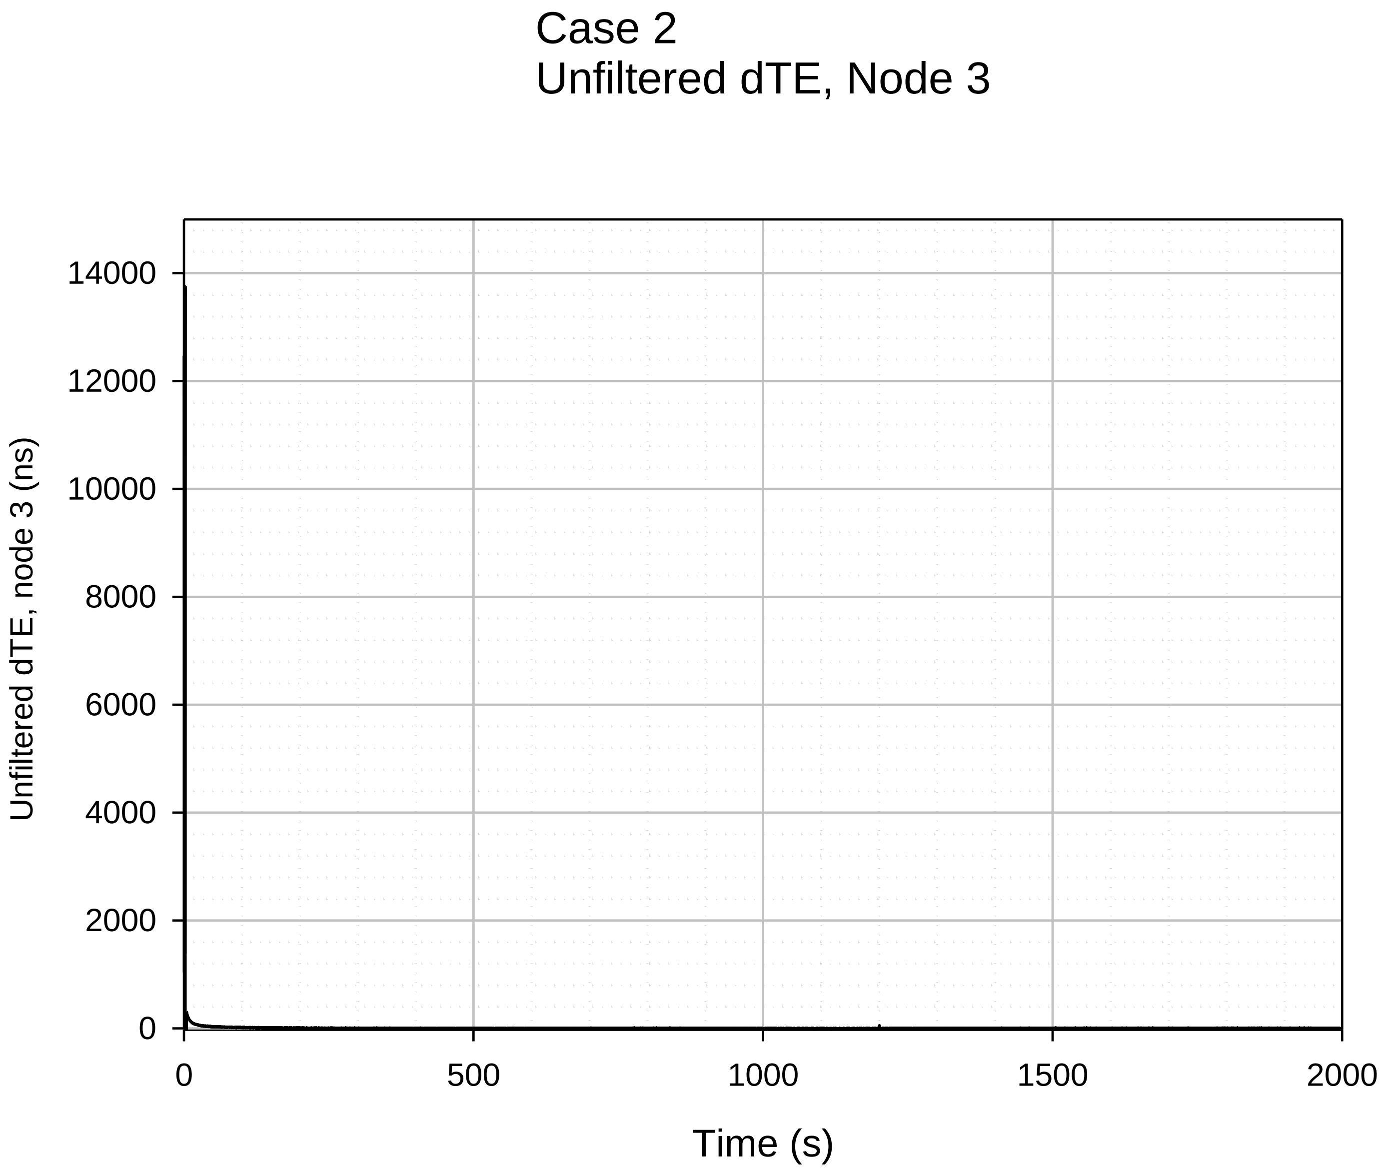

Case 2 - Unfiltered dTE (Detail of 50 s - 2000 s)

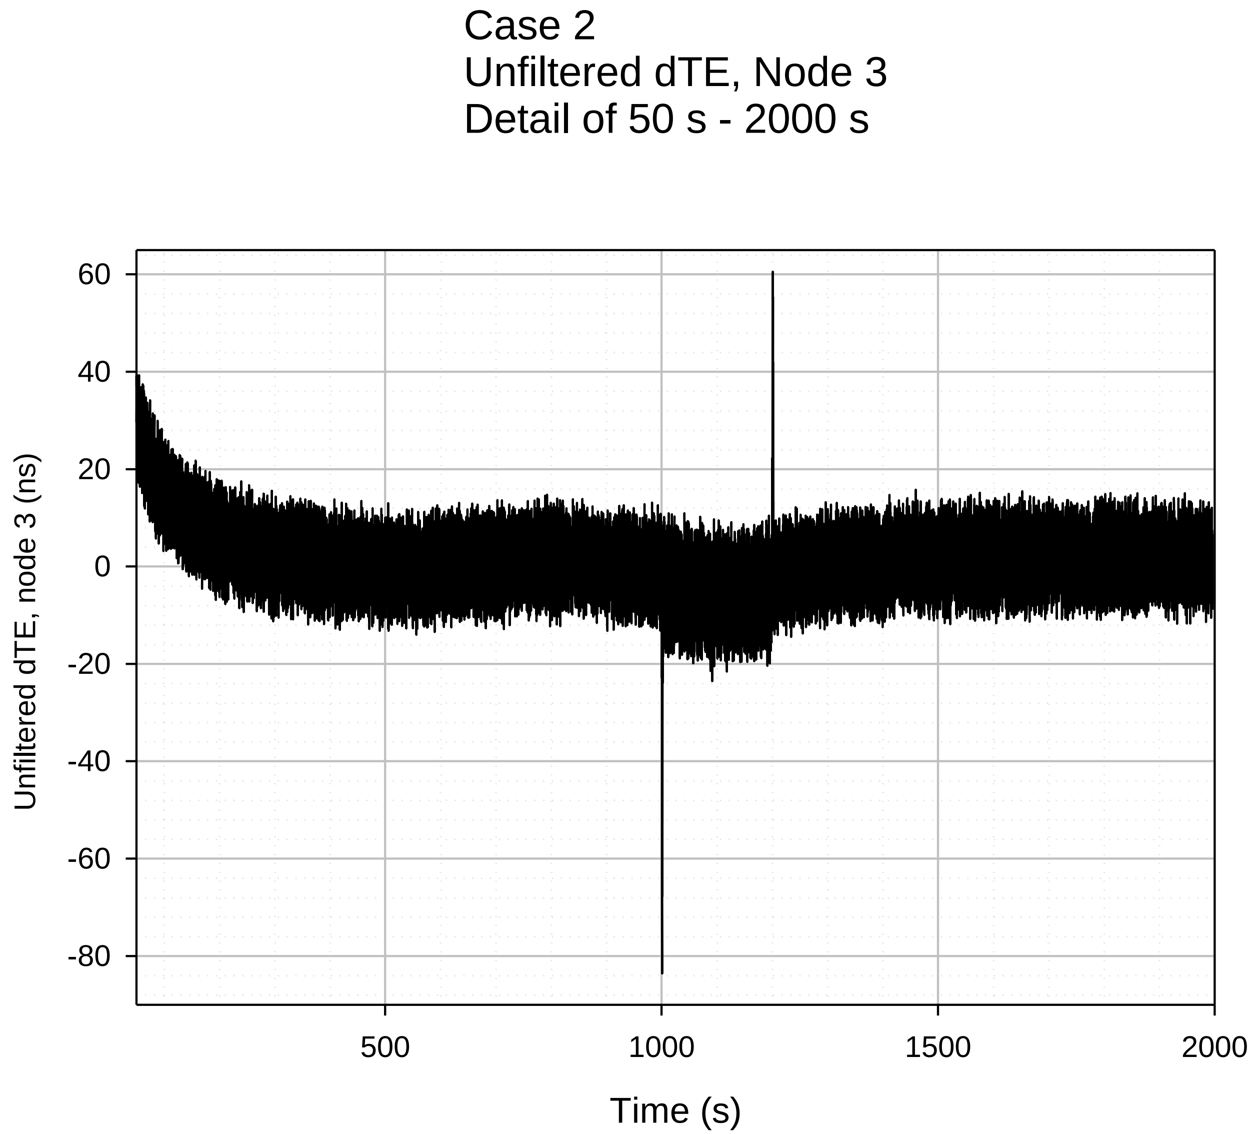

Case 2 - PreciseOriginTS+corrF-WorkingClock at GM

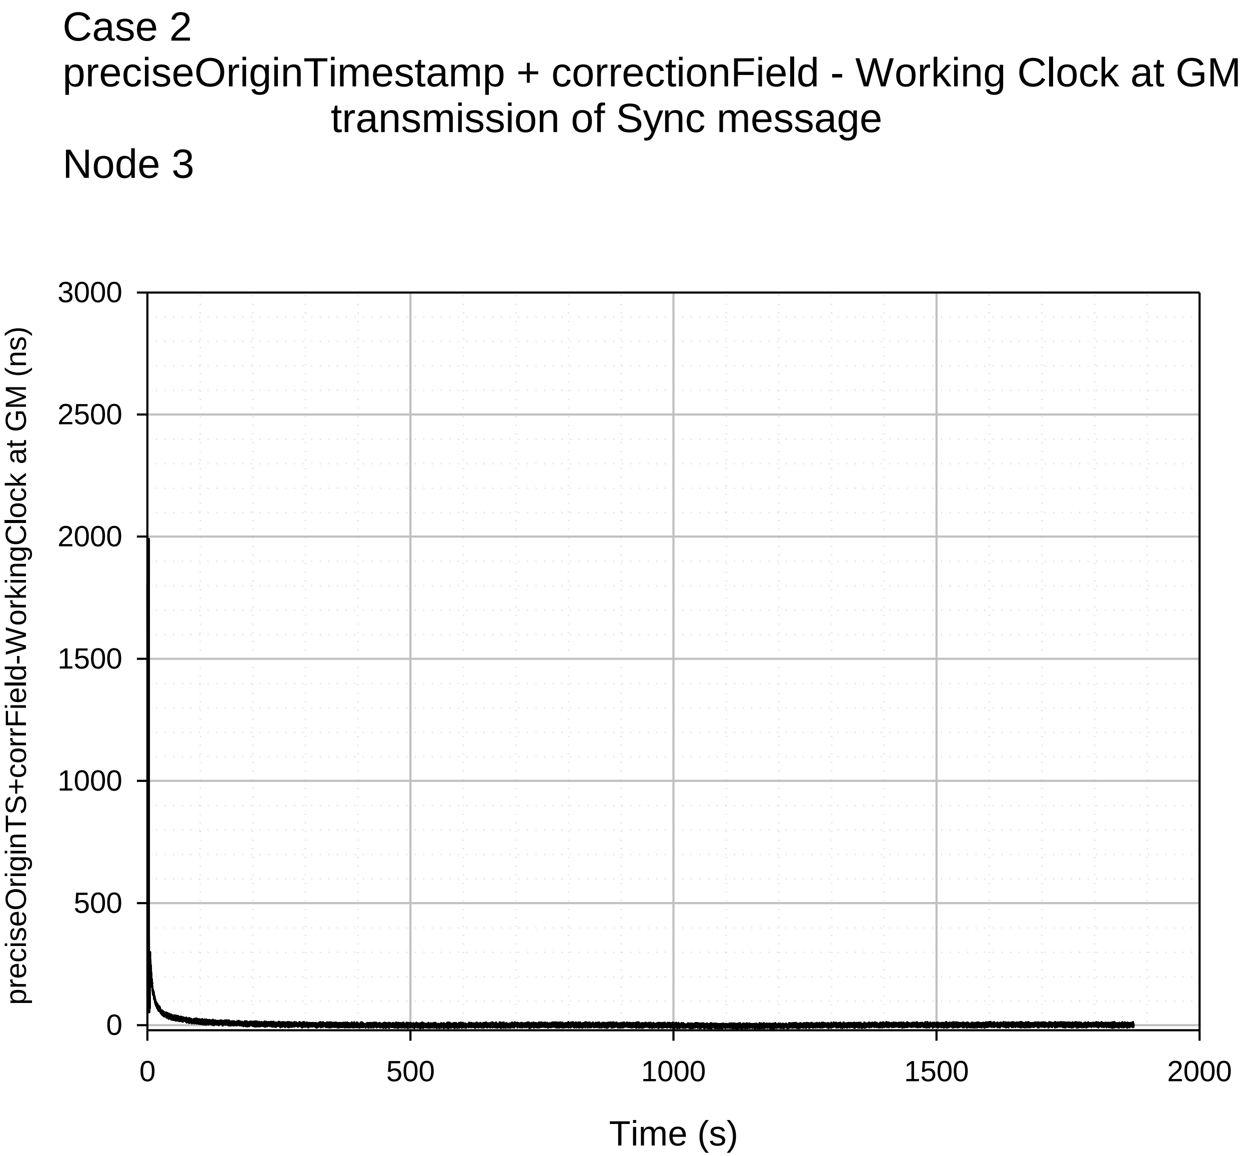

Case 2 - PreciseOriginTS+corrF-WorkingClock at GM (detail of 50 s - 2000 s)

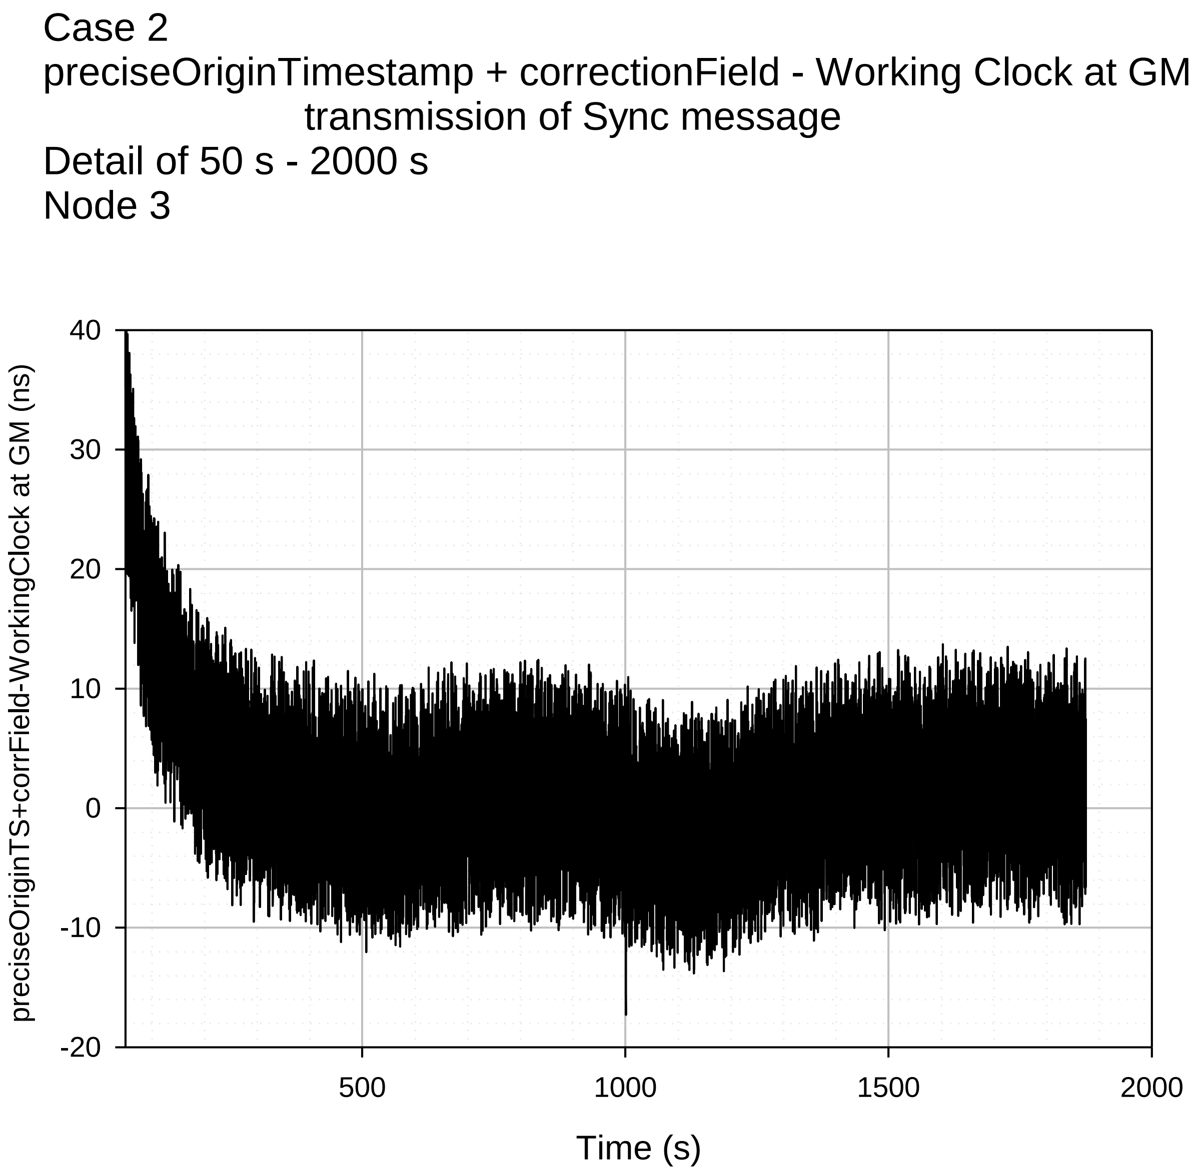

Case 2 - RR field of Sync Minus Actual RR

Case 2
rateRatio field of transmitted Sync message by DUT minus actual rate ratio of DUT working clock relative to Working Clock at GM at transmission of Sync message
Node 3

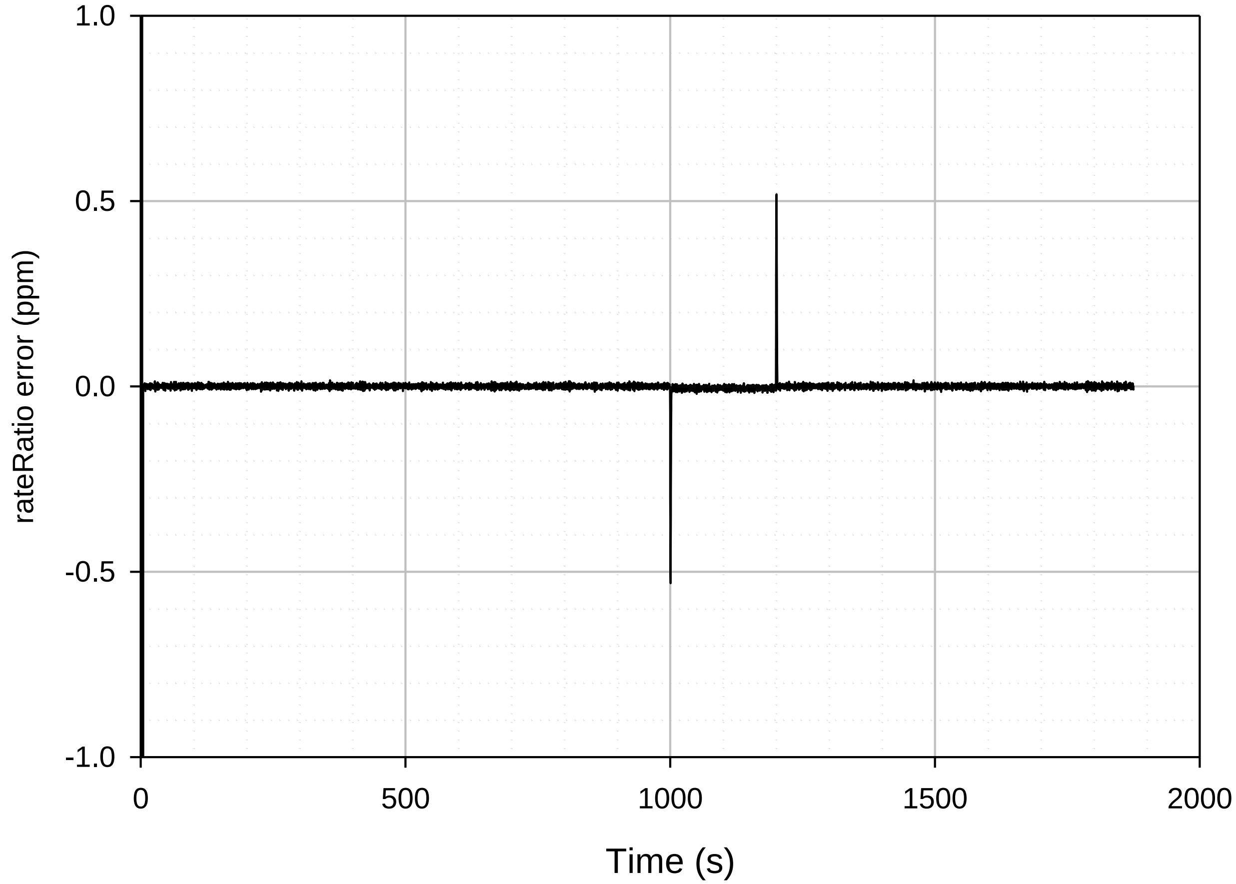

Case 2 - RR field of Sync Minus Actual RR (detail of small ppm)

Case 2
rateRatio field of transmitted Sync message by DUT minus actual rate ratio of DUT working clock relative to Working Clock at GM at transmission of Sync message
Detail of -0.02 ppm to +0.02 ppm
Node 3

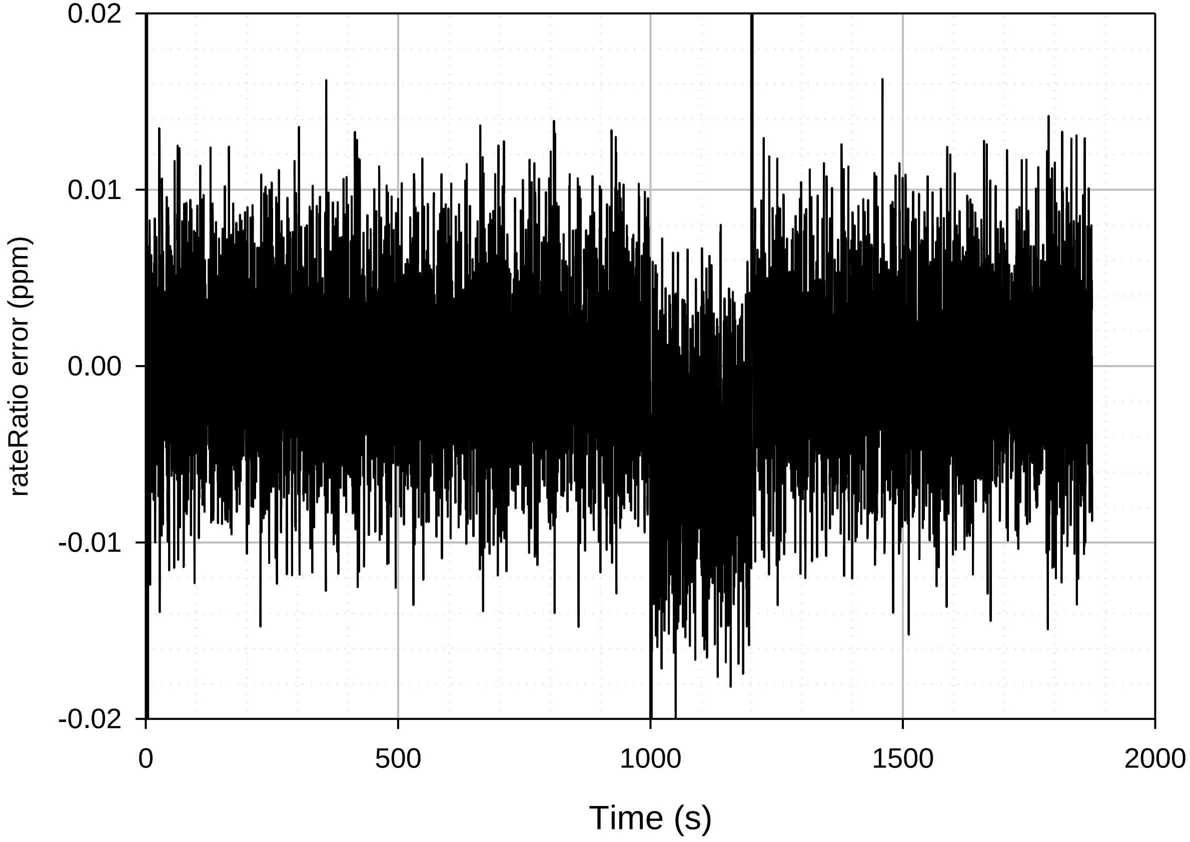

Case 2 – RR drift field of Sync Minus Actual RR Drift

Case 2
rateRatioDrift field of transmitted Sync message by DUT minus actual rate ratio drift rate of DUT working clock relative to Working Clock at GM at transmission of Sync message
Node 3

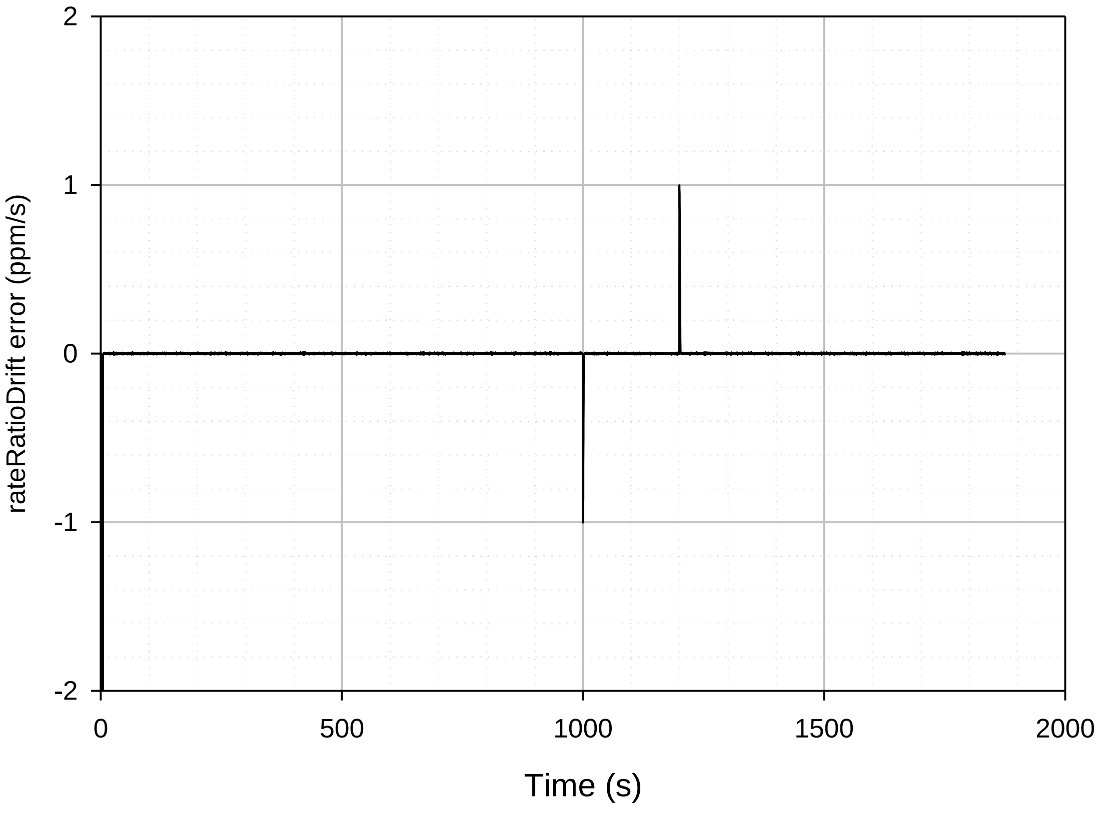

Case 2 – RR drift field of Sync Minus Actual RR Drift (detail of small ppm/s)

Case 2
rateRatioDrift field of transmitted Sync message by DUT minus actual rate ratio drift rate of DUT working clock relative to Working Clock at GM at transmission of Sync message
Detail of -0.02 ppm/s to +0.02 ppm/s
Node 3

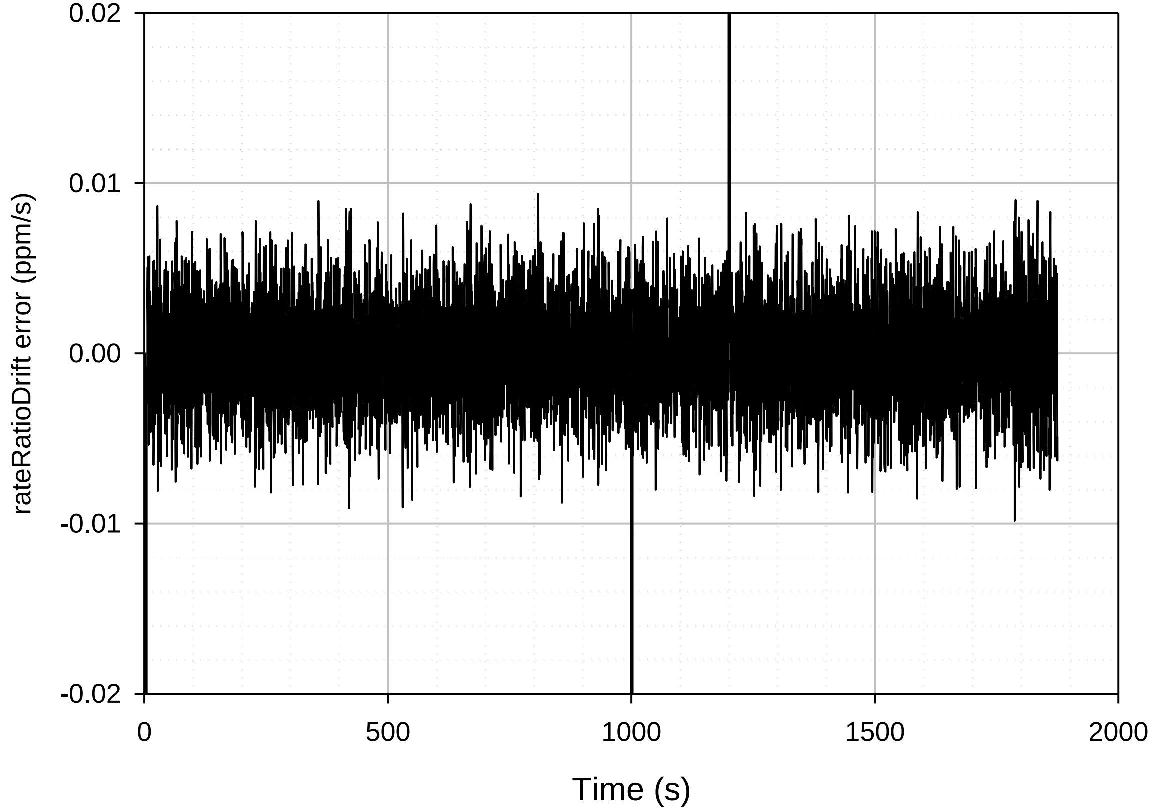

Case 2 - Filtered dTE (PTP End Instance)

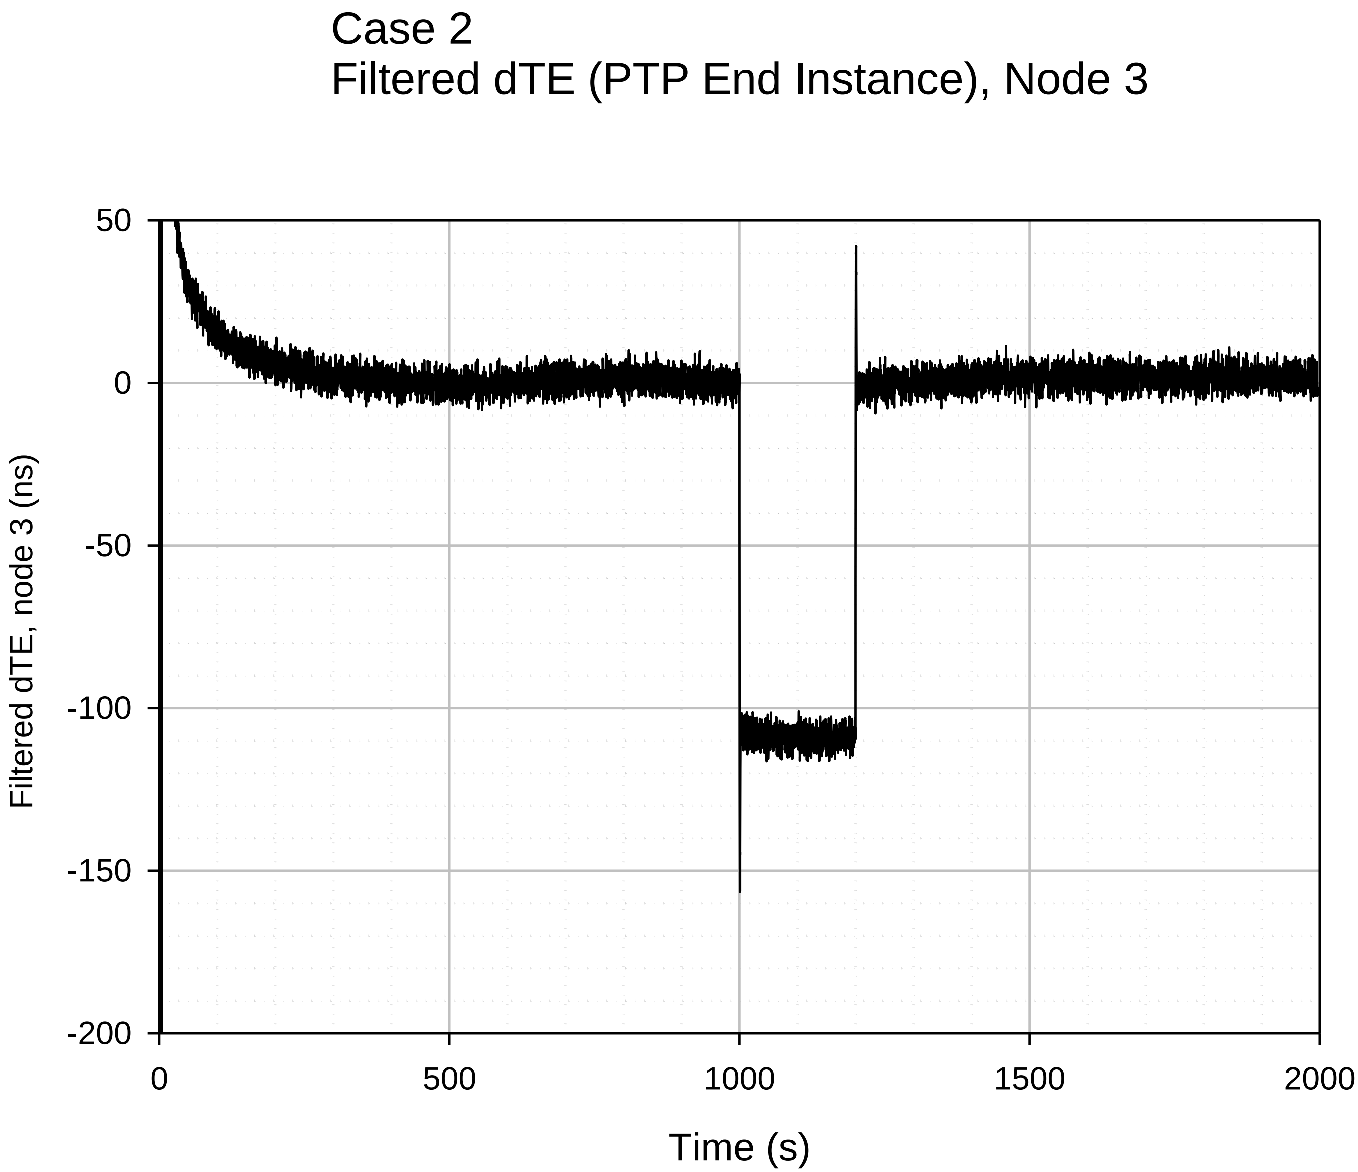

Case 2N - Unfiltered dTE (Detail of 50 s - 2000 s)

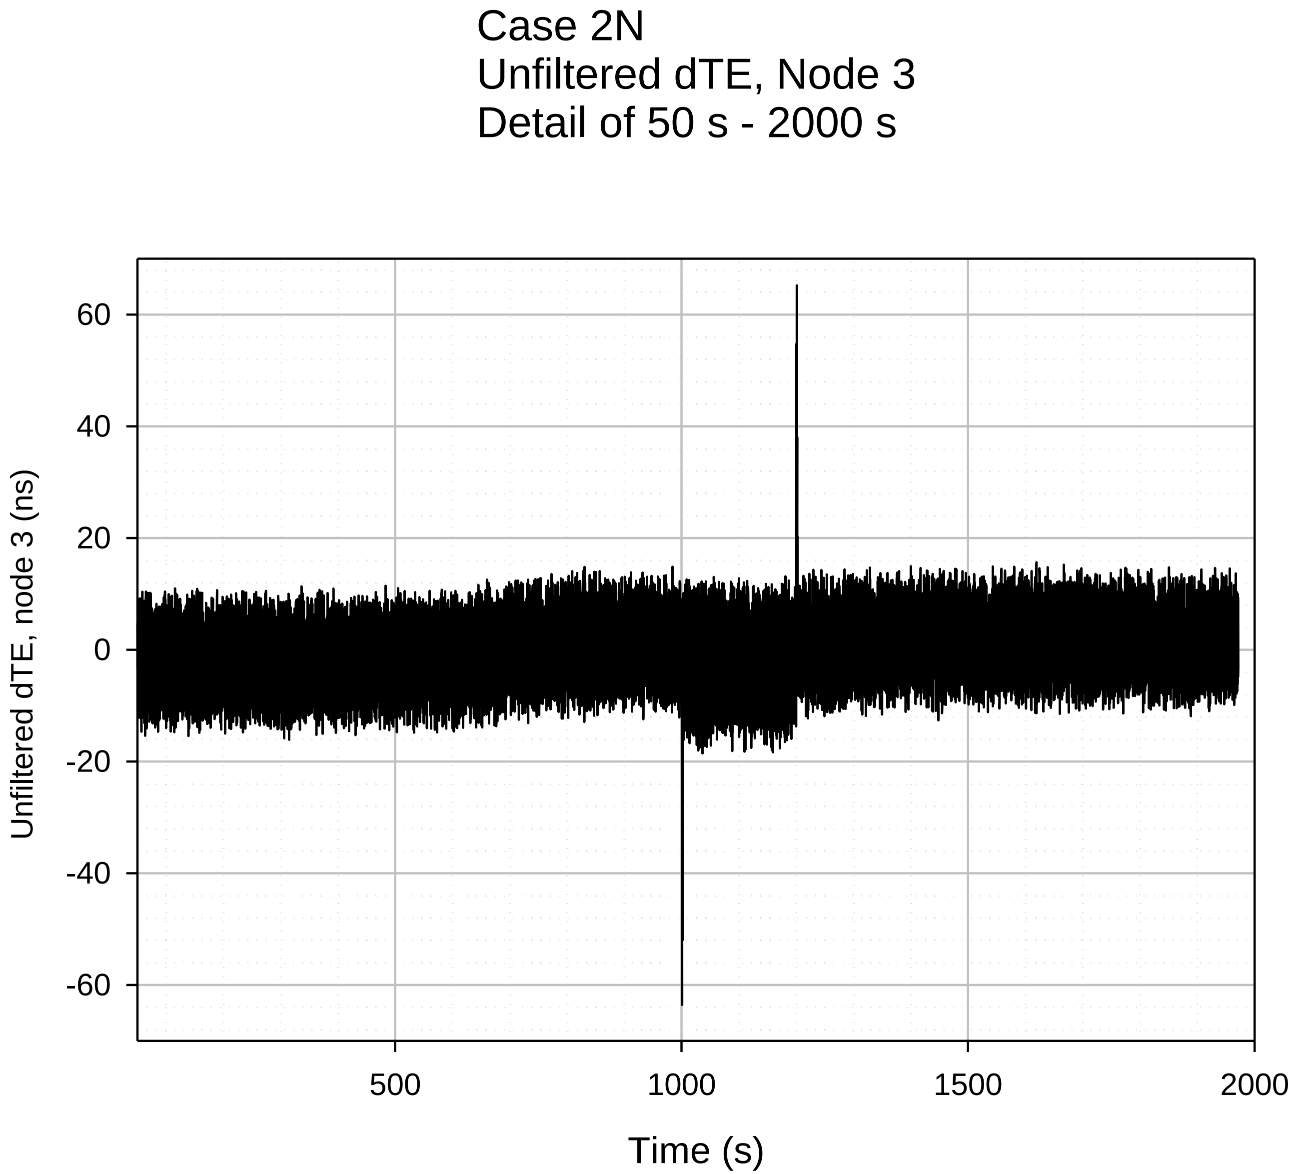

Case 2N - PreciseOriginTS+corrF-WorkingClock at GM (detail of 50 s - 2000 s)

Case 2N
preciseOriginTimestamp + correctionField - Working Clock at GM at transmission of Sync message
Detail of 50 s - 2000 s
Node 3

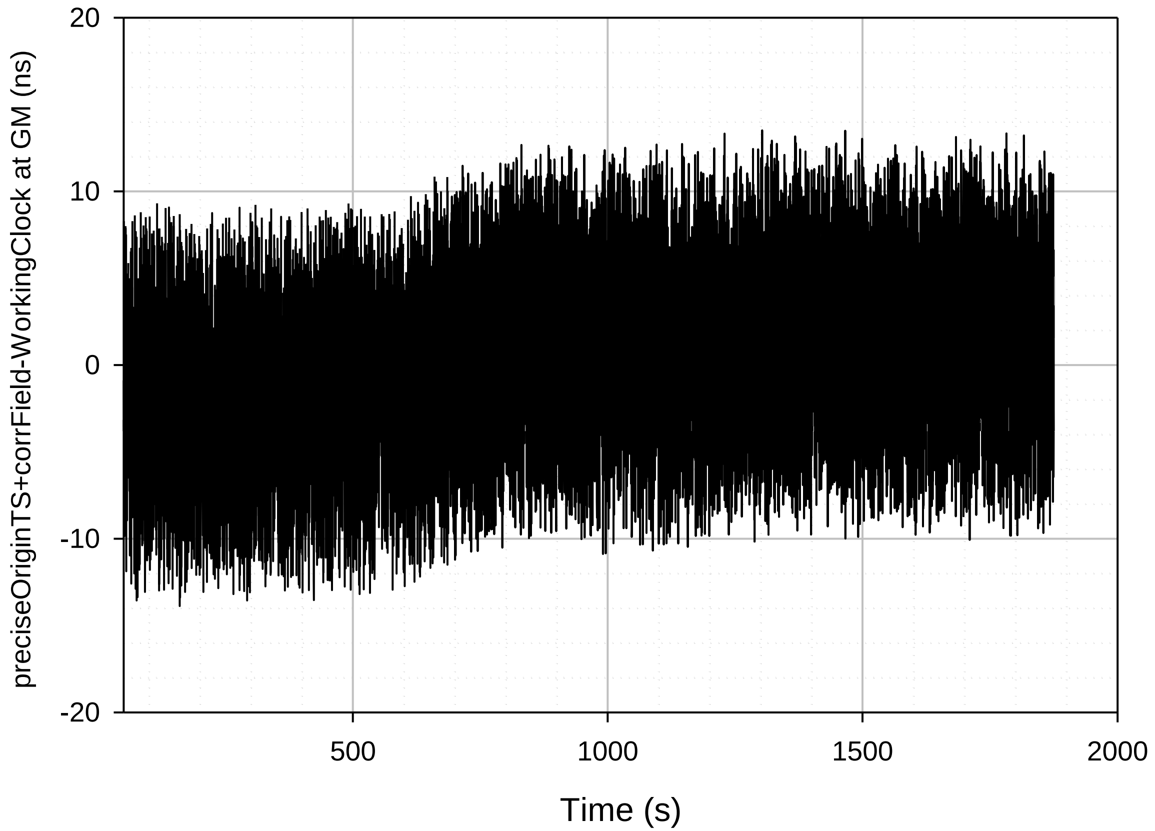

Case 2N - RR field of Sync Minus Actual RR (detail of small ppm)

Case 2N
rateRatio field of transmitted Sync message by DUT minus actual rate ratio of DUT working clock relative to Working Clock at GM at transmission of Sync message
Detail of -0.03 ppm to +0.03 ppm
Node 3

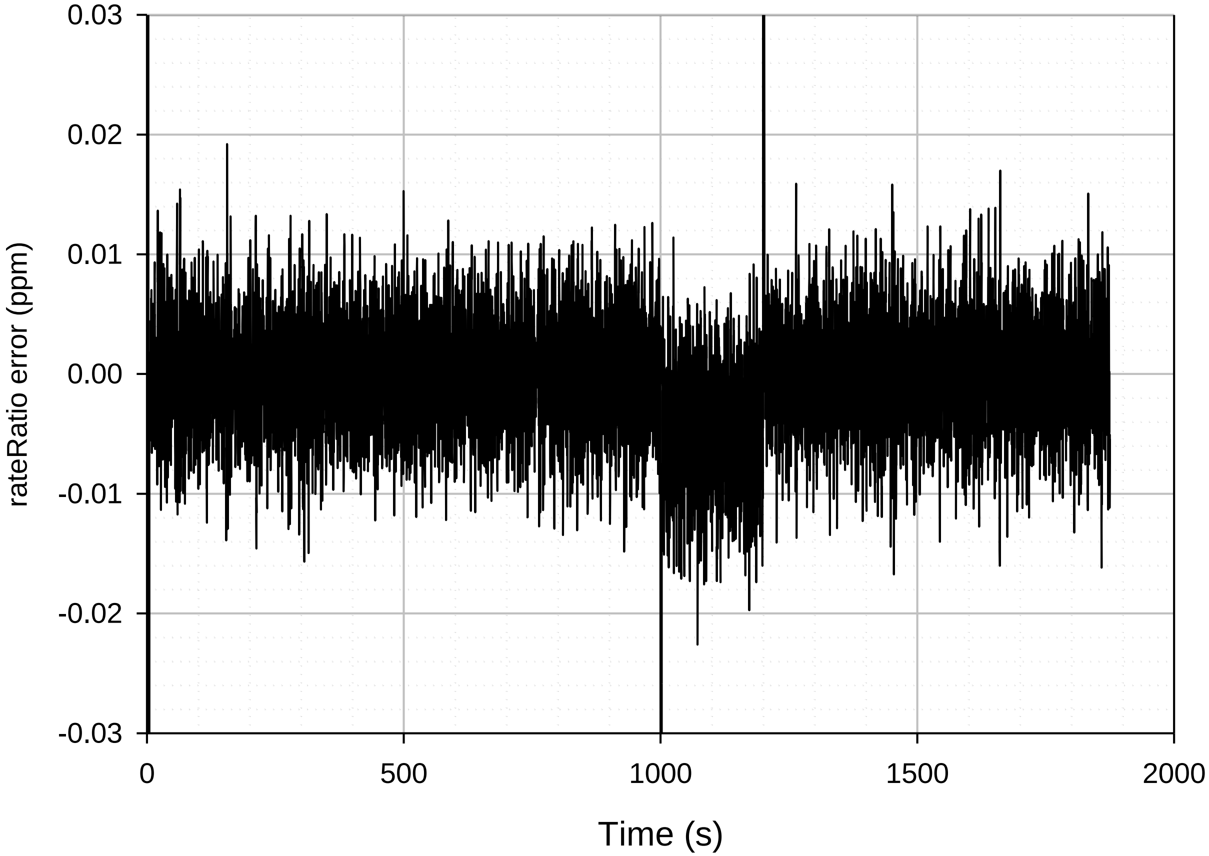

Case 2N – RR drift field of Sync Minus Actual RR Drift (detail of small ppm/s)

Case 2N
rateRatioDrift field of transmitted Sync message by DUT minus actual rate ratio drift rate of DUT working clock relative to Working Clock at GM at transmission of Sync message
Detail of -0.02 ppm to +0.02 ppm
Node 3

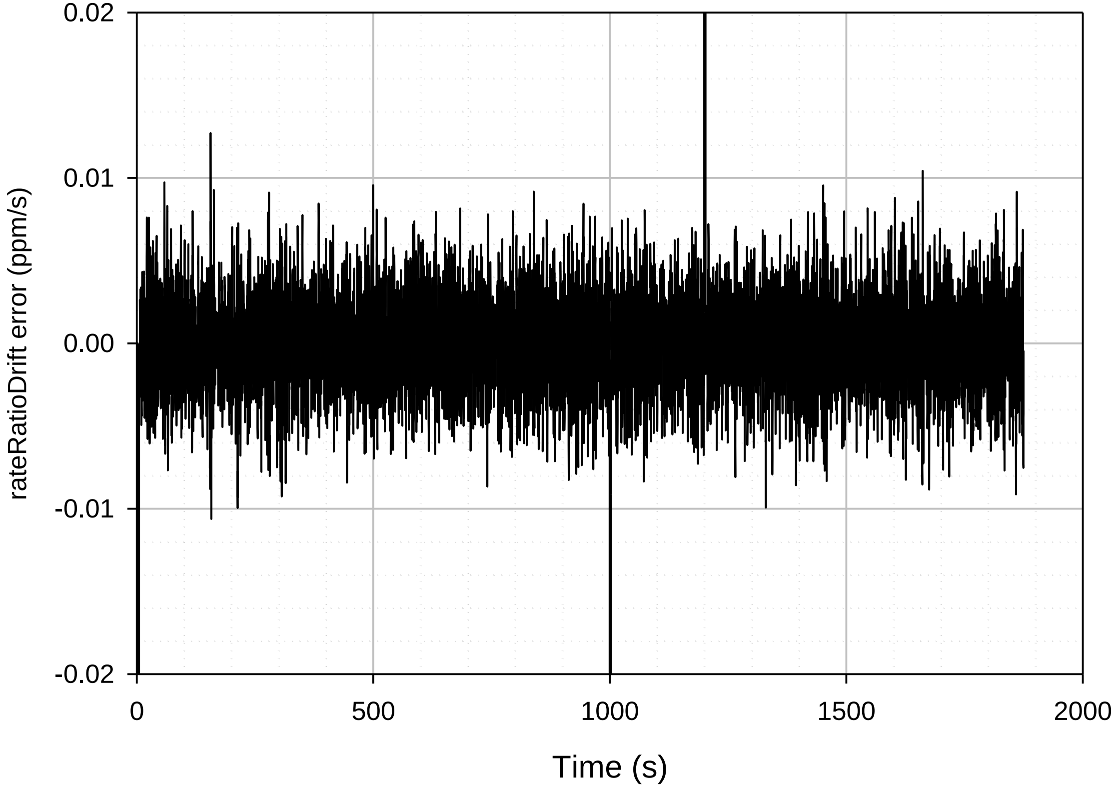

Case 2N - Filtered dTE (PTP End Instance)

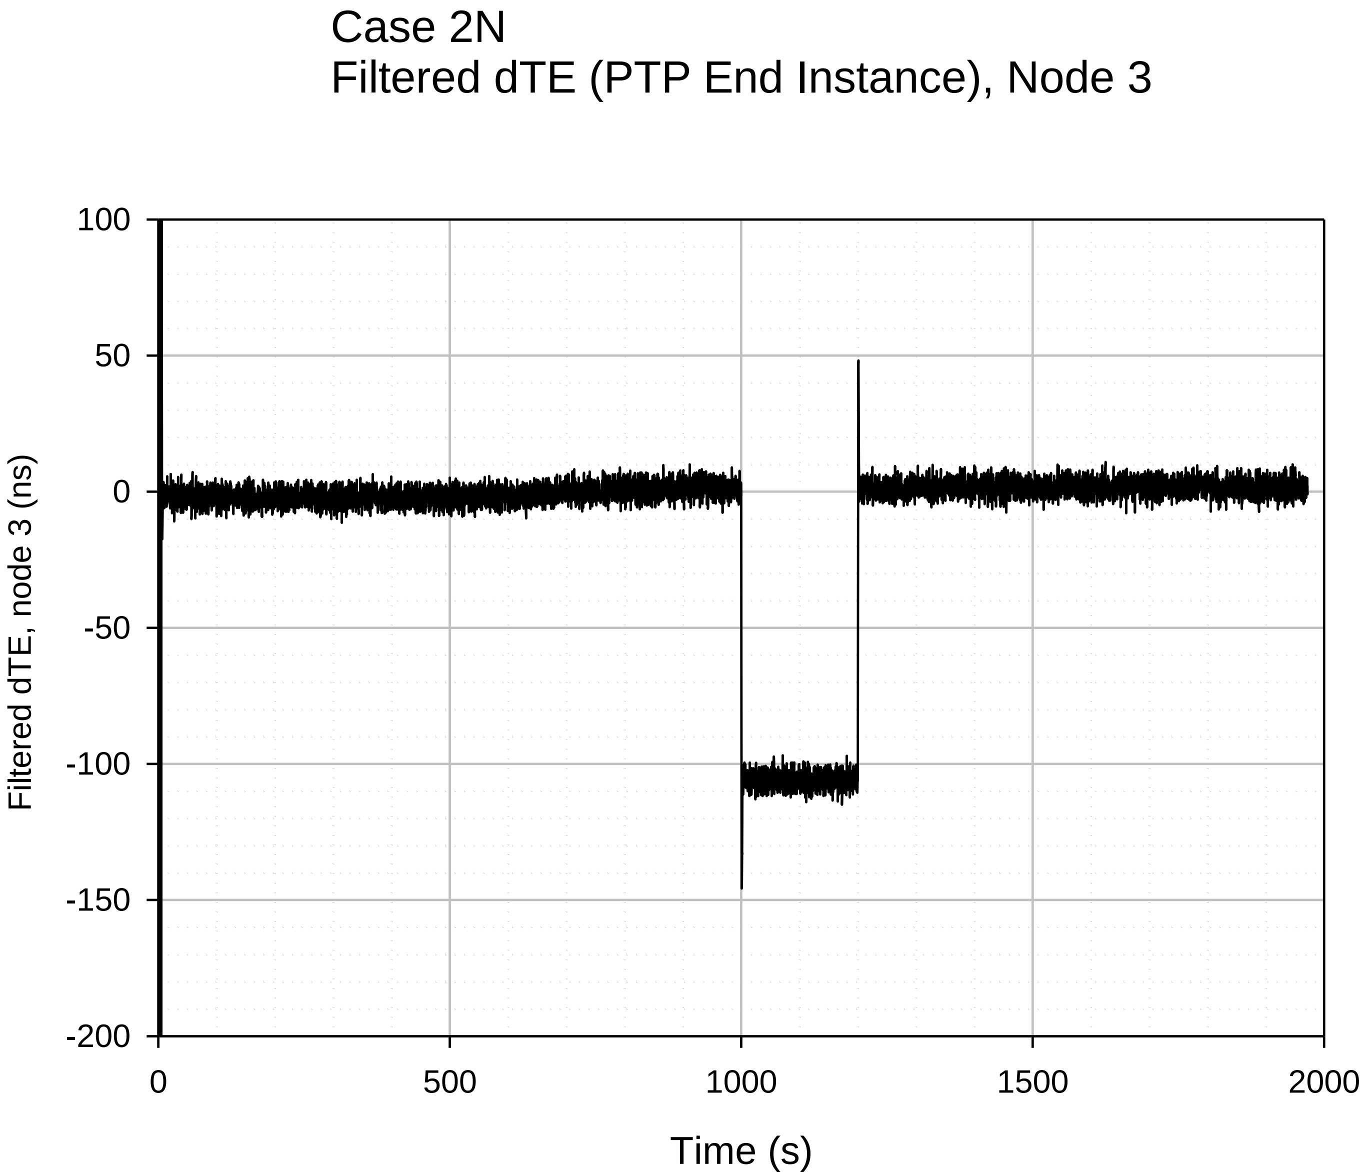

Numerical Results for Cases 2 and 2N

Case 2

Quantity	Mean	Standard Deviation	Minimum	Maximum
dTE_{unfilt} (ns)	0.1633	5.341	-83.60	60.57
M (ns)	0.609	5.075	-17.30	13.72
N (ppm)	-6.666×10^{-4}	1.980×10^{-2}	-0.5309	0.5179
P (ppm/s)	1.602×10^{-6}	3.892×10^{-2}	-1.002	0.9981
dTE_{filt} (ns)	-13.68	37.42	-156.5	42.12

Case 2N

Quantity	Mean	Standard Deviation	Minimum	Maximum
dTE_{unfilt} (ns)	0.5938	5.094	-63.57	65.24
M (ns)	1.027	5.048	-13.20	13.51
N (ppm)	-6.707×10^{-4}	1.952×10^{-2}	-0.5125	0.5156
P (ppm/s)	7.579×10^{-6}	3.858×10^{-2}	-0.9951	1.004
dTE_{filt} (ns)	-13.23	36.43	-145.8	48.22

Cases 2 and 2N Results – Discussion - 1

- Cases 2 and 2N differ mainly in that transients are more pronounced in Case 2 due to the presence of residence time at node 2 and Pdelay turnaround time at nodes 1 and 2
  - As for the results for cases 1 and 1N, results here for minimum and maximum values and standard deviations, for corresponding parameters, are similar
  - Results for mean values for cases 1 and 1N differ, most likely due to statistical variation
- Results for unfiltered dTE, N, and P show transient effect (spikes) at 1000 s and 1200 s (i.e., when the clock drift at node 1 begins and ends
- Results for M show transient effect at 1000 s, but not at 1200 s, in case 2, and no transient effects at 1000 s and 1200 s in case 2N
- Results for filtered dTE show transient effects (spikes) at both 1000 s and 1200 s for both cases; in addition, both cases show a phase offset of approximately -100 ns during the clock drift

Cases 2 and 2N Results - Discussion - 2

- The -100 ns phase offset is due to the response of the second-order filter to the 1 ppm/s frequency drift
- To see this, note that the phase drift corresponding to an $A = 1$ ppm/s $= 1000$ ns/s frequency drift is $0.5At^2$, where t is the time in seconds
  - The Laplace transform of this waveform is $U(s) = A/s^3$
  - The steady-state value of the filter output due to this drift is obtained from the final value theorem:

$$\text{steady-state-response} = \lim_{s\to 0} sH(s)U(s) = \lim_{s\to 0} s\cdot\frac{s^2}{s^2+2\varsigma\omega_n s+\omega_n^2}\cdot\frac{A}{s^3} = \frac{A}{\omega_n^2}$$

where

$\omega_n = \text{undamped natural frequency} = 3.1011 \text{ rad/s}$

$\varsigma = \text{ damping ratio } = \text{ } 2.1985 \text{ dB}$

  - Then the steady-state response is

$$\frac{A}{\omega_n^2} = \frac{1000 \text{ ns/s}}{(3.1011 \text{ rad/s})^2} = 104 \text{ ns}$$

  - This is in agreement with the steady-state filtered dTE during the 1 ppm/s clock drift for cases 2 and 2N

Case 2 and 2N Results - Discussion - 3

- Some, but not all, of the requirements of Tables 12 and 13 of [1] are met:
  - Mean of M is within range of -2 ns to +2 ns, and minimum and maximum values are within ±20 ns range (row 1 of Table 12 of [1])
  - Mean of N is within range of -0.1 ppm to +0.1 ppm, and standard deviation is within ±0.02 ppm (row 2 of Table 12 of [1])
  - Mean of P is within range of -0.1 ppm/s to +0.1 ppm/s; however, **standard deviation exceeds ±0.02 ppm/s** (row 5 of Table 12 of [1])
  - dTE_{filt} **exceeds the range of -15 ns to +15 ns** (row 1 of Table 13 of [1]), mainly due to the transients (spikes) and offset caused by the 1 ppm/s clock drift

Case 2 and 2N Results – Discussion - 2

- Like cases 1 and 1N, the results for unfiltered dTE and M are similar but not identical; M values are approximately equal to dTE values computed at times that Sync messages are transmitted by the DUT, while dTE values are also computed at times in between these transmit times
- The results for minimum and maximum values, and standard deviation, for M, are in reasonable agreement with results for dTE for Test Type 1 on p.16 of [3]; however, the mean values are larger by several orders of magnitude. This is likely due to the variation in measured mean link delay error obtained here.
  - As for cases 1 and 1N, means and standard deviations are computed here as time averages, but it has not been established the these have converged, or even that the processes are ergodic (i.e., roughly that time averages converge to ensemble averages). It would be desirable to at least run multiple replications of the simulations to consider this further.
- Results for minimum and maximum values, and standard deviation for N and P are larger than corresponding results for rateRatio and rateRatioDrift given for Test Type 1 on p.16 of [3]; this is likely due to variation in measured mean link delay error and explicit modeling of the time dependence of clock drift.

Case 3 - Unfiltered dTE (Detail of 50 s - 2000 s)

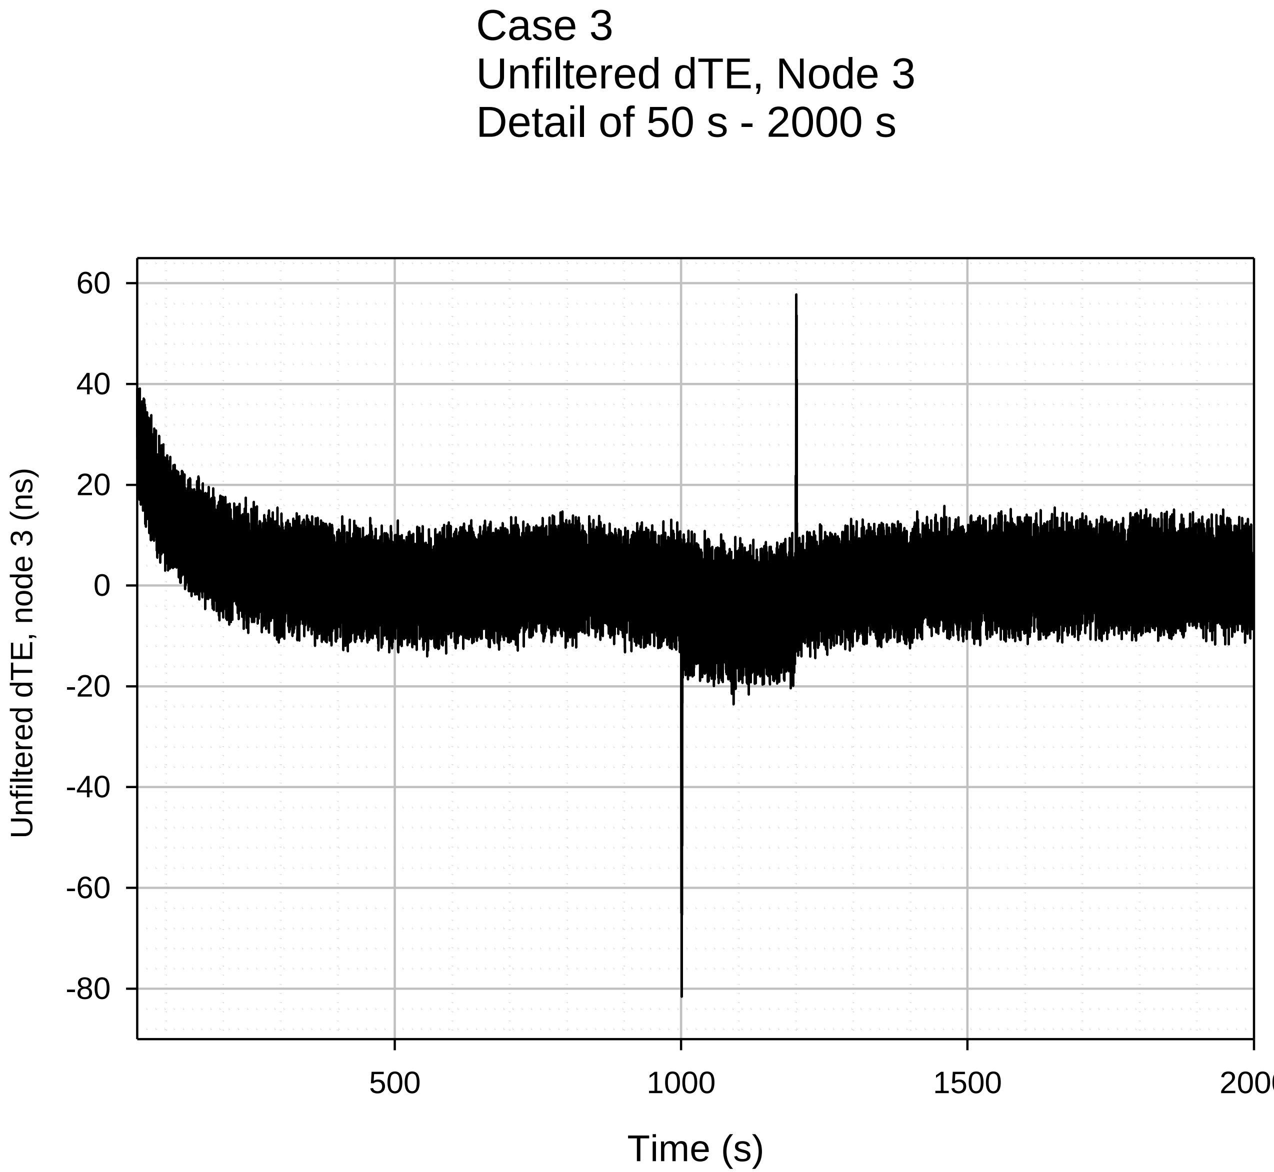

Case 3 – PreciseOriginTS+corrF-WorkingClock at GM (detail of 50 s – 2000 s)

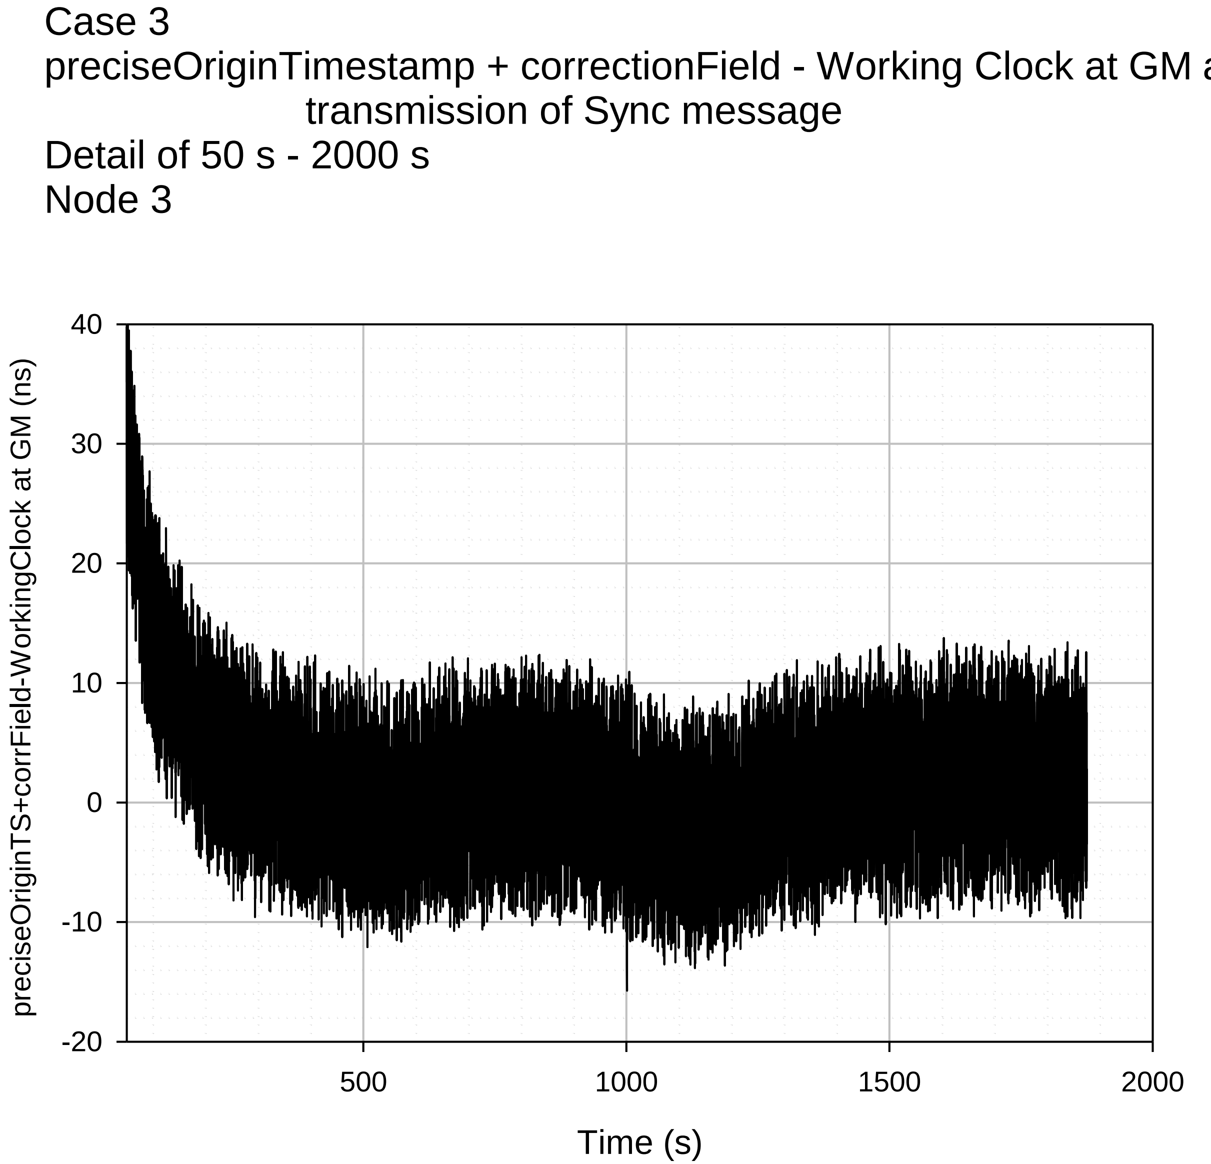

Case 3 - RR field of Sync Minus Actual RR (detail of small ppm)

Case 3
rateRatio field of transmitted Sync message by DUT minus actual rate ratio of DUT working clock relative to Working Clock at GM at transmission of Sync message
Detail of -0.02 ppm to +0.02 ppm
Node 3

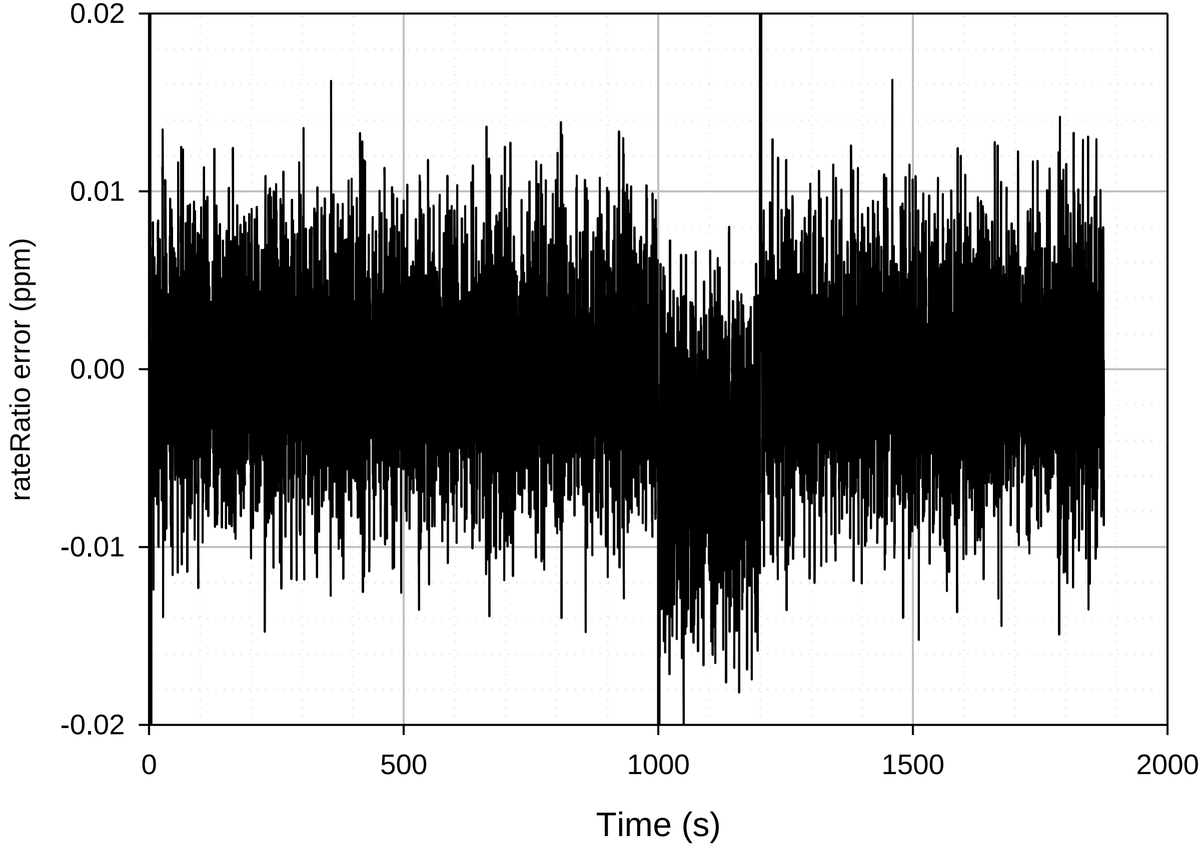

Case 3 – RR drift field of Sync Minus Actual RR Drift (detail of small ppm/s)

Case 3
rateRatioDrift field of transmitted Sync message by DUT minus actual rate ratio drift rate of DUT working clock relative to Working Clock at GM at transmission of Sync message
Detail of -0.02 ppm to +0.02 ppm
Node 3

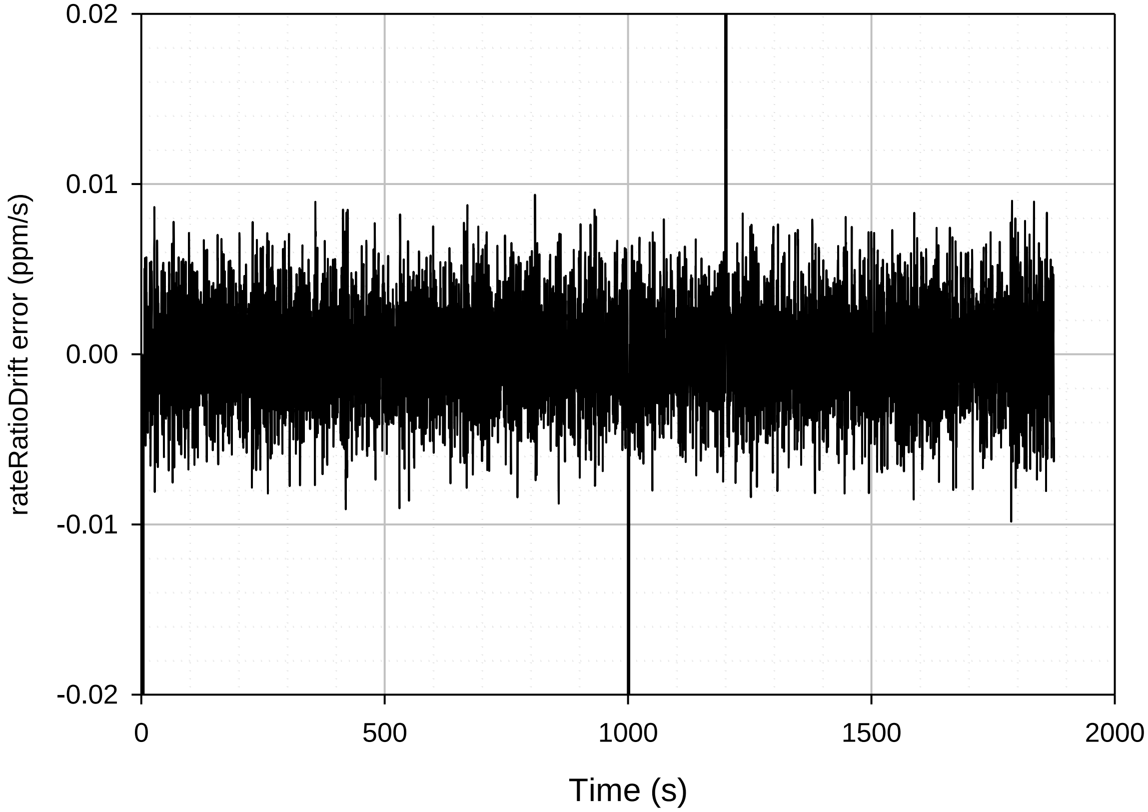

Case 3 – Filtered dTE (PTP End Instance)

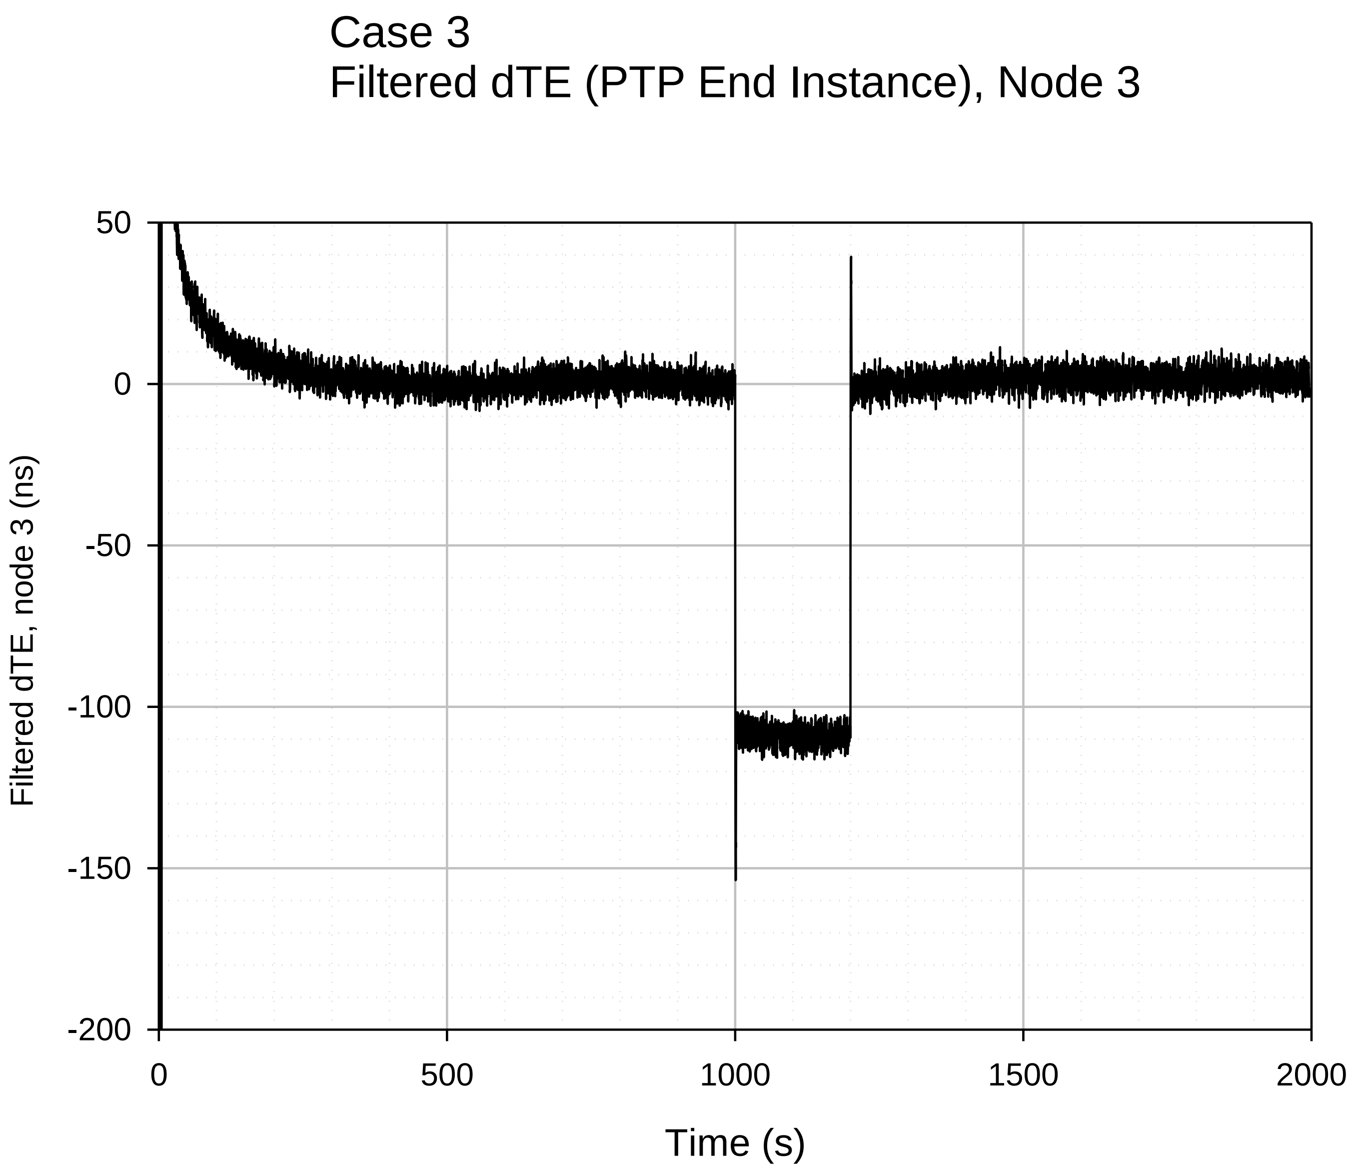

Case 3N - Unfiltered dTE (Detail of 50 s - 2000 s)

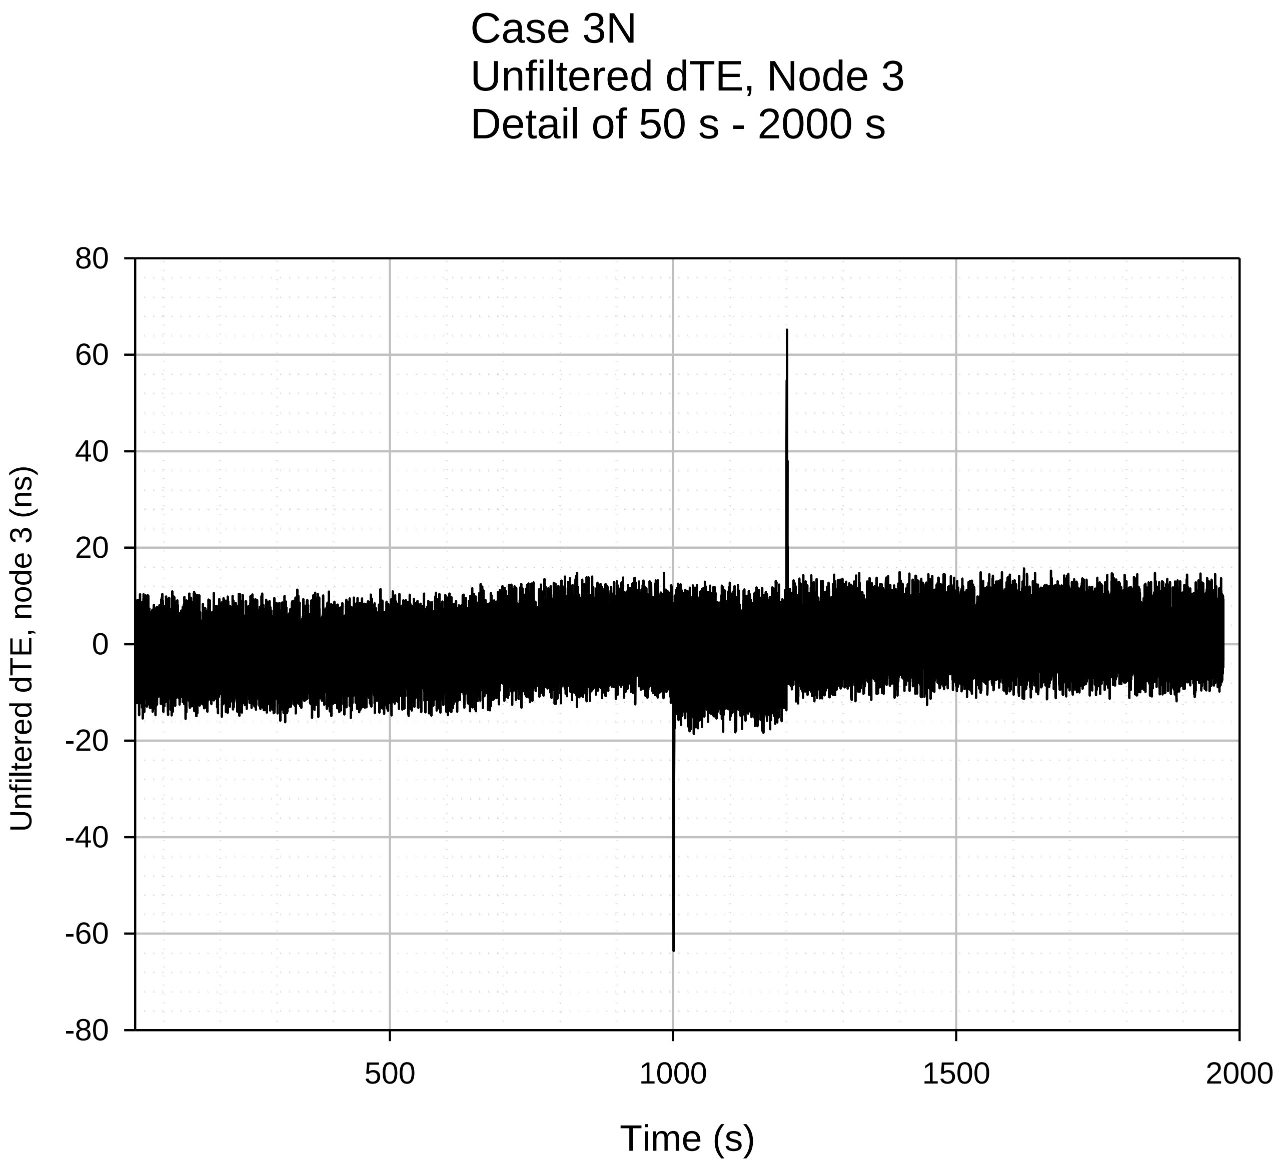

Case 3N - PreciseOriginTS+corrF-WorkingClock at GM (detail of 50 s - 2000 s)

Case 3N
preciseOriginTimestamp + correctionField - Working Clock at GM at transmission of Sync message
Detail of 50 s - 2000 s
Node 3

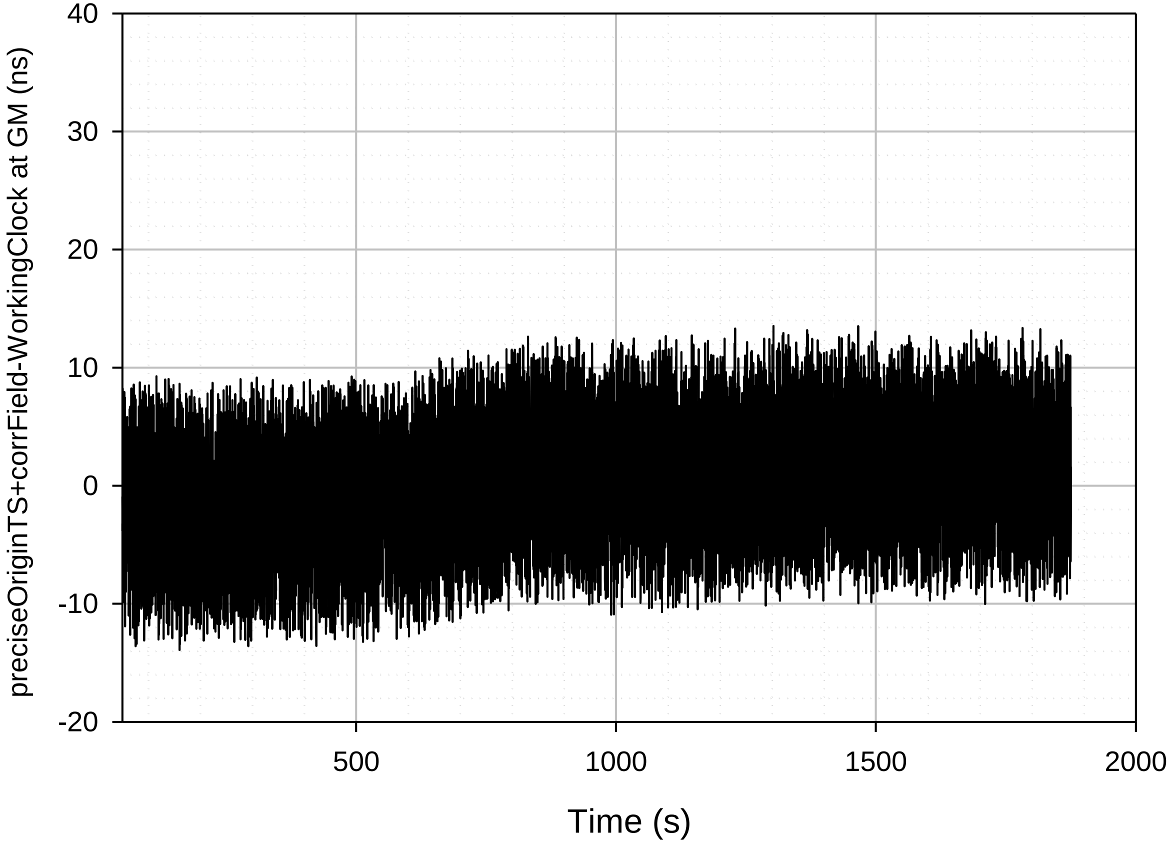

Case 3N – RR field of Sync Minus Actual RR (detail of small ppm)

Case 3N
rateRatio field of transmitted Sync message by DUT minus actual rate ratio of DUT working clock relative to Working Clock at GM at transmission of Sync message
Detail of -0.03 ppm to +0.03 ppm
Node 3

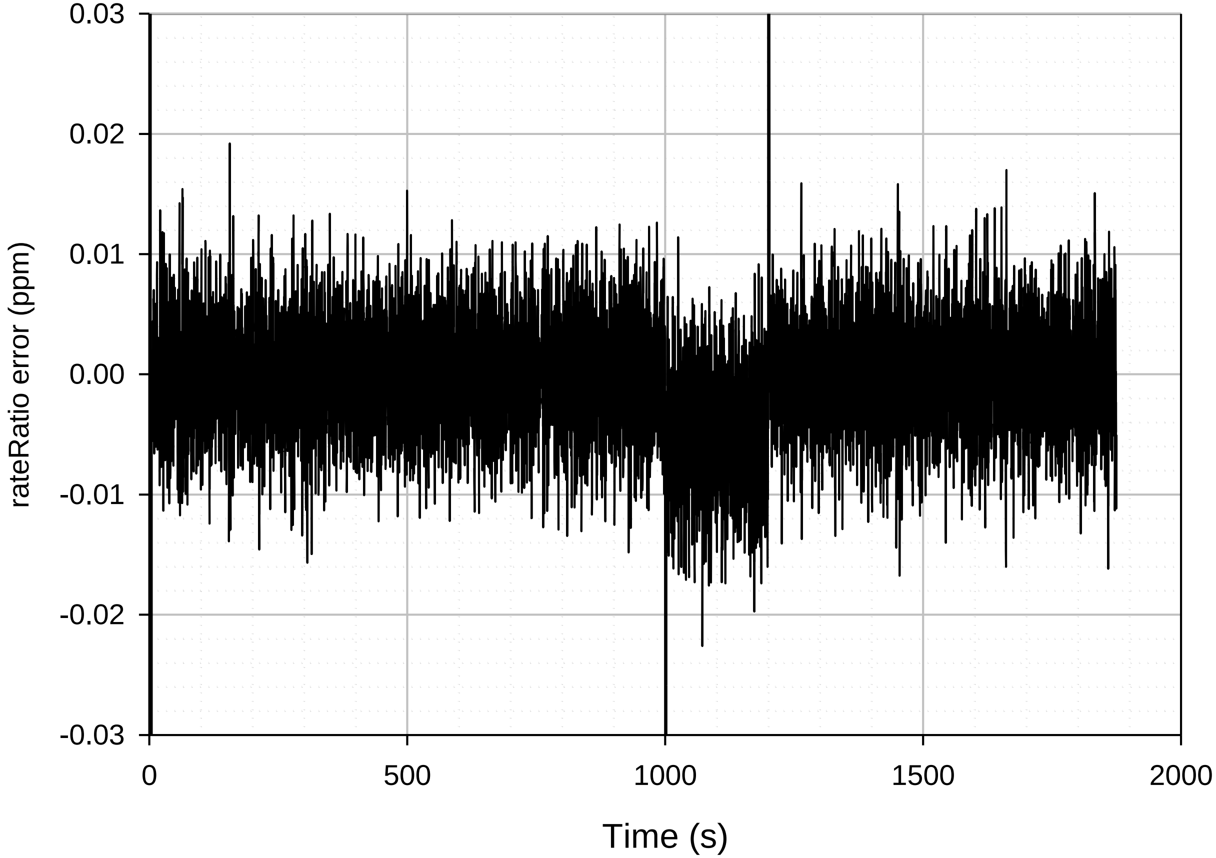

Case 3N – RR drift field of Sync Minus Actual RR Drift (detail of small ppm/s)

Case 3N
rateRatioDrift field of transmitted Sync message by DUT minus actual rate ratio drift rate of DUT working clock relative to Working Clock at GM at transmission of Sync message
Detail of -0.02 ppm to +0.02 ppm
Node 3

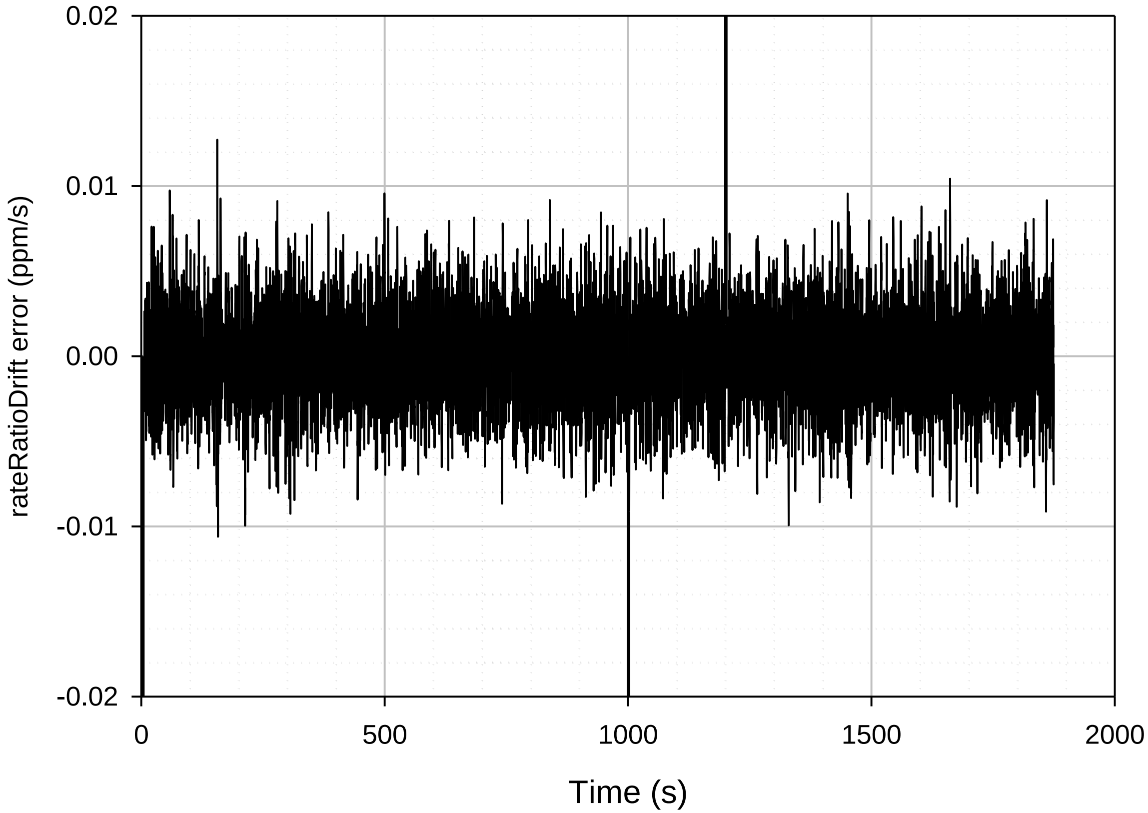

Case 3N - Filtered dTE (PTP End Instance)

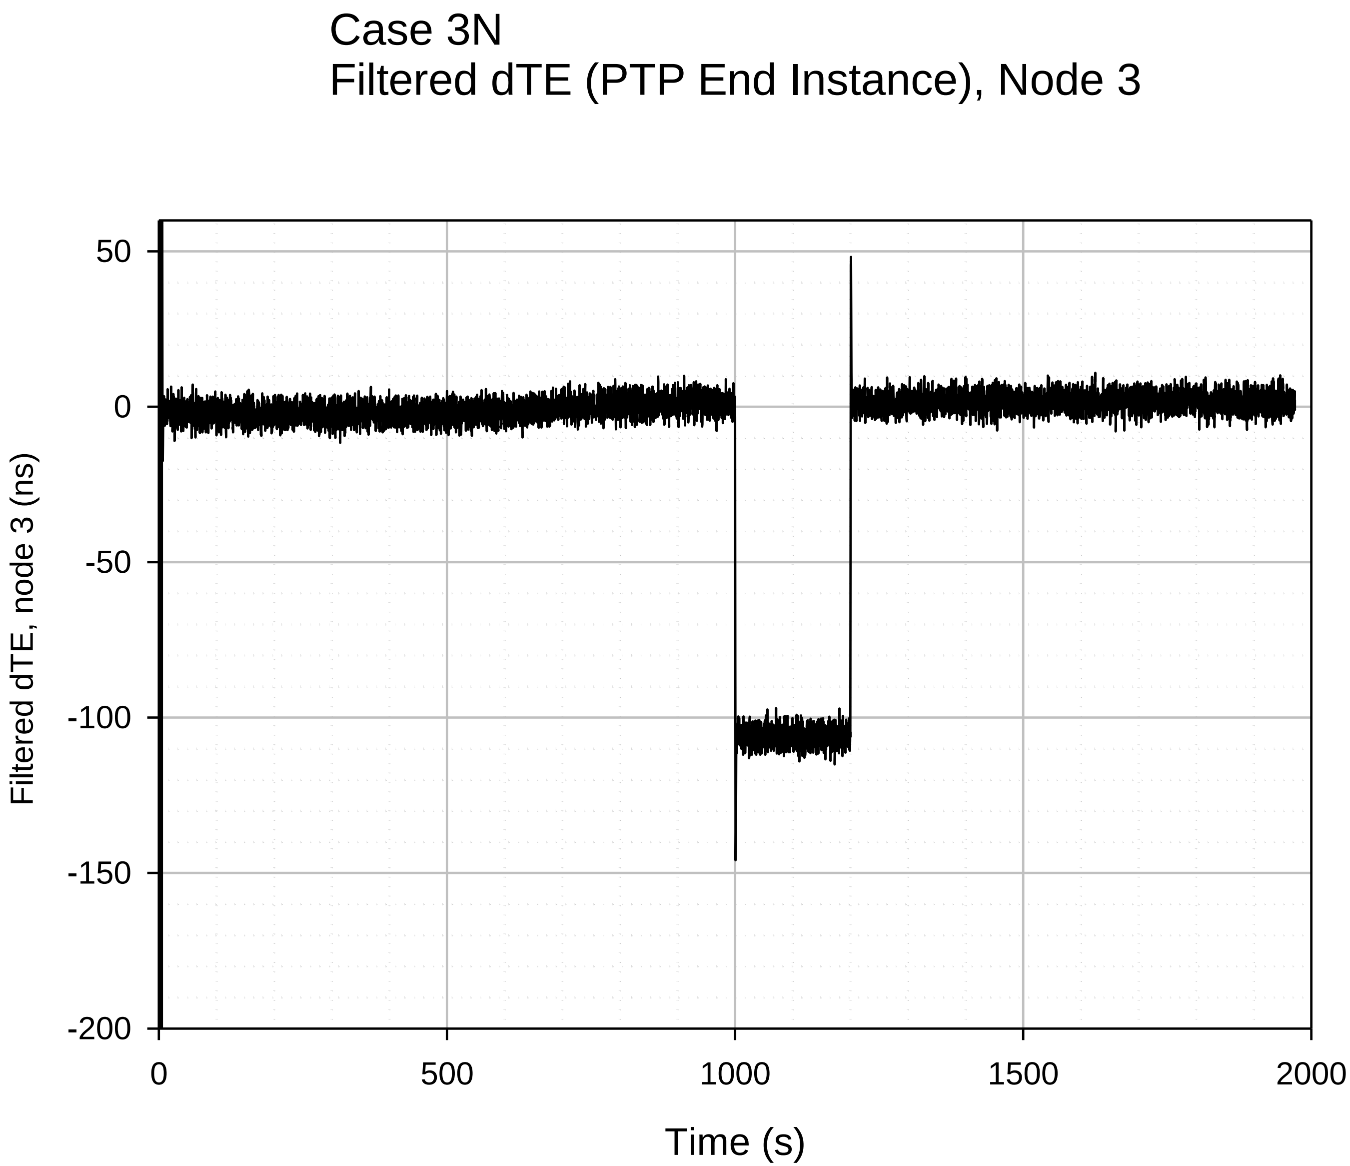

Numerical Results for Cases 3 and 3N

Case 3

Quantity	Mean	Standard Deviation	Minimum	Maximum
dTE_{unfilt} (ns)	0.1629	5.327	-81.61	57.80
M (ns)	0.611	5.078	-15.74	13.76
N (ppm)	-6.668×10^{-4}	1.966×10^{-2}	-0.5271	0.5144
P (ppm/s)	1.493×10^{-6}	3.884×10^{-2}	-1.002	0.9980
dTE_{filt} (ns)	-13.68	37.43	-153.8	39.37

Case 3N

Quantity	Mean	Standard Deviation	Minimum	Maximum
dTE_{unfilt} (ns)	0.5939	5.102	-63.61	65.23
M (ns)	1.029	5.053	-13.24	13.52
N (ppm)	-6.708×10^{-4}	1.952×10^{-2}	-0.5125	0.5156
P (ppm/s)	7.579×10^{-6}	3.858×10^{-2}	-0.9951	1.004
dTE_{filt} (ns)	-13.23	36.44	-145.9	48.21

Cases 3 and 3N Results - Discussion - 1

- Results for cases 3 and 3N are very similar to corresponding results for cases 2 and 2N
- Therefore, the discussion of the case 2 and 2N results also is valid for cases 3 and 3N

Conclusion - 1

- While some of the requirements of Tables 12 and 13 of [1] are not met for cases 2, 2N, 3, and 3N, this does ***not*** indicate a fundamental problem, because previous simulations for overall network performance [3], [4] indicated that, subject to the assumptions of Annex D of [1], the overall objective of 500 ns for max|dTE_{rel}| can be met.
- However, some of the Annex D, Table 12 and 13 error generation requirements and/or tests should be revisited
- In particular, the Table 13 requirement for max|dTE_{filt}| requirement or test needs to be revisited, in view of the effect of the endpoint filter on entry into the 1 ppm/s drift, steady-state response of the 1 ppm/s drift, and exit from the 1 ppm/s drift (the entry and exit conditions are, in effect, transients)
  - Note that the response of the endpoint filter on max|dTE_{filt}| will be larger for a 0.7 Hz bandwidth by a factor of 1.0/0.7 = 1.429

Conclusion - 2

- The effect of the clock drift transients on rateRatio and rateRatioDrift also indicate that the Table 12 requirements for these items, for Test 2 and 3, need to be looked into further, as in some cases the requirements were exceeded
- It would be useful to run multiple replications of the simulations, to get better estimates of the results, especially for those cases that are revisited
  - However, given that the results were very similar for case with and without residence time and Pdelay turnaround time at nodes 1 and 2, it is not necessary to run simulations for both sets of cases
  - It should simply be decided whether to include or not include residence time and Pdelay turnaround time at nodes 1 and 2, and run only those cases
  - This means that multiple replications need to be run only for 3 cases
  - Since the 2000 s runs for cases 2 and 3 required a few seconds to run, and the 100,000 s run for case 1 required 9 minutes (this included post-processing), it should be possible to run many more than the 300 replications run for full networks (i.e., 101 nodes)

Thank you

References - 1

[1] IEC/IEEE 60802 Time-Sensitive Networking Profile for Industrial Automation, Draft 2.1, October 2023.

[2] David McCall, *IEC/IEEE 802 Contribution – Annex D (Informative), Time Synchronization Informative Annex, Version 9*, October 2023 (available at https://www.ieee802.org/1/files/public/docs2023/60802-McCall-Time-Sync-Informative-Annex-Clean-1023-v09.pdf)

[3] David McCall, *60802 Time Sync – Error Generation Normative Requirements & Next Steps, Version 2*, IEC/IEEE 60802 presentation, June 2023, (available at https://www.ieee802.org/1/files/public/docs2023/60802-McCall-Error-Generation-Normative-Requirements-1123-v02.pdf).

[4] Geoffrey M. Garner, *New Multiple Replication 60802 Time Domain Simulation Results for Cases with Drift Tracking Algorithms and PLL Noise Generation*, Revision 1, October 20, 2023 (available at https://www.ieee802.org/1/files/public/docs2023/60802-garner-new-time-domain-simul-results-with-drift-tracking-algorithms-and-PLL-noise-generation-multiple-replic-1023-v01.pdf)

References - 2

[5] ITU-T Rec. G.8251, *The control of jitter and wander within the optical transport network*, ITU-T, Geneva, November 2022

[6] ITU-T Series G Supplement 65, *Simulations of transport of time over packet networks*, ITU-T, Geneva, October 2018

[7] John Rogers, Calvin Plett, Foster Dai, *Integrated Circuit Design for High-Speed Frequency Synthesis*, Artech House, 2006.