

60802 Time Sync – Monte Carlo Simulations with RR & NRR Drift Tracking and Compensation – Results Part 2

David McCall – Intel Corporation

Version 1

References

- 1 – David McCall “[60802 Time Synchronisation – Monte Carlo Analysis: 100-hop Model, “Linear” Clock Drift, NRR Accumulation – Overview & Details, Including Equations – v2](#)”, contribution to IEC/IEEE 60802, September 2022
- 2 – David McCall “[60802 Time Sync – Monte Carlo Simulations with RR & NRR Drift Tracking and Compensation –Initial Results – v2](#)”, contribution to IEC/IEEE 60802, July 2023
- 3 – David McCall & Kevin Stanton “[60802 Dynamic Time Sync Error – Error Model & Monte Carlo Method Analysis](#)”, contribution to IEC/IEEE 60802, November 2021


Contents

- Initial Results, Hypotheses
- More Simulation Results – Questions & Answers
- Summary & Proposed Configuration for Time Series Simulations

Initial Results, Hypotheses and Questions


Default: Quarter-Sinusoidal; 125 s Temp Ramp; GM Scale 1

How true is this?




- Mostly a Normal distribution, but with extra-long tails
- Measured Mean: 0.33
- Measured σ : 171.58
- Min/Max -1300 to 1000 ns (approx.)
- 100,000 Runs

Normal distribution with matching mean and standard deviation




- Main part of bell-curve is slightly wider than DTE_SUM but with narrower tails
- Mean: 0.33
- σ : 171.58
- Min/Max ± 870 ns (approx)
- 100,000 samples

Normal distribution with matching mean and standard deviation



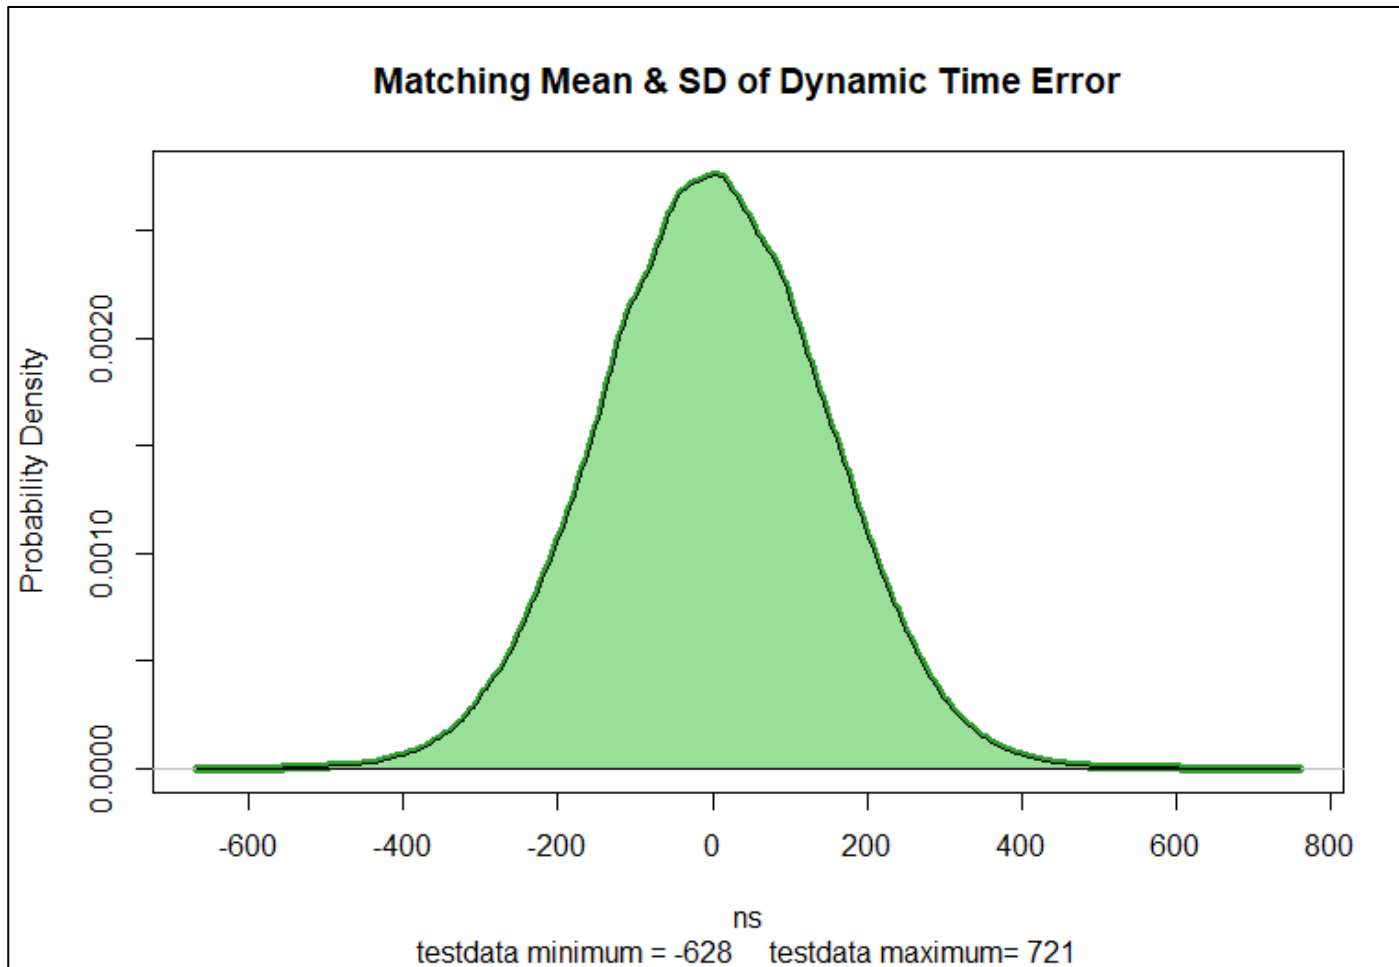
- Plotted on top of DTE_SUM.

Normal distribution to fit with main part of DTE_SUM bell curve




- Mean: 0.33
- σ : 144.5
- 100,000 samples

Normal distribution to fit with main part of DTE_SUM bell curve



- Mean: 0.33
- σ : 144.5
- Min/Max ± 670 ns (approx)
- 100,000 samples

Default: Quarter-Sinusoidal; 125 s Temp Ramp; GM Scale 1





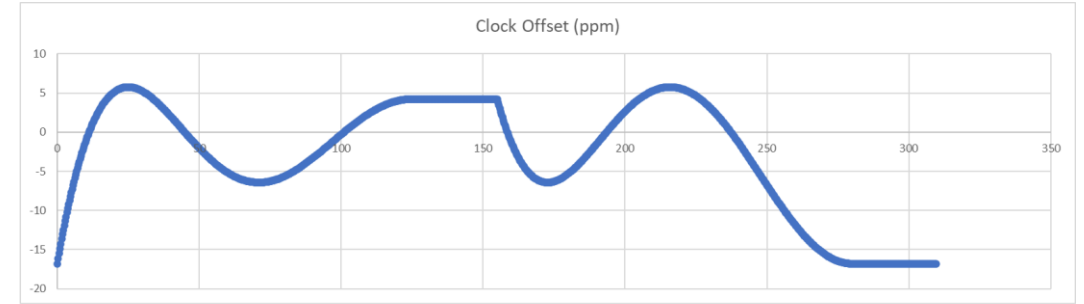
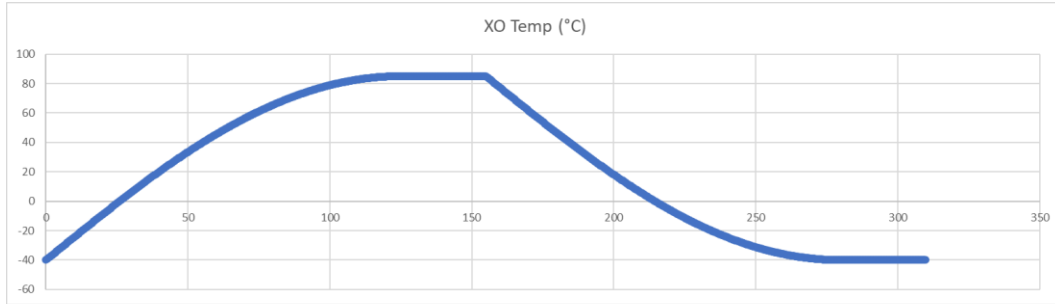
- **True:** mostly a Normal distribution, but with extra-long tails
- Specifically the main normal distribution would suggest Min/Max of ± 670 ns rather than -1300 to +1000 ns (approx)

More Simulation Results – Questions & Answers

Hypothesis: ppm/s discontinuities are causing the “long tails”

- The “Quarter Sinusoidal” temperature ramp includes two points which transition instantaneously from zero ppm/s to either maximum positive or negative ppm/s drift.
 - Clock Drift maximum: 2.12 ppm/s; minimum -1.35 ppm/s
- If the GM starts drifting immediately after sending a Sync message, the rest of the chain will be entirely unaware until the next Sync message.
 - Approx 500 ms for the Sync message to pass down the chain to the End Station, plus 125 ms until the next Sync message arrives at the End Station
- The tracking & compensation algorithms at Relay and End Station instances take up to 2 seconds to respond to changes in ppm/s.

Clock Drift Example – Quarter-Sinusoidal Temperature Ramp: 125s \updownarrow



Inputs	
Temp Max	85°C
Temp Min	-40°C
Temp Ramp Period	125s
Temp Hold	30s



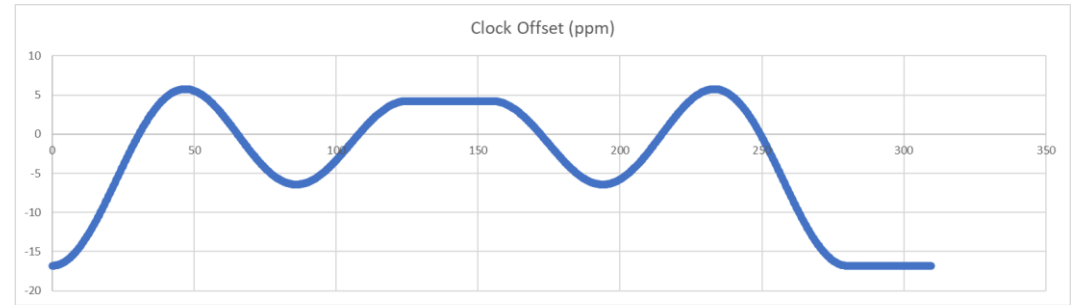
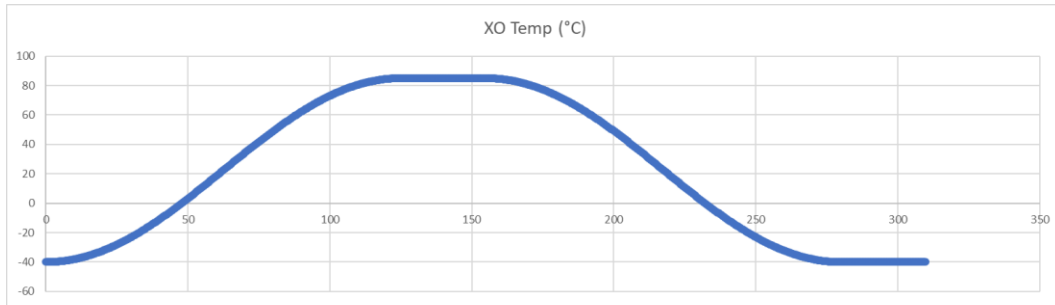
Temp Rate of Change	
MAX	1.57°C/s
MIN	-1.57°C/s

Clock Drift	
MAX	2.12ppm/s
MIN	-1.35ppm/s

Hypothesis: a temperature curve without discontinuities will remove the long tails

- The “Half Sinusoidal” temperature ramp starts and ends the ramp slowly
 - But, for same 125s ramp, Clock Drift maximum: 0.76 ppm/s; minimum -0.76 ppm/s.
 - May need to have faster ramp to match ± 1 ppm/s normative requirement?

Clock Drift Example – Half-Sinusoidal Temperature Ramp: 125s \updownarrow




Inputs	
Temp Max	85°C
Temp Min	-40°C
Temp Ramp Period	125s
Temp Hold	30s

Temp Rate of Change	
MAX	1.57°C/s
MIN	-1.57°C/s

Clock Drift	
MAX	0.76ppm/s
MIN	-0.76ppm/s

Half-Sinusoidal; 125 s Temp Ramp; GM Scale 1




- Switching to Half-Sinusoidal Temp Ramp (no ppm/s discontinuities) minimises / removes the long tails.
- Min/Max ± 615 ns (approx)

Hypothesis: using a GM with lower drift will further reduce time sync error

- GM clock drift has an outsized effect on time sync error
 - Consequent error doesn't tend to cancel out node-to-node as it does for errors due to clock drift at Relay instances
 - Resulting error in mRR survives and “pushes” time sync error off, in the same direction, down the entire chain of nodes, including the End Station.
- GMscale of 0.5 (equivalent of GM oscillator $\pm 25\text{ppm}$ vs usual $\pm 50\text{ppm}$) results in simulation using half the error due to clock drift at the GM

Half-Sinusoidal, 125 s Temp Ramp, GMscale 0.5




- Marginal improvement to Min/Max ± 608 ns (approx)
- Clock Drift error from other nodes and other errors may be dominating?

Hypothesis: if the temperature ramp is less aggressive, time sync error will be further reduced

- Easiest way to do this is to increase the ramp time from 125s to, for example, 250s

Half-Sinusoidal, 250 s Temp Ramp GMscale 0.5



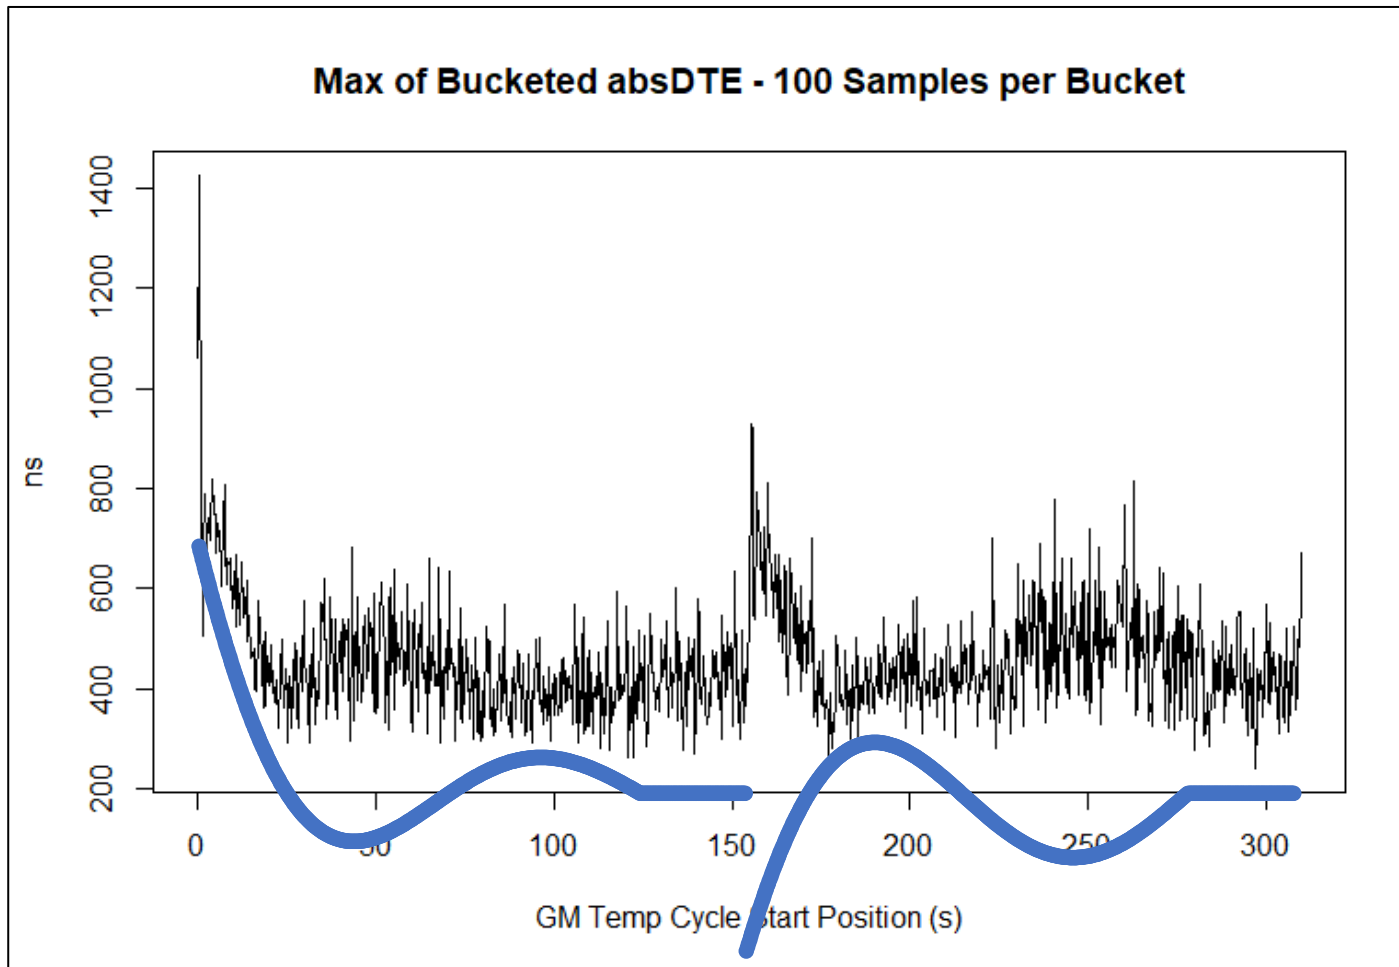
- Noticeable improvement!
- Min/Max ± 480 ns (approx)
- But Min/Max ± 0.38 ppm/s drift

More Results – Questions & Answers

Question: Are the long tails really due to the discontinuities?

- Investigate by...
 - Simulation where GM Temp Cycle start position is not random, but instead in sequence across Temp Cycle range, equally spaced
 - 100,000 runs across 310 s means, for each run, GM start point is incremented by 3.1 ms
 - Bucket absolute DTE_SUM results with each bucket containing 100 results
 - Find MAX value of dTE in each bucket
 - Plot MAX values over time
- Default values: Quarter-Sinusoidal; 125 s Temp Ramp; GM Scale 1

Question: Are the long tails really due to the discontinuities?

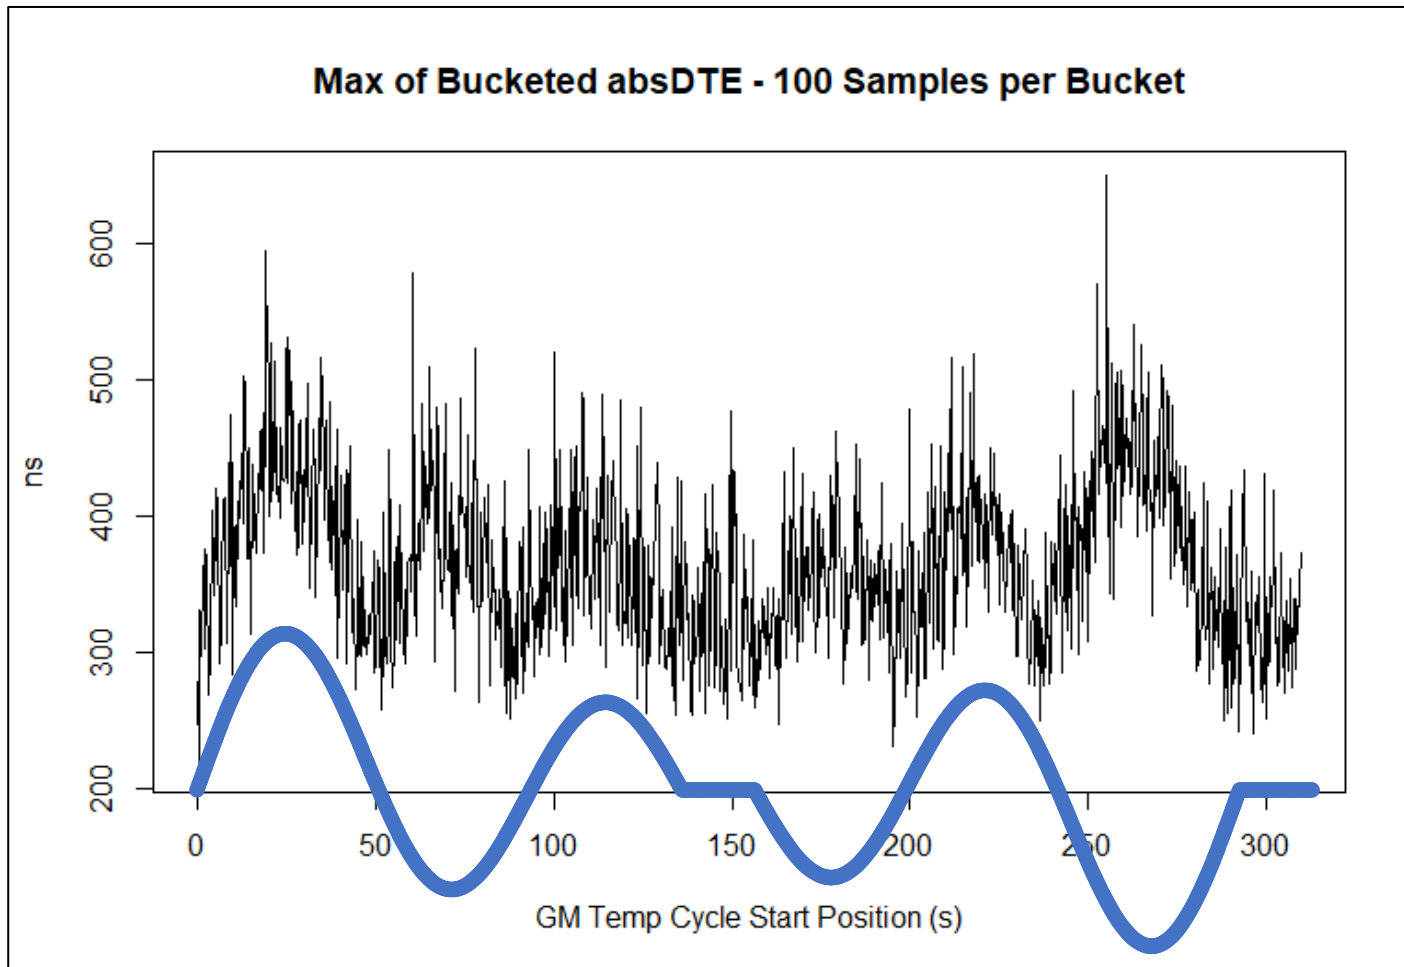


- Yes. Spikes immediately after discontinuities are clear.
- Errors are higher in period after discontinuity due to high ppm/s drift
- Useful to overlay ppm/s drift curve

Question: What does the equivalent graph look like for the Half-Sinusoidal Temp Ramp?

- Repeat with Half-Sinusoidal temp ramp
 - Half-Sinusoidal; 125 s Temp Ramp; GM Scale 1

Question: What does the equivalent graph look like for the Half-Sinusoidal Temp Ramp?




- Large spikes are absent.
- Again, useful to overlay ppm/s drift curve

Question: How much does a better oscillator at the GM buy you?

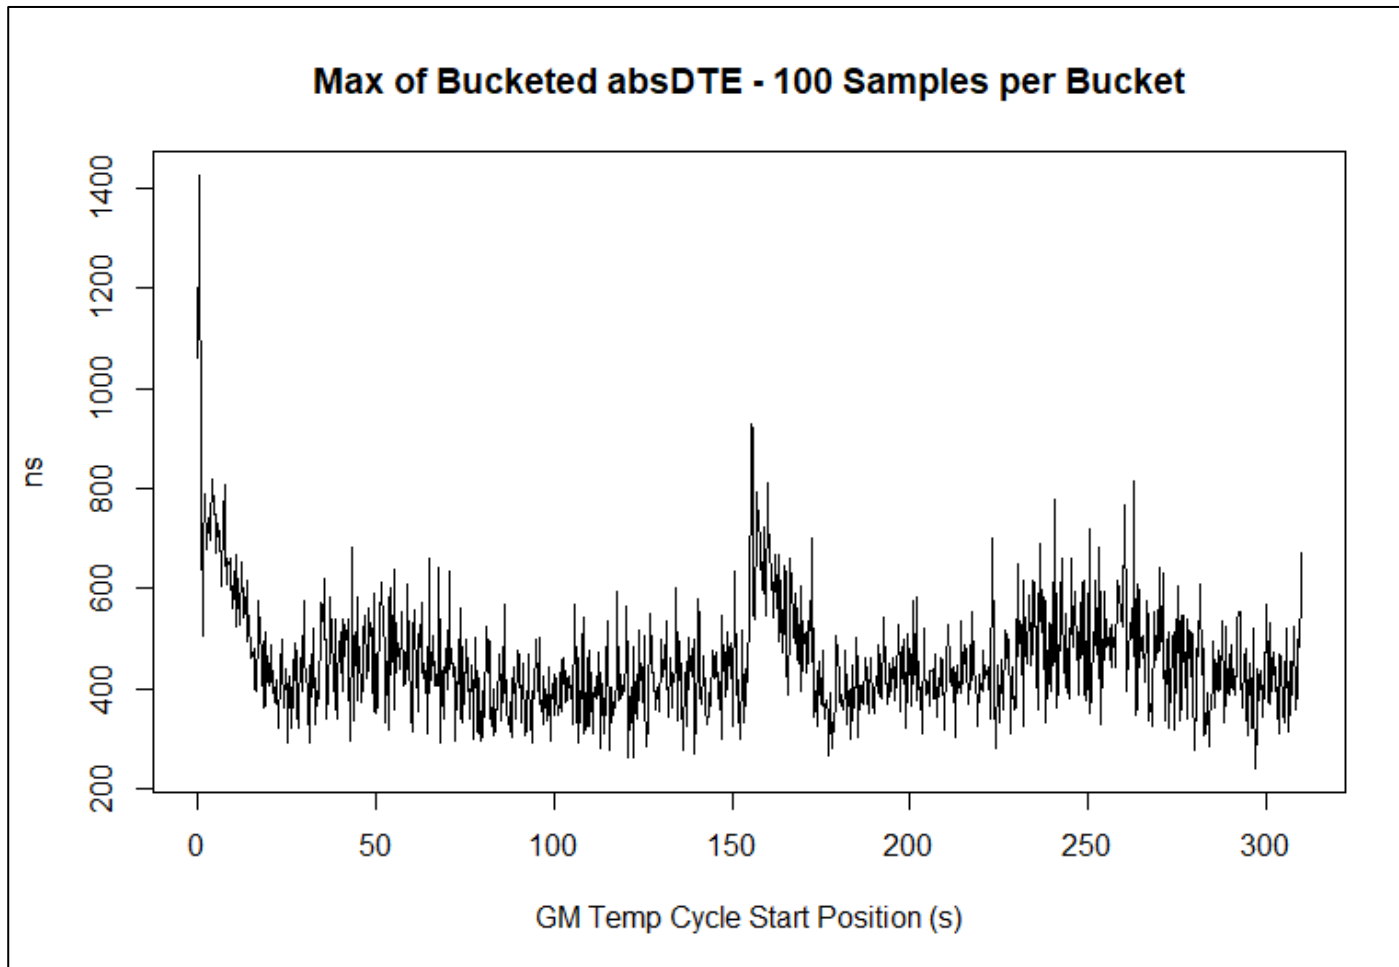
- Run the same analysis for both Quarter- and Half-Sinusoidal temperature ramps, but GM Scale of 0.5 (equivalent to ± 25 ppm offset range).

Default: Quarter-Sinusoidal; 125 s Temp Ramp; GM Scale 1




- Rerun with GM temp cycle start positions in order, equally space, across Temp Cycle
 - May be slightly different than random positions
- Min/Max -1400 to +900 ns (approx.)

Default: Quarter-Sinusoidal; 125 s Temp Ramp; GM Scale 1




- When GM clock is continuous (no discontinuities), worst case absolute error is between 600 and 800 ns

Quarter-Sinusoidal; 125 s Temp Ramp; GM Scale 0.5




- Min/Max -1100 to +800 ns (approx.)
 - Min improved by 300 ns; Max improved by 100 ns

Quarter-Sinusoidal; 125 s Temp Ramp; GM Scale 0.5




- When GM clock is continuous (no discontinuities), worst case absolute error is between 600 and 800 ns
 - No significant change from GM Scale 1

Half-Sinusoidal; 125 s Temp Ramp; GM Scale 1




- Min/Max -650 to +590 ns (approx.)

Half-Sinusoidal; 125 s Temp Ramp; GM Scale 1




Half-Sinusoidal; 125 s Temp Ramp; GM Scale 0.5



- Min/Max -650 to +600 ns (approx.)
 - No significant change

Half-Sinusoidal; 125 s Temp Ramp; GM Scale 0.5



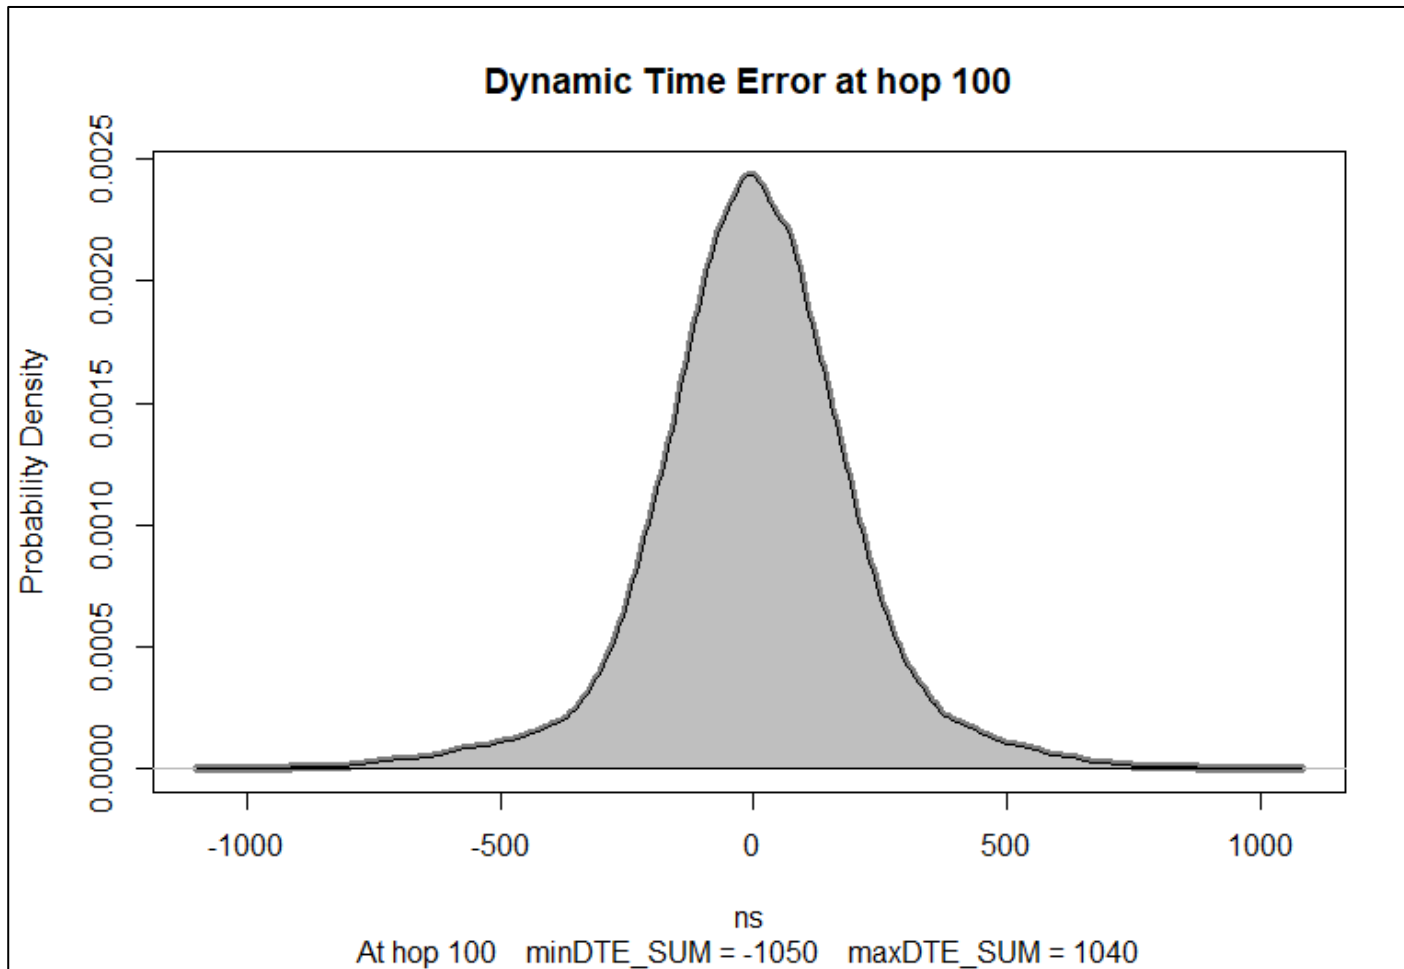
Conclusion: a better GM oscillator mainly buy improvements when there are discontinuities

- Discontinuities where the worst cases occur. A better crystal reduces the size of the discontinuities.
- If temperature ramp doesn't include discontinuities, the benefit is minimal.
- For the configurations we're looking at, discontinuities cause Time Sync error to exceed goal (dTE worst case is greater than total TE budget)

Question: how much does NRR & RR drift tracking and error compensation buy you?

- Run the same analysis for both Quarter- and Half-Sinusoidal temperature ramps, but without the tracking and compensation algorithms.
 - Quarter-Sinusoidal with GM Scale 0.5
 - Half-Sinusoidal with GM Scale 1

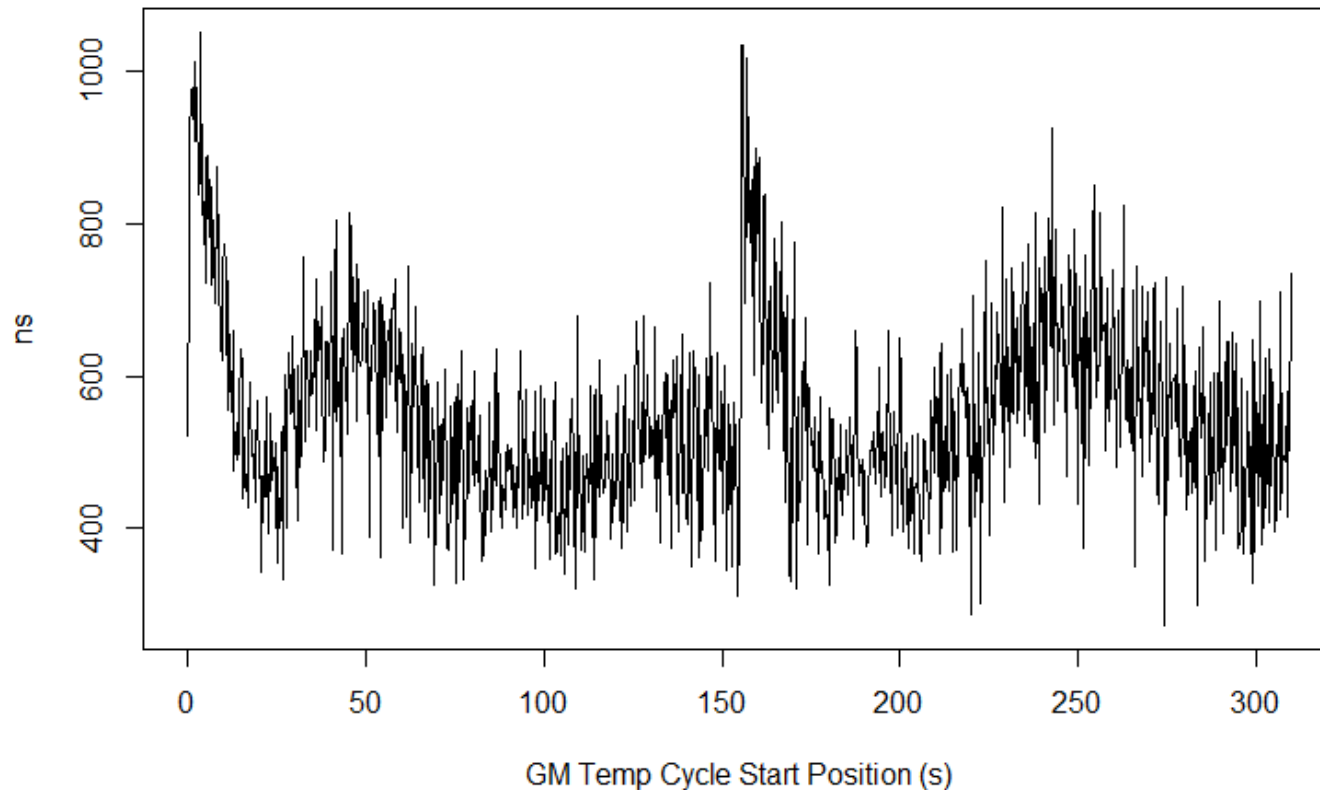
Quarter-Sinusoidal; No NRR/RR Tracking; 125 s Temp Ramp; GM Scale 0.5



- Min/Max ± 1050 ns
 - With tracking was -1100 to +800 ns


Quarter-Sinusoidal; No NRR/RR Tracking; 125 s Temp Ramp; GM Scale 0.5

Max of Bucketed absDTE - 100 Samples per Bucket



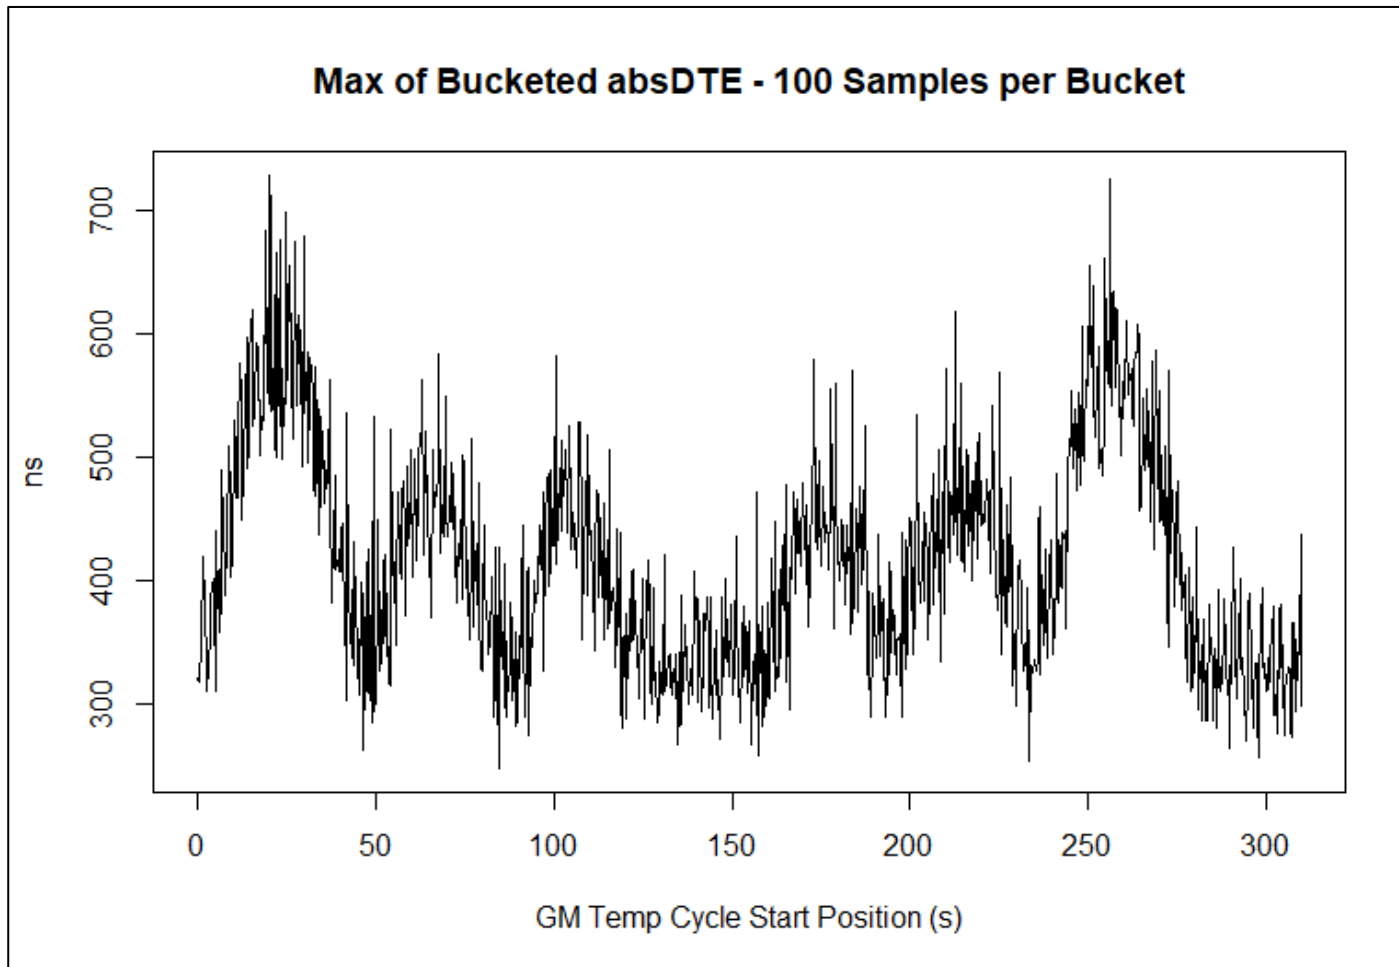
- Appears to be 100 to 200 ns worse across the entire GM Temp Ramp

Half-Sinusoidal; No NRR/RR Tracking; 125 s Temp Ramp; GM Scale 1



- Min/Max ± 725 ns (approx.)
 - With tracking was -650 to +590 ns

Half-Sinusoidal; No NRR/RR Tracking; 125 s Temp Ramp; GM Scale 1




- Appears to be 100 to 150 ns worse across the entire GM Temp Ramp

Question: Switching off NRR / RR tracking and error compensation makes things worse. Can increasing it make things better?


- Run simulations with mNRRcompNAP = 6 and then 8
 - Half-Sinusoidal; 125 s Temp Ramp; GM Scale 1

Half-Sinusoidal; mNRRcompNAP 6; 125 s Temp Ramp; GM Scale 1




- Min/Max -550 to +610 ns (approx.) a 40 ns (approx.) improvement.
 - With mNRRcompNAP 4 was -650 to +590 ns

Half-Sinusoidal; mNRRcompNAP 6; 125 s Temp Ramp; GM Scale 1




- Confirms approx. 50 ns improvement across the board.

Half-Sinusoidal; mNRRcompNAP 8; 125 s Temp Ramp; GM Scale 1



- Min/Max -539 to +550 ns (approx.) an additional 20 to 50 ns (approx.) improvement.
 - With mNRRcompNAP 6, was -550 to +610 ns

Half-Sinusoidal; mNRRcompNAP 8; 125 s Temp Ramp; GM Scale 1




- Does appear to shave off some of the more extreme errors.

Question: Any benefit to increasing mNRRsmoothingNA?


- Run simulations with mNRRsmoothing = 6 and then 8
 - Half-Sinusoidal; 125 s Temp Ramp; GM Scale 1; mNRRcompNAP 8

Half-Sinusoidal; mNRRcompNAP 8;
mNRRsmoothingNA 6; 125 s Temp Ramp; GM Scale 1




- Min/Max -518 to +538 ns (approx.) a 15 to 20 ns (approx.) improvement.
 - With mNRRsmoothingNA 4 was -539 to +550 ns

Half-Sinusoidal; mNRRcompNAP 8;
mNRRsmoothingNA 6; 125 s Temp Ramp; GM Scale 1




- Confirms approx. 50 ns improvement across the board.

Half-Sinusoidal; mNRRcompNAP 8;
mNRRsmoothingNA 8; 125 s Temp Ramp; GM Scale 1



- Min/Max -515 to +554 ns (approx.) which is slightly worse, although probably within margin of error
 - With mNRRsmoothingNA 6 4 was -518 to +538 ns

Half-Sinusoidal; mNRRcompNAP 8;
mNRRsmoothingNA 8; 125 s Temp Ramp; GM Scale 1



- No visible difference vs. mNRRsmoothing 6


Simulations Suggest Optimised Configuration...

- mNRRcompNAP 8, so for NRR drift measurement...
 - For first measurement...
 - Look back 8 Sync messages for initial NRR calculation (N)
 - Take average of past 8 calculations (A)
 - For second measurement, do the same thing but...
 - Start 16 Sync messages in the past (P)
- mNRRsmoothingNA 8, so for compensated mNRR measurement...
 - Look back 8 Sync messages for initial calculation (N)
 - Take average of past 8 calculations
 - Allows reuse of initial calculations from NRR drift measurement

Question: Do these settings make the system less robust against potential discontinuities?


- mNRRcompNAP 8 means using data from past 32 Sync messages, i.e. past 4 seconds (with Sync Interval 125 ms)
- Although the simulation with Half-Sinusoidal temp ramp does not include discontinuities, there may be concerns that this optimisation makes the system less robust against behaviours the simulation does not fully cover.
- Run simulation with optimised settings, but Quarter-Sinusoidal temperature ramp and GM Scale 0.5, and see how it compares with initial simulation results

Quarter-Sinusoidal; mNRRcompNAP 8;
mNRRsmoothingNA 8; 125 s Temp Ramp; GM Scale 0.5



- Min/Max -2150 to +1500 ns (approx.) which is approx. 2x worse than with mNRRcompNAP 4 & mNRRsmoothingNA 4
 - with mNRRcompNAP 4, mNRRsmoothingNA 4 was -1100 to +800 ns

Quarter-Sinusoidal; mNRRcompNAP 8; mNRRsmoothingNA 8; 125 s Temp Ramp; GM Scale 0.5



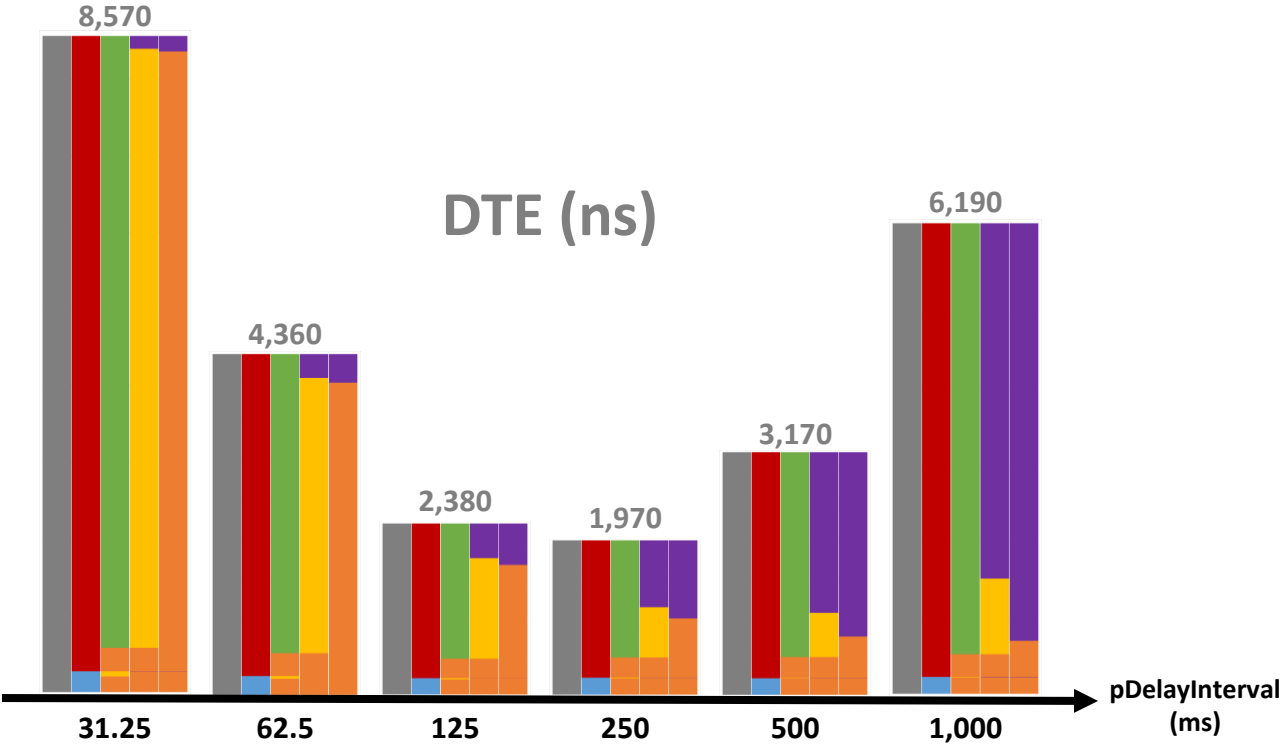
- Increase appears to be driven entirely by spikes following discontinuities
- But spikes appeared to be half the size (and the same size) for both...
 - mNRRcompNAP 4
mNRRsmoothingNAP 4
 - mNRRcompNAP 4
mNRRsmoothingNAP 0 (no compensation)
- So...what is going on?

Question: What is going on?

- Key observation: doubling `mNRRsmoothingNA` (from 4 to 8) almost exactly doubled the size of the spikes in Time Sync Error following a discontinuity.
- Increasing `mNRRsmoothingNA` is (almost) equivalent to increasing the `Pdelay` Interval when using `Pdelay_Resp` to measure NRR as far as the balance between error due to Timestamp vs. error due to Clock Drift
 - See slide 72 from [3] (way back in November 2021)...

From [3]

pDelayInterval Sensitivity Analysis




Question: What is going on?

- Key observation: doubling `mNRRsmoothingNA` (from 4 to 8) almost exactly doubled the size of the spikes in Time Sync Error following a discontinuity.
- Increasing `mNRRsmoothingNA` is (almost) equivalent to increasing the `Pdelay Interval` when using `Pdelay_Resp` to measure NRR as far as the balance between error due to Timestamp vs. error due to Clock Drift
 - See slide 72 from [3] (way back in November 2021)
- Balance is different when drift tracking and error compensation algorithm reduces the relative contribution to overall time error from clock drift (purple in the diagram), but...
- Immediately after a discontinuity the drift tracking and error compensation algorithm is ineffective and behaviour is dictated by the amount of error due to clock drift
 - If the interval (now Sync Interval, not `Pdelay Interval`) and `mNRRsmoothingNA` means that error due to clock drift is comparable to error from Timestamp error, doubling `mNRRsmoothingNA` will double Time Sync Error.

Hypothesis:


- When the drift tracking and compensation algorithm is working well, larger mNRRsmoothingNA values lead to lower time sync error.
 - Diminishing returns for drift tracking accuracy beyond mNRRcompNAP 8
 - For mNRRcomp8, values of mNRRsmoothing 8 or 6 appear optimal (at least for the temperature ramp we are using)
- When the drift tracking and compensation algorithm isn't working well, larger mNRRsmoothingNA lead to higher time sync error.
- Question: is there a lower limit for mNRRsmoothingNA below which Timesync Error dominates?
 - Run additional simulations to investigate.

Quarter-Sinusoidal; No NRR/RR Tracking; mNRRsmoothingNA 3; 125 s Temp Ramp; GM Scale 0.5




- Min/Max ± 750 ns, an improvement of approx. 300 ns.
 - With mNRRsmoothingNA 4 was ± 1050 ns

Quarter-Sinusoidal; No NRR/RR Tracking;
mNRRsmoothingNA 3; 125 s Temp Ramp; GM Scale 0.5




- Most of the noise curve appears to be 100 to 150ns lower.

Quarter-Sinusoidal; No NRR/RR Tracking; mNRRsmoothingNA 2; 125 s Temp Ramp; GM Scale 0.5



- Min/Max -620 to + 700 ns (approx.), an improvement of approx. 90 ns.
 - With mNRRsmoothingNA 3 was ± 750 ns

Quarter-Sinusoidal; No NRR/RR Tracking;
mNRRsmoothingNA 2; 125 s Temp Ramp; GM Scale 0.5




- No correlation with GM Temp Cycle Start Position, which indicates that Timestamp error is dominating.


Question: What happens if you add Drift Tracking & Error Compensation Back In?

- Run simulations with...
 - Quarter-Sinusoidal; mNRRcompNAP 8; mNRRsmoothingNA 3; 125 s Temp Ramp; GM Scale 0.5
 - Quarter-Sinusoidal; mNRRcompNAP 8; mNRRsmoothingNA 2; 125 s Temp Ramp; GM Scale 0.5

Quarter-Sinusoidal; mNRRcompNAP 8;
mNRRsmoothingNA 3; 125 s Temp Ramp; GM Scale 0.5



Quarter-Sinusoidal; mNRRcompNAP 8;
mNRRsmoothingNA 3; 125 s Temp Ramp; GM Scale 0.5



Summary & Proposed Configuration for Time Series Simulations

Summary

- For Half-Sinusoidal temperature ramp (\updownarrow 125 s) optimal values are...
 - mNRRcompNAP 8
 - mNRRsmoothingNA 8 (or 6, but 8 allows reuse of calculations)
- For increased robustness against any discontinuities...
 - mNRRsmoothingNA 4

...won't incur much of a penalty when there are no discontinuities
- NRR/RR drift tracking and error compensation buys significant benefits
 - From 200 ns cTE error budget to 400+ ns cTE error budget
- Note: this is all “unfiltered” dTE. Time Series simulation required to determine “filtered” dTE.

Recommended Time Series Simulation Configurations

- Based on “default” values in Backup
- Half-Sinusoidal Temp Ramp; mNRRcompNAP 8; mNRRsmoothingNA ?
- Any investigation? Or just one configuration with multiple replications?

Thank you!

“Man with binoculars” icon is from [Icon Fonts](#) under CC BY 3 license.

Backup

Default Configuration

```
hops <- 100 # Minimum 1 hop
runs <- 100000
#
# Input Errors, Parameters & Correction Factors
driftType <- 4 # 1 = DO NOT USE - Historical - Uniform Probability Distribution between MIN & MAX ppm/s
             # 2 = Probability Based on Linear Temp Ramp
             # 3 = Probability Based on Half-Sinusoidal Temp Ramp
             # 4 = Probability Based on Quarter-Sinusoidal Temp Ramp
# Clock Drift Probability from Temp Curve & XO Offset/Temp Relationship
tempMax <- +85 # degC - Maximum temperature
tempMin <- -40 # degC - Minimum temperature
tempRampRate <- 1 # degC/s - Drift Rate for Linear Temp Ramp
tempRampPeriod <- 125 # s - Drift Period for Sinusoidal & Half-Sinusoidal Temp Ramps
tempHold <- 30 # s - Hold Period at MIN and MAX temps before next temp ramp down or up
GMscale <- 1 # Ratio of GM stability vs. standard XO. 1 is same. 0 is perfectly stable.
nonGMscale <- 1 # Ratio of non-GM (and non-ES) node stability vs. standard XO. 1 is same. 0 is perfectly stable.
```

Default Configuration

TSGEtx <- 4 # +/- ns - Error due to Timestamp Granularity on TX

TSGErx <- 4 # +/- ns - Error due to Timestamp Granularity on RX

DTSEtx <- 6 # +/- ns - Dynamic Timestamp Error on TX

DTSErx <- 6 # +/- ns - Dynamic Timestamp Error on RX

pDelayInterval <- 125 # ms - Nominal Interval between two pDelay measurements

PDImax <- 1.3 # Max factor for Tpdelay2pdelay (uniform linear distribution max of pDelayInterval x PDImax)

PDImin <- 0.9 # Min factor for Tpdelay2pdelay (uniform linear distribution min of pDelayInterval x PDImin)

syncInterval <- 125 # ms - Nominal Interval between two Sync messages

SImode <- 3 # Mode for generating Tsync2sync *HARD CODED to MODE 3*

1 = Gamma Distribution, defaulting to 90% of Tsync2sync falling within 10% of the nominal syncInterval. Truncated at SImax (higher values above are reduced to SImax)

No truncation of low values

2 = Gamma Distribution, defaulting to 90% of Tsync2sync falling within 10% of the nominal syncInterval. Truncated at SImax (higher values are reduced to SImax)

Truncated at SImin (lower values are increased to SImin)

3 = Uniform, linear distribution between syncInterval x SImin and syncInterval x SImax

Default Configuration

```
SlScale <- 1 # Scaling factor for Mode 1 & 2 Tsync2sync vs regular distribution.  
          # Scaling factor of 3 would mean 90% of Tsync2sync falling within 30% of the nominal syncInterval  
Slmax <- 1.048 # For mode 1 & 2, Max truncation factor (e.g. 2x syncInterval) limit for Tsync2sync; higher values reduced to Slmax  
          # For mode 3, upper limit of uniform linear distribution  
Slmin <- 0.952 # For mode 1 & 2, Min truncation factor (e.g. 0.5 x syncInterval) limit for Tsync2sync; higher values reduced to Slmin  
          # For mode 3, lower limit of uniform linear distribution  
pDelayTurnaround <- 15 # ms - TpdelayTurnaround maximum; higher values truncated  
pathDelayMin <- 5 # ns - 1m cable = 5ns path delay  
pathDelayMax <- 500 # ns - 100m cable = 500ns path delay  
PDTmin <- 1 # TpdelayTurnaround minimum; lower values truncated  
PDTmean <- 10 # TpdelayTurnaround mean  
PDTsd <- 1.8 # TpdelayTurnaround sigma; 3.4ppm will fall outside 6-sigma either side of the mean  
residenceTime <- 15 # ms - TResidenceTime maximum; higher values truncated  
RTmin <- 1 # TResidenceTime minimum; lower values truncated  
RTmean <- 5 # TResidenceTime mean  
RTsd <- 1.8 # TResidenceTime sigma; 3.4ppm will fall outside 6-sigma either side of the mean
```

Default Configuration

```
mLinkDelayAverage <- 50 # Number of Path Delay calculations, from Pdelay_Req & _Resp messages
    # that are averaged to generate mLinkDelay
mNRRsmoothingNA <- 4 # Whole Number >=1 - Combined N & A value for "smoothing" calculated mNRR (mNRRc)
    # Calculate mNRR using timestamps from Nth Sync message in the past
    # Then take average of previous A mNRRcalculations.
mNRRcompNAP <- 4 # Whole Number >=1
    # For NRR drift rate error correction calculations, take two measurements, mNRRa and mNRRb.
    # Both use timestamps from Nth Sync message in the past, then take average of previous A calculations.
    # Calculation mNRRb starts P calculations in the past from mNRRa, where P = mNRRcompNAP * 2.
    # If 0, there is no NRR drift rate error correction.
```