



The 60802 challenge of meeting time accuracy goals across long daisy-chains using 802.1AS™-2020

An analysis and a proposed path forward

David McCall, Kevin Stanton, Greg Schlechter (Intel)
Jordan Woods, Tom Weingartner (Analog Devices)

September 2021 Interim of 802.1 TSN / 60802

Updates

- **v2 → v3**
 - Fixed minor typos and consistency issues.
 - Updated speaker notes to match revised slide content.
 - Changed top level references from Path Delay, $pDelay_{error}$ and $pDelay_{errorCorrection}$ to Link Delay, $meanLinkDelay_{error}$ and $meanLinkDelay_{errorCorrection}$.
- **v1 → v2**
 - In NRR_{error} and RR_{error} sections, changed to reflect that $pDelay_{resp}$, t_3 & t_4 are used to measure NRR (not $pDelay_{req}$, t_1 & t_2)
 - In NRR_{error} section, fixed calculation of NRR to represent NRR_{AB} (was previously, erroneously, NRR_{BA}). Also changed graphic to make it clear that NRR_{AB} is the ratio measured from A of its clock rate relative to B.
 - Fixed $pDelay_{errorNRR}$ equation (should be /2)
 - Fixed $RR_{errorDrift}$ equations (should be $clockDrift_{GM} - clockDrift(n)$ not $clockDrift(n) - clockDrift_{GM}$)
 - Fixed discussion of NRR_{error} in most recent time series simulations
 - NRR is measured across seven pDelay Intervals prior to taking the median of the past seven measurements. (Previously, only taking the median was discussed.)
 - Completed formulae in Backup and added slides for relevant formulae at the end of each section
 - Formulae in Backup include consideration of the units for each factor and secondary errors which will not be included in the initial model. Formulae in the main presentation do not.
 - Changed some “pDelay” references to “Path Delay” where the latter is more appropriate (i.e. the reference isn't to a pDelay message or used in a formula).

Abstract

Industrial Automation Systems require microsecond-accurate time across long daisy-chains of devices using IEEE Std. 802.1AS™-2020 as specified by IEEE/IEC 60802.

Simulated protocol and system parameters have thus far either been judged impractical or have failed to meet the time-accuracy requirement.

In this contribution we:

- 1. Summarize the assumptions, requirements and results to date**
- 2. Challenge certain assumptions**
- 3. Propose an approach that we believe could lead to a solution**

Some Requirements & Assumptions to Date

- **Requirements:**

- Topology: 64-100 network hops (i.e. 2 endstations + 63-99 devices in a chain)
- Time Accuracy: Not to exceed +/-1us with respect to the GM at any node
 - Measured at each Working Clock, e.g. used by 802.1Qbv (i.e. GCL) and the Application Clock

- **Assumptions:**

- → Local Clock reference frequency (in each system) is an XO, i.e. a quartz crystal oscillator
- → The 802.1AS GM does not have better PPM stability than the other nodes
- → No averaging of the pDelay measurement is performed
- Temperature gradients induce PPM gradients, resulting in time-duration-measurement errors
- Various Ambient Temperature profiles (T_{Ambient}) have been used
 - With various transfer functions of $\text{PPM} = f(T_{\text{Crystal}})$
- → Zero thermal resistance between T_{Ambient} and T_{Crystal} (0°K/W) ←

Note: Assumptions that we recommend revisiting are marked with “→”

Other Principles

- **Initial 60802 standard should not trigger new silicon design / spins**
 - Requirements should be met by commodity switch and endpoint silicon today that were designed to support IEEE Std. 802.1AS™-2020
- **The solution should not unnecessarily limit the deployment of IEEE Std. 1588™-2019 security mechanisms**
 - With today's 802.1AS-2020-capable silicon
- **Once viable 802.1AS-2020 parameters are chosen, the worst-case PPM/sec allowed per time-aware system should be derived**
- **The function of PPM vs T_{ambient} is a property of the electrical, mechanical and thermal design of the board / chassis / system**
 - Thus the PPM / T_{Ambient} requirement should be specified, but not the method by which the PPM / T_{Ambient} requirement must be met by the system design

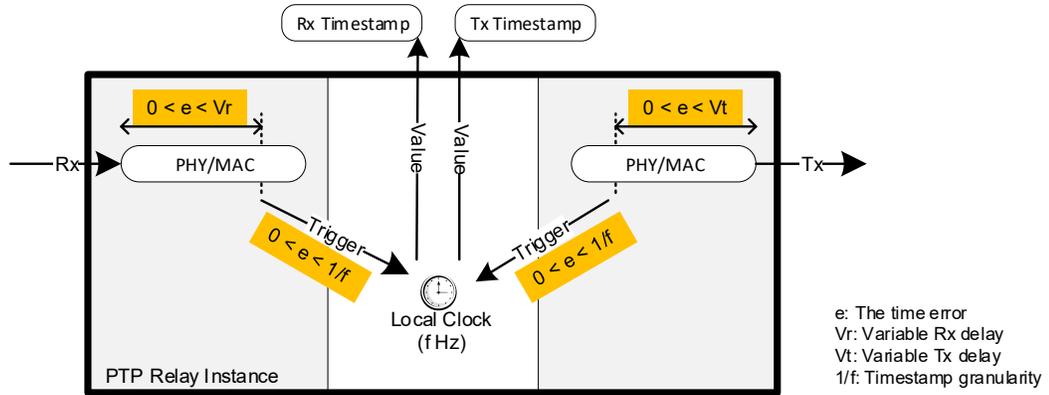
Viabie Parameter Selections

- Sync Interval
- Sync-to-Follow-up delay
- pDelay Interval
- Timestamp Granularity
- Rx/Tx Timestamp Asymmetry:
- Residence Time (in addition to Sync-to-Follow-up delay)

→ A simple analytical Model is needed to more quickly choose scenarios for more detailed time-series simulation

Quickly iterate a Model over minutes, then invest the >1 week for the detailed simulation

Variable Time Error (vTE): Timestamp Model



Motivation & Results (so far)

- Develop a method of iterating faster
 - Input: key 802.1AS parameters
 - Output: time sync accuracy over 64 and 100hops
- Approach by modelling the errors and how they accumulate
 - Requires an analysis of the errors first
- Analysis of the errors has already suggested some assumptions inherent in the current time series simulation that should be revisited
- This presentation walks through the error analysis and the most important insights it reveals
 - Goal is to introduce the overall approach share the insights
 - There are some areas that require further analysis; will present in a later session
- Next steps:
 - Build the model
 - Complete the additional analysis

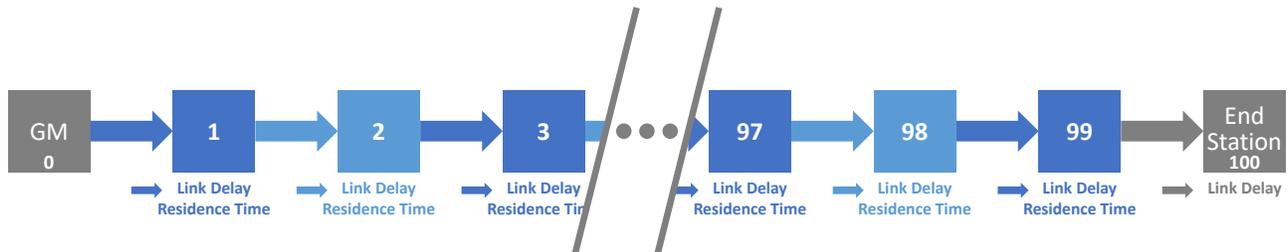
This is a very graphical presentation, but the uploaded version includes extensive speaker notes (which you are reading now!).

There is a lot of information in this presentation. The main section I'm about to go through is nearly over 100 slides, but many of them are builds, so we should get through them fairly quickly.

Summary of Main Insights

- meanLinkDelay should be an average of pDelay measurements.
 - Should nearly eliminate meanLinkDelay_{error}
- NRR should not be averaged.
 - Provided (Timestamp Errors / pDelay Interval) is negligible in terms of ppm
- There are significant benefits to the GM having greater ppm stability (lower ppm/s).
- It is hard (impossible?) to eliminate Timestamp granularity and Dynamic Timestamp error effects from Residence Time
 - Need more work to see how far these elements can realistically be reduced
- Synchronising pDelayReq to arrive just before Sync would minimise the effect Clock Drift induced errors it NRR and RR.
- There may be opportunities to compensate for NRR_{error}
 - And therefore RR_{error}

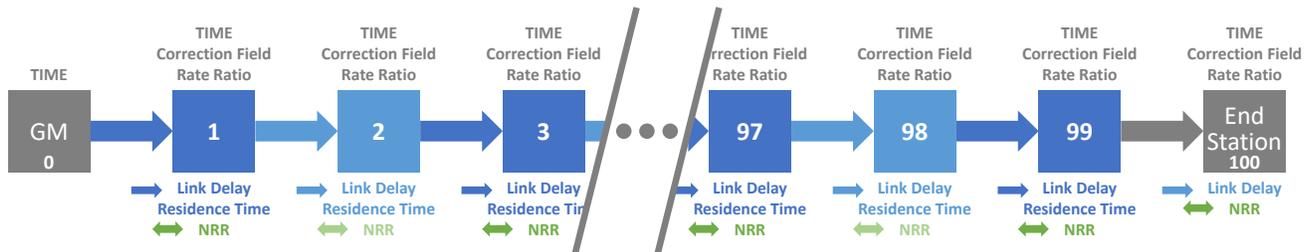
Time Sync



This is a simple view of how time sync works.

The basics are: measure Link Delays and Residence Times; add them to a Sync message as it is passed from the GM to the End Station.

Time Sync



IEEE 802.1 TSN Interim, Virtual, September 2021

Available at <http://www.ieee802.org/1/files/public/docs2021>

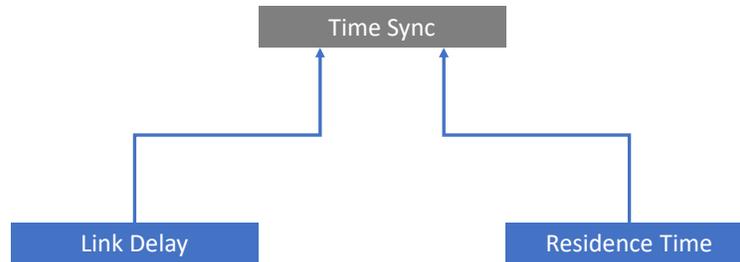
McCall, Stanton with Woods, Weingartner, and Schlechter p. 11

Adding a bit more detail, this is how that's done in practice.

The Sync message contains the TIME of transmission from the GM, a Correction Field of the delay that has accumulated since it was transmitted, and a measure of the Rate Ratio, i.e. the ratio between the current clock speed and the GM's clock speed.

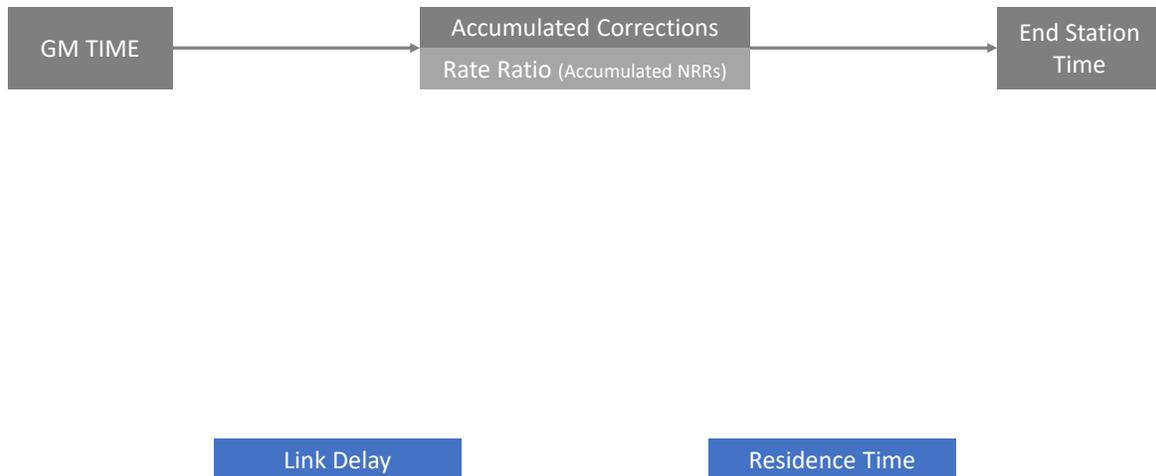
It's worth pointing out that each device is responsible for updating the Correction Field based on the **incoming** Link Delay plus **it's own** Residence Time. Also a device measures the Neighbor Rate Ratio between itself and the upstream device. It then adds that to the incoming Rate Ratio. And finally uses the new Rate Ratio to modify the measured Link Delay and the Residence Time prior modifying the Correction Field, and passing on everything to the next device. Rate Ratio is calculated and used for this step because the Correction Field is in terms of the GM clock, not the local clock – which is used to measure Link Delay and Residence Time - and multiplying by the Rate Ratio accomplishes that translation.

Time Sync – Elements & Relationships



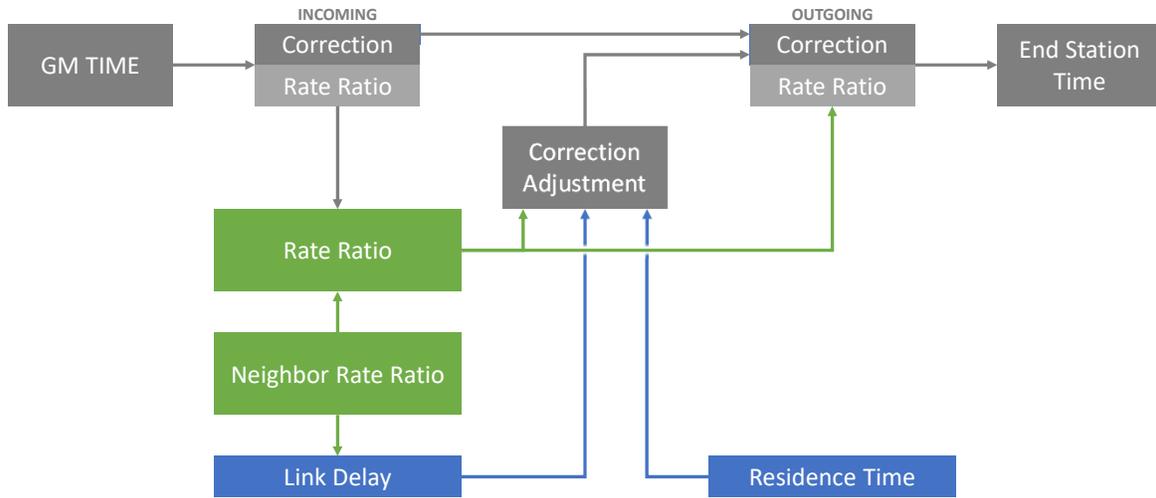
Building up the logical elements and connection of this process...the simple model looks like this...

Time Sync – Elements & Relationships



As mentioned, the Time Sync is actually accomplished by accumulating the delays (i.e. the corrections) and measuring the rate ratio between the GM and the End Station.

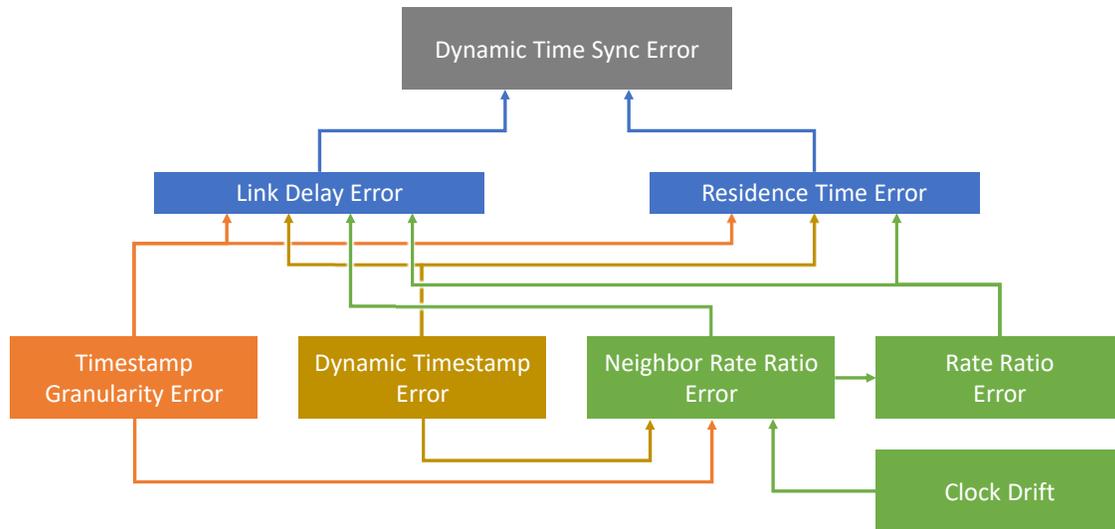
Time Sync – Elements & Relationships



At each bridge there is a process for taking the incoming Correction Field and Rate Ratio and adding effects of the local pDelay (between the local and upstream device) and Residence Time.

I'll be referring back to this diagram later, but when it comes to analysing the errors involved in this process, it's helpful to take a slightly different view...

Time Sync – Elements & Relationships



This is the start of the error analysis.

Timestamp Granularity Error, is related to the maximum resolution of Timestamps.

Dynamic Timestamp Error is in addition to Granularity and can be very implementation dependant.

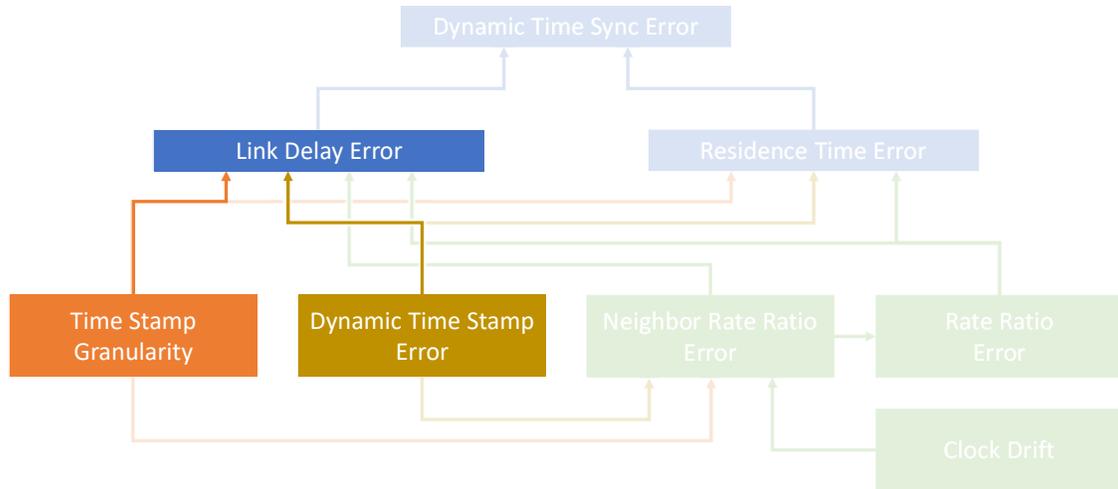
Clock Drift is just the way that each device has it's own local clock which is running a slightly different speed...and also changing that speed over time (mostly due to temperature changes; or at least those are the changes that have the greatest impact on our analysis).

All other errors are all rooted in these three.

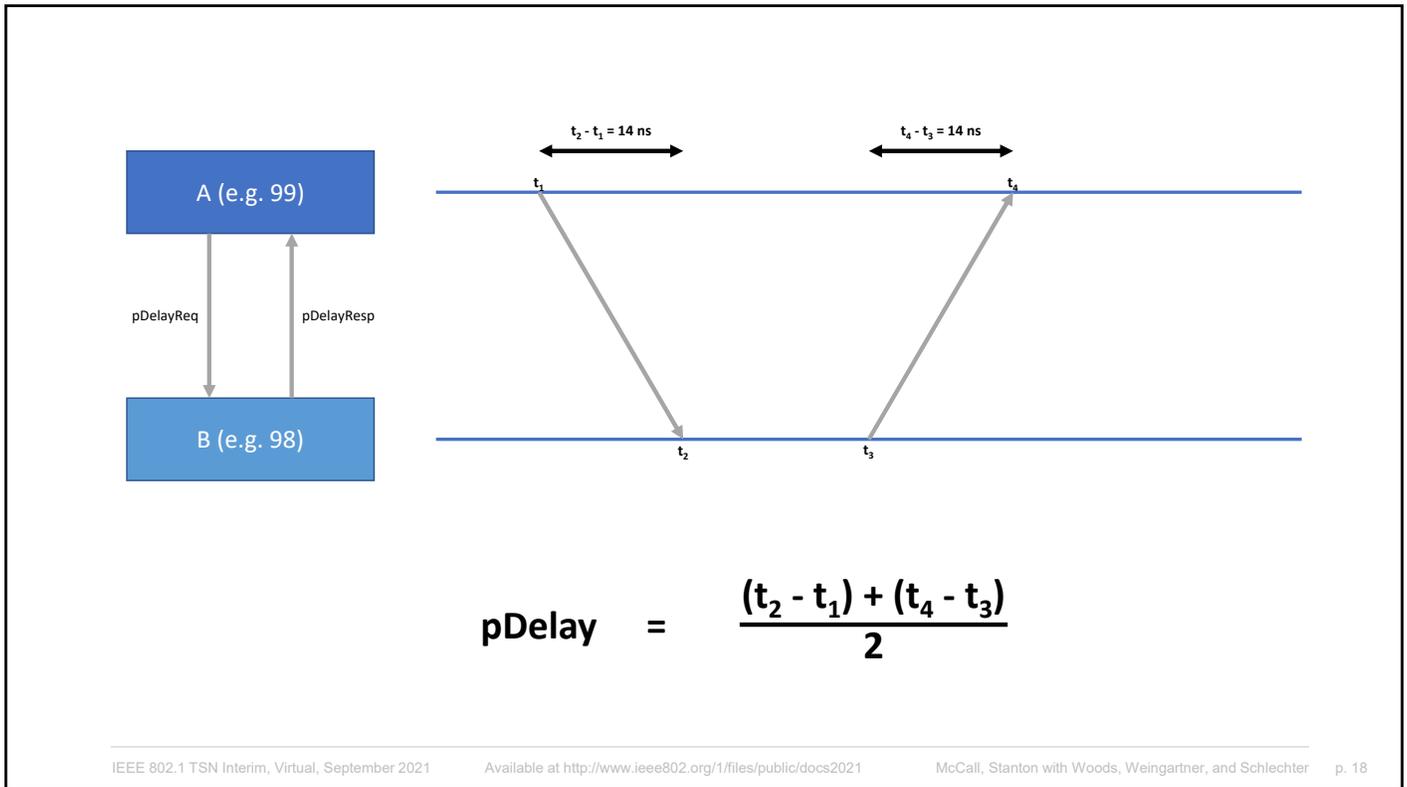
Not all of the relationships need to be modelled. Figuring out which can be ignored is part of the analysis.

Link Delay Error – Timestamp Granularity & Dynamic Time Stamp Error

Time Sync – Elements & Relationships



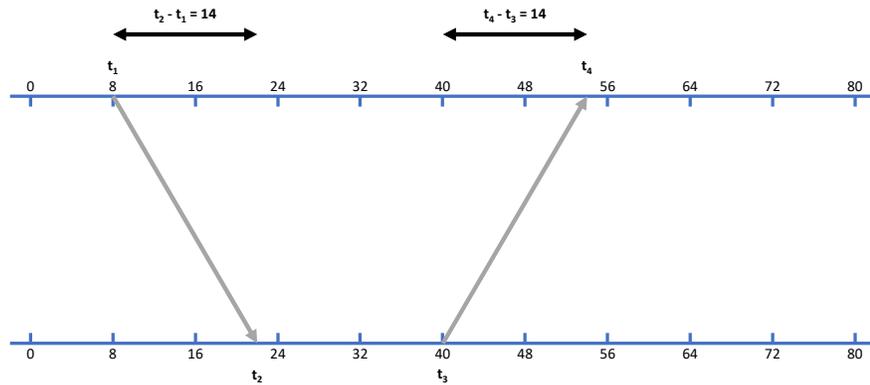
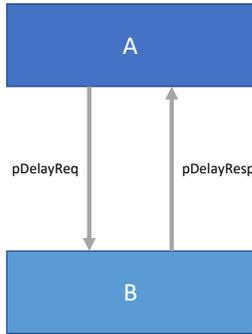
To be clear, we're only looking at these three elements in this section. We aren't looking at any of the others.



This is the idealised view...the Platonic ideal of how pDelay is measured.

Note that “A” is device 99. It initiates a pDelayReq to the upstream device (98)...which responds with a pDelayResp...after the pDelay Response Time.

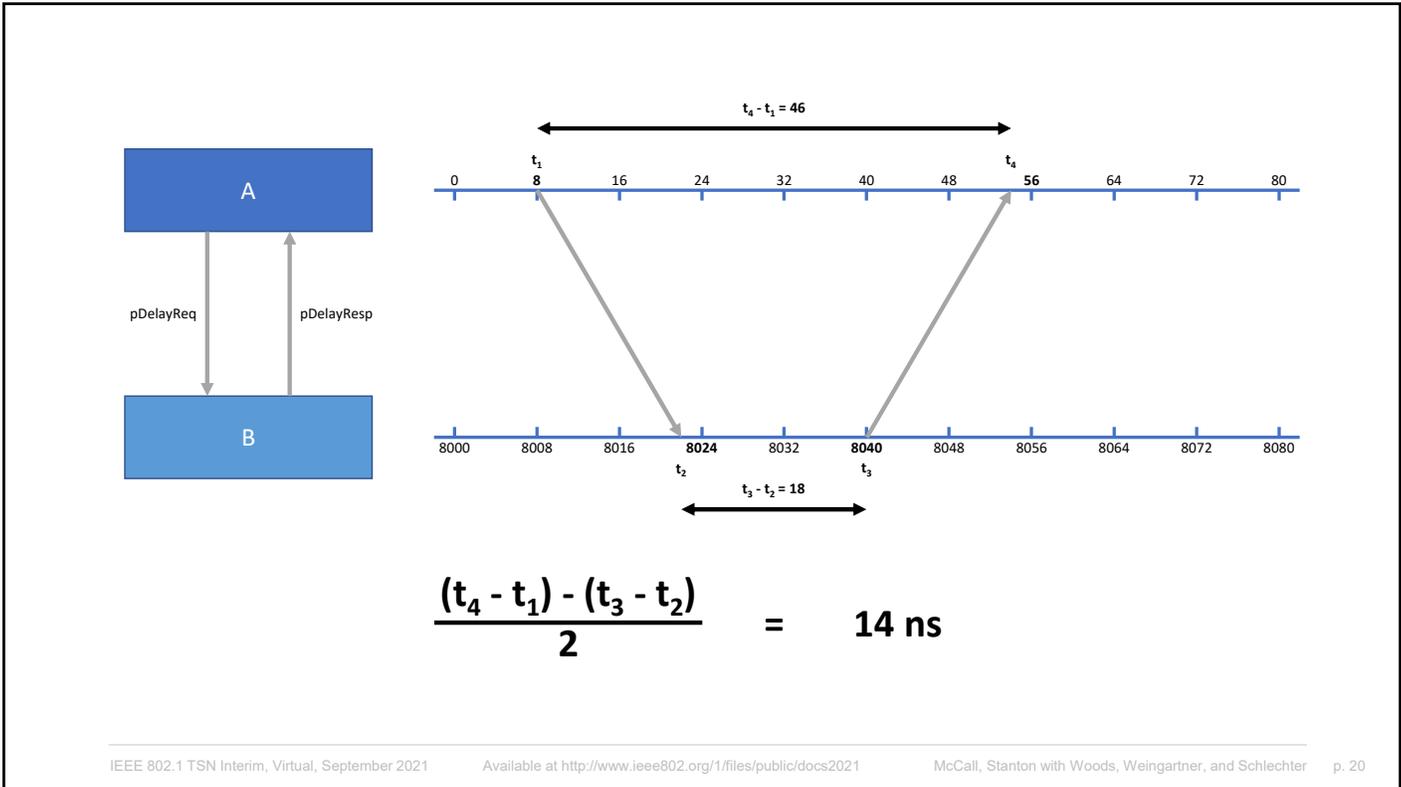
The default maximum for pDelay Response is 10ms, which is much longer than the pDelay which in this case is 14ns...indicating a cable length of 4.2m...but for the moment I’m going to model this shorter pDelay Response Time to make things legible and illustrate some principles...



$$\frac{(t_2 - t_1) + (t_4 - t_3)}{2} = 14 \text{ ns}$$

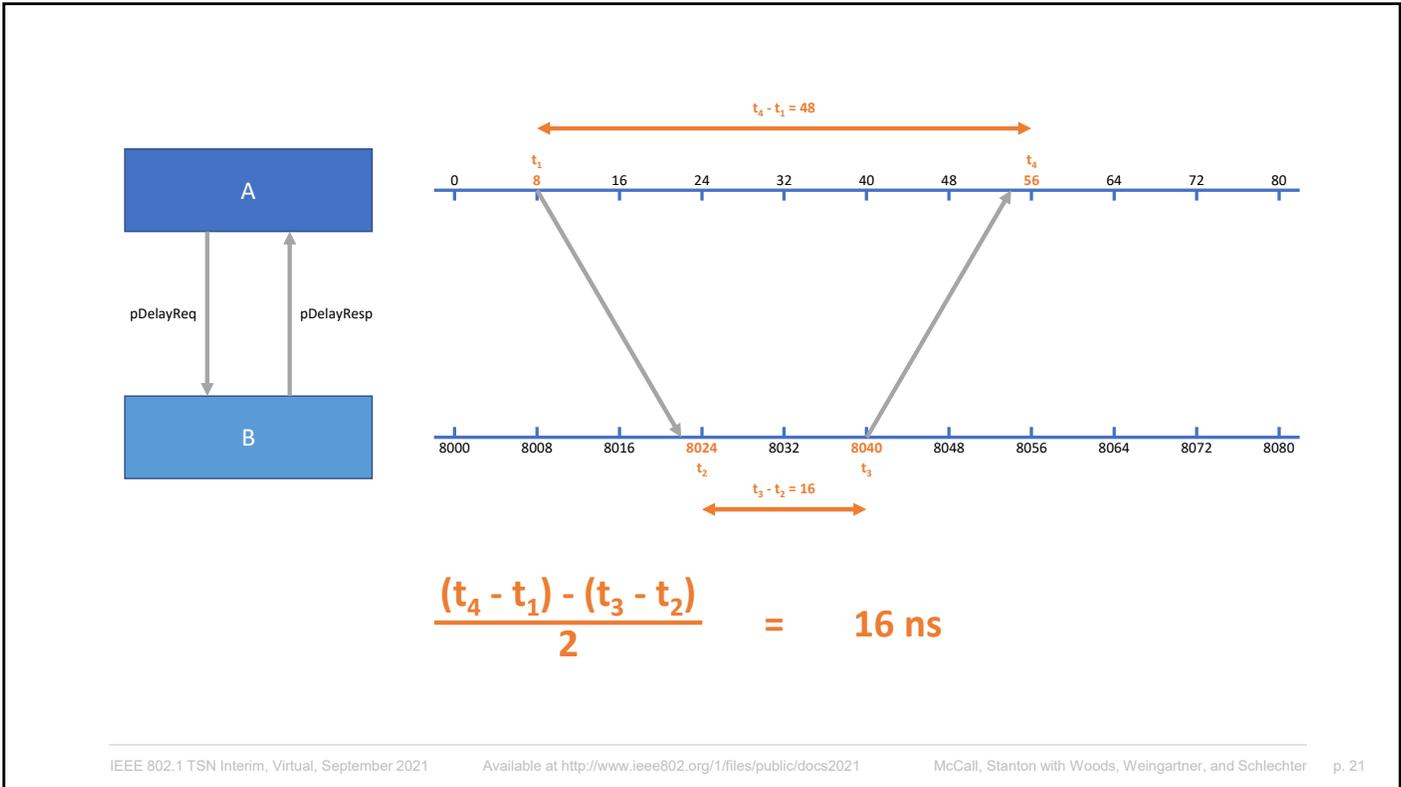
Let's put some numbers on the time line. We'll use a resolution of 8ns, which is what you'd get with a clock speed of 125MHz.

But...clocks A & B aren't the same...they aren't counting the same numbers at the same time.



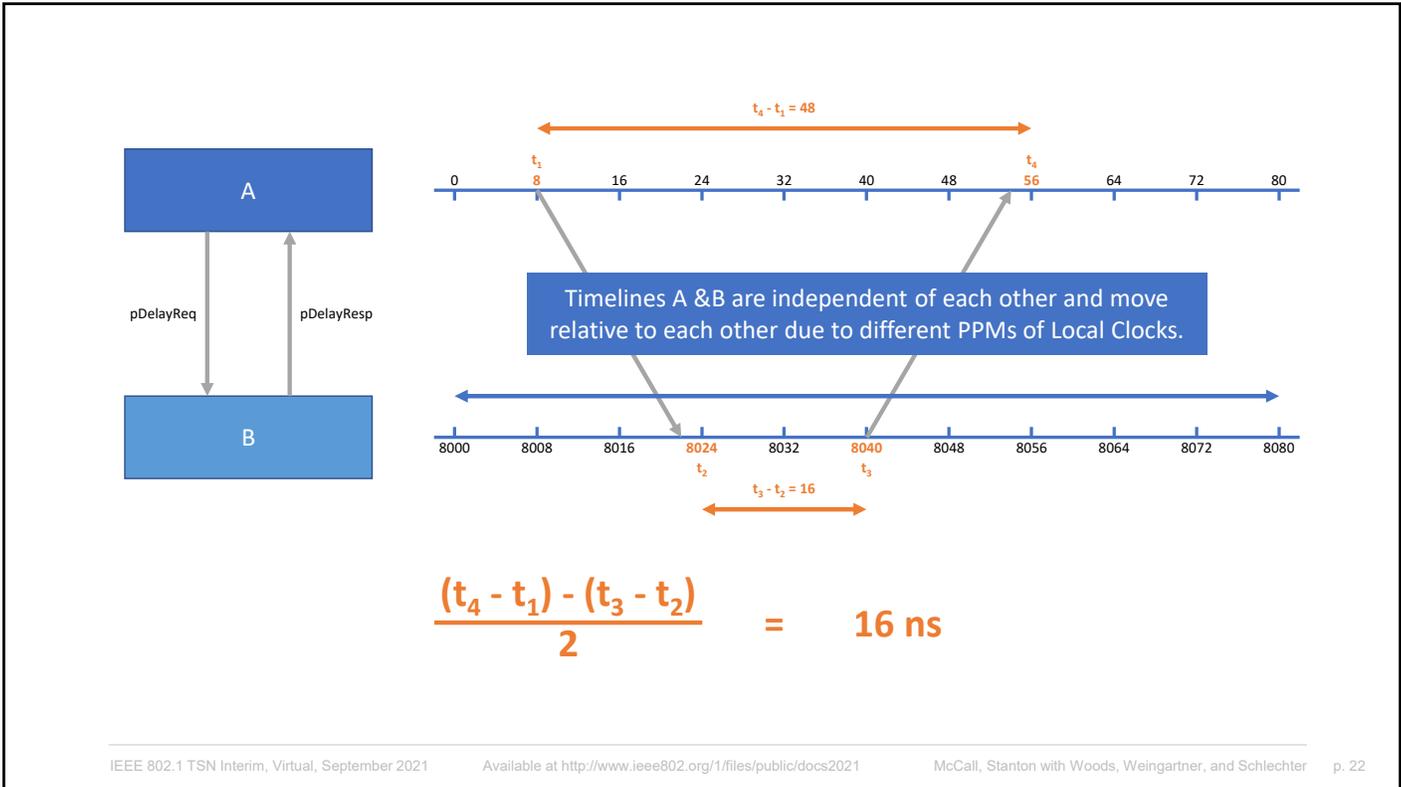
We'll model this by just adding a nice, round 8,000 to the lower line.

This requires changing the equation to subtract B's pDelay Response time (on the lower timeline) from the time between A sending the Request and receiving the Response. It's mathematically equivalent...and also allows compensation for differing clock speeds (via Neighbor Rate Ratio...which we are **not** considering in this section).

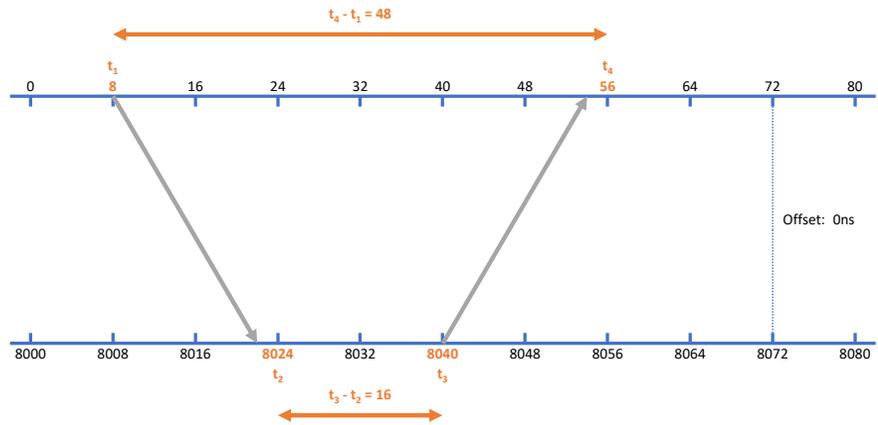
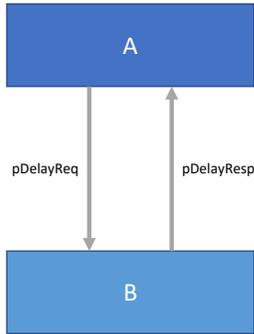


8ns is our Timestamp Granularity...which means the pDelay Req could arrive at any time between clock ticks and the Timestamp will actually reflect the next clock tick.

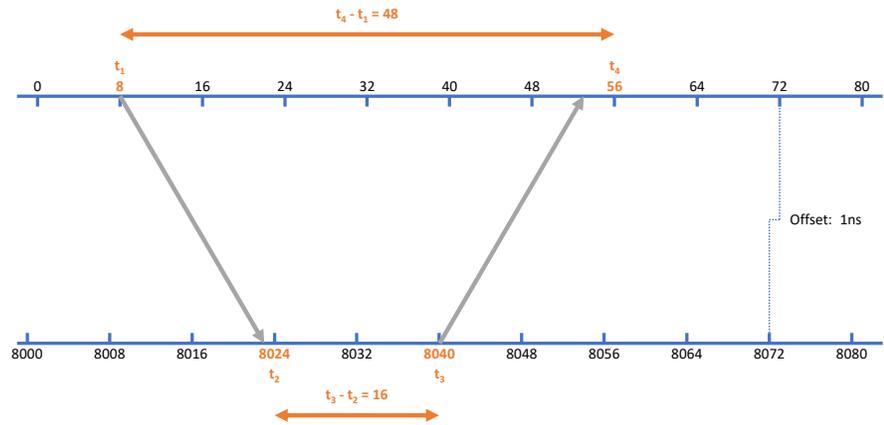
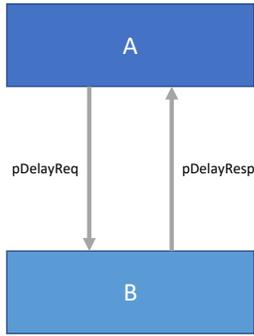
I'm keeping things simple here by assuming the TX time is exactly on a clock tick (and measured as such)...but you can see the effect this has. The measurement of pDelay is 16ns, not 14ns.



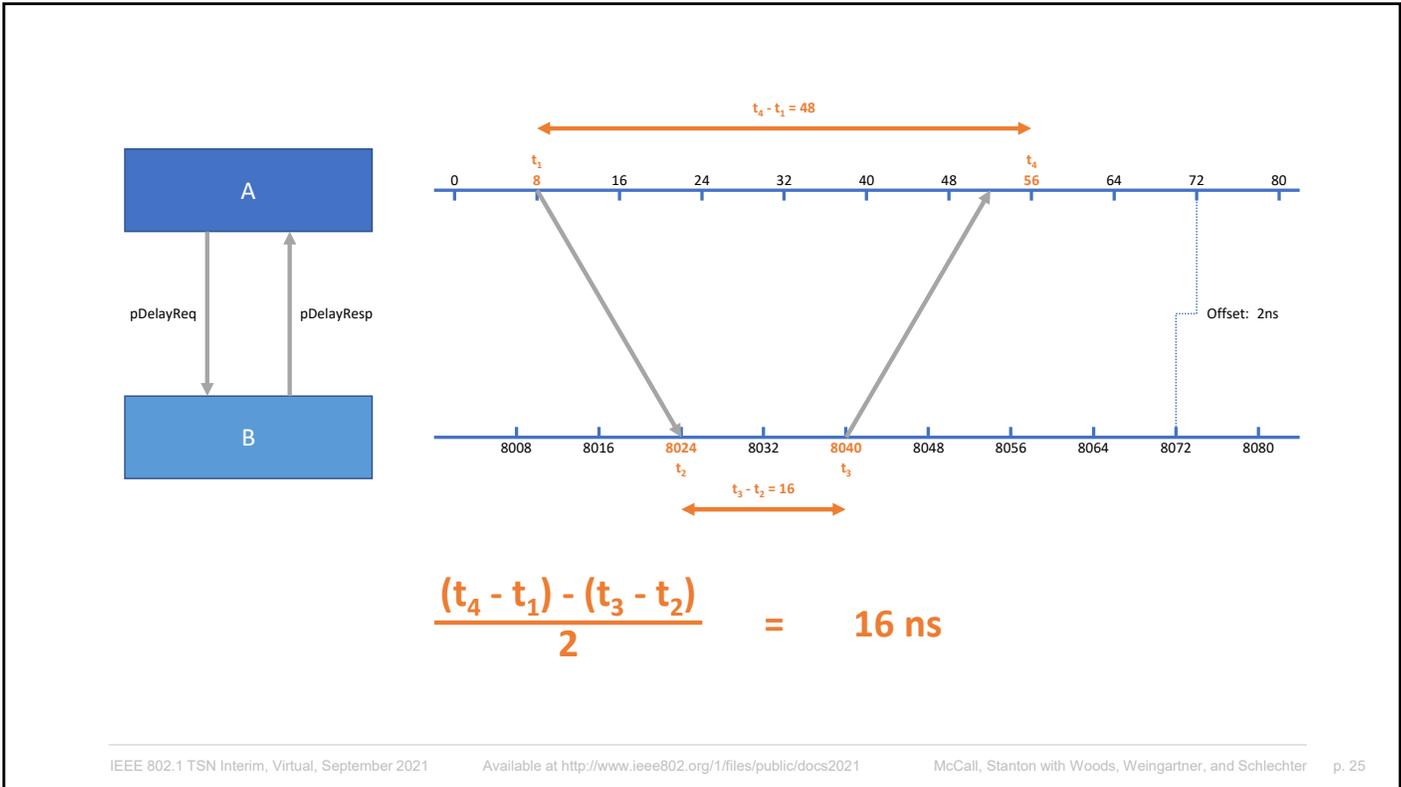
I'll deal with the movement relative to each other in a moment. For now, let's just look at this measurement and the fact that the timelines are independent. This means the clock ticks don't necessarily line up exactly. They might...but it's more likely there will be an offset.



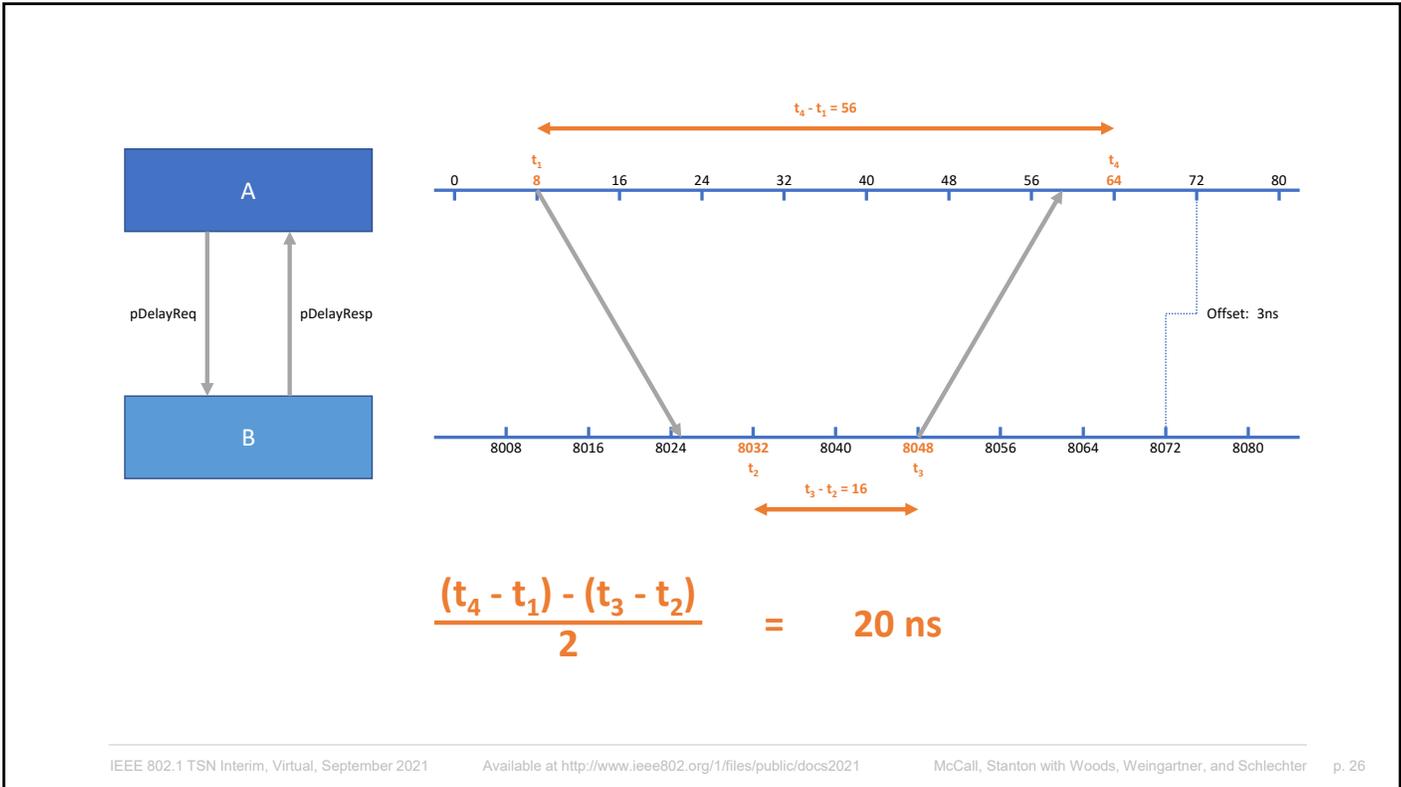
$$\frac{(t_4 - t_1) - (t_3 - t_2)}{2} = 16 \text{ ns}$$



$$\frac{(t_4 - t_1) - (t_3 - t_2)}{2} = 16 \text{ ns}$$

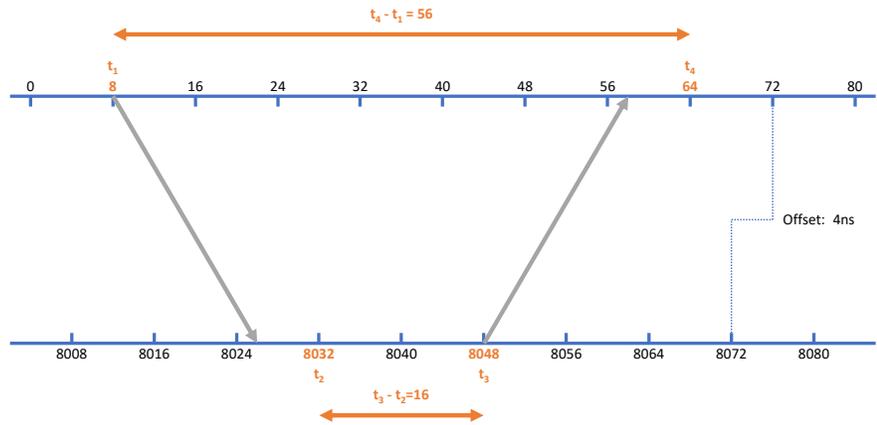
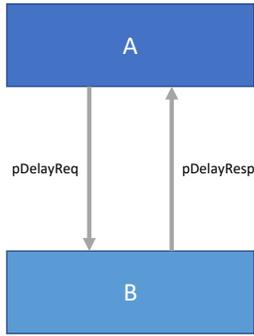


At an offset >2ns, the measurement is no longer 16ns...

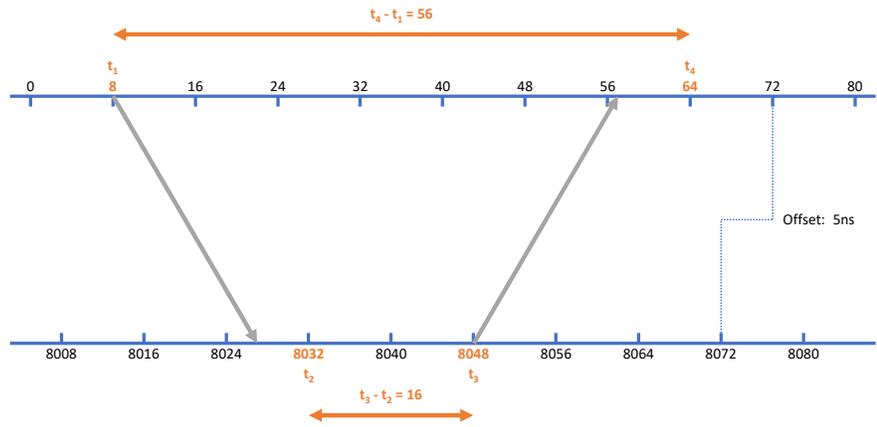
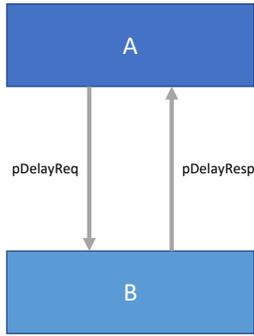


...it's worse. 20ns.

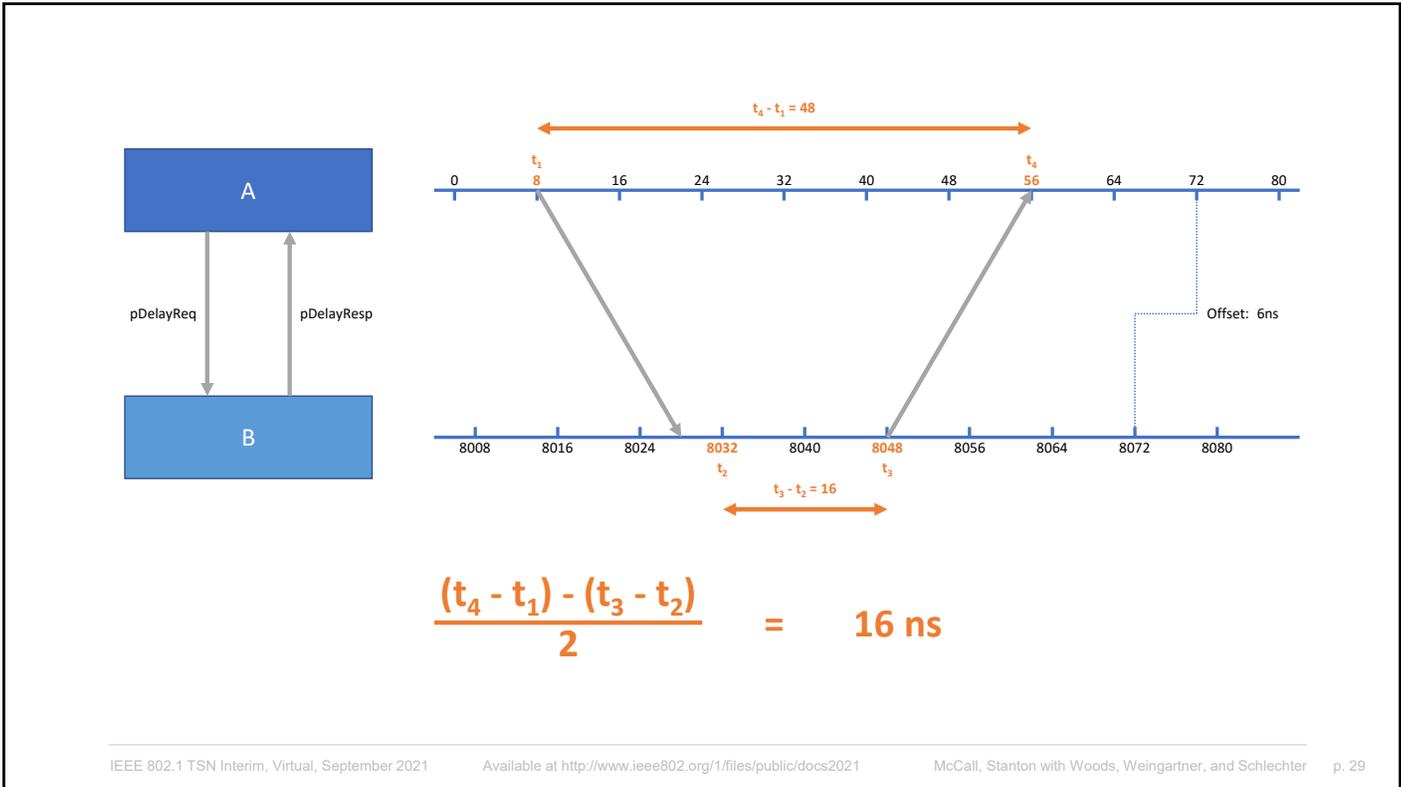
This, by the way, is to be expected. With a Timestamp Granularity of 8ns, and a divide by 2 in the equation, we can expect a measurement granularity of 4ns.



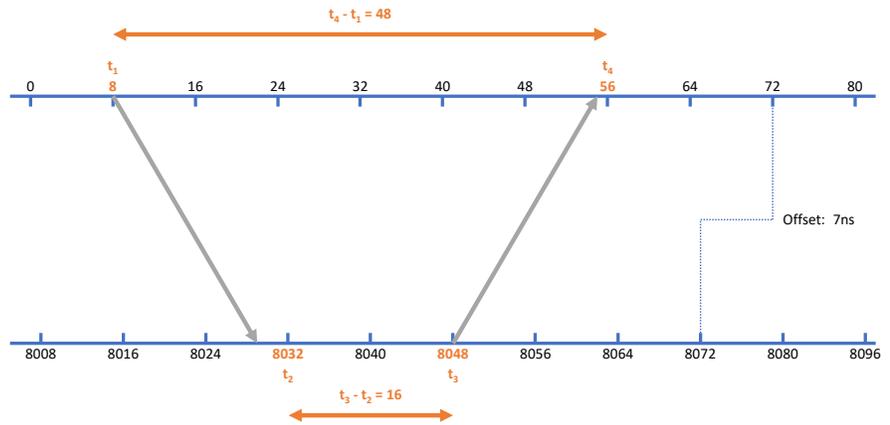
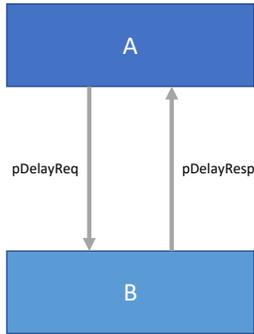
$$\frac{(t_4 - t_1) - (t_3 - t_2)}{2} = 20 \text{ ns}$$



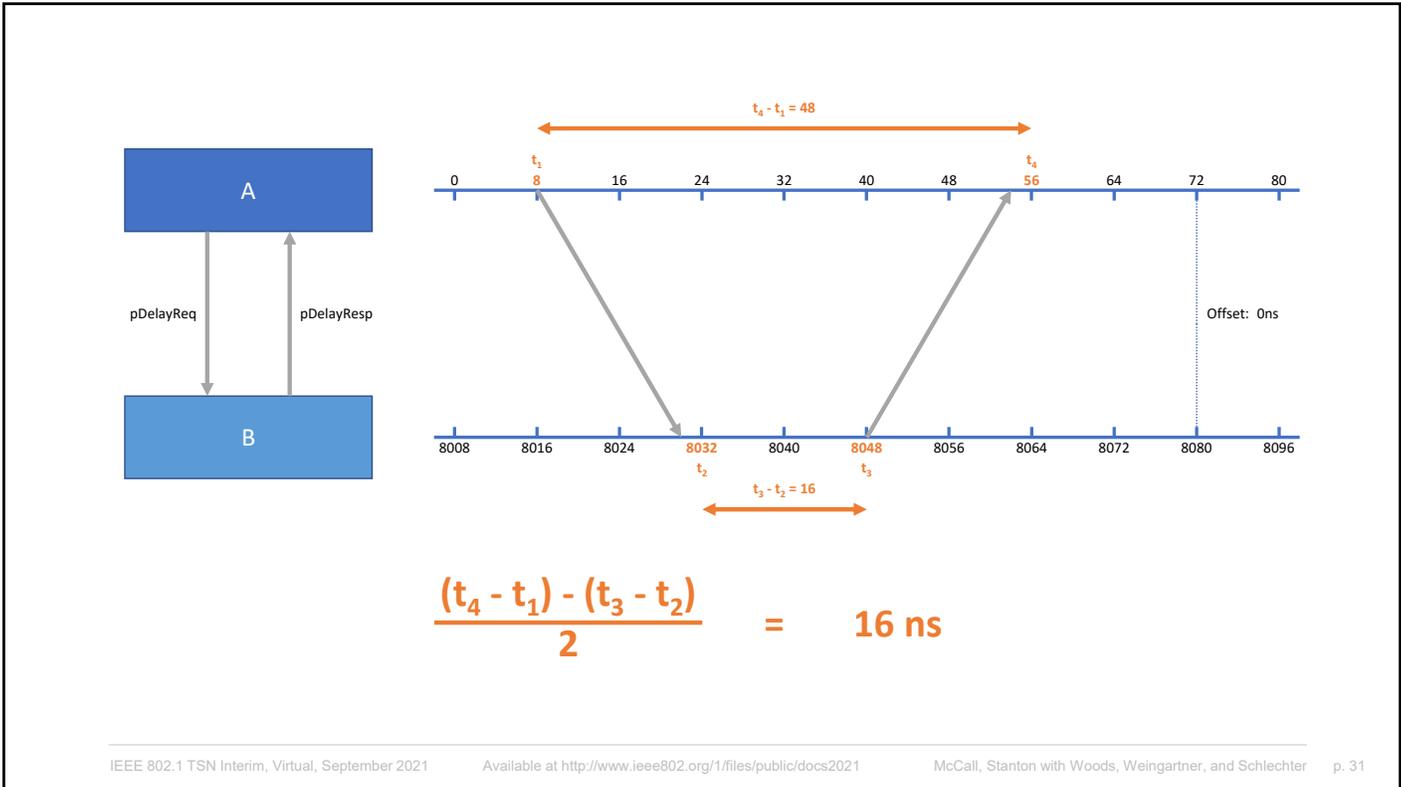
$$\frac{(t_4 - t_1) - (t_3 - t_2)}{2} = 20 \text{ ns}$$



At 6ns the measurement gets back to 16ns.



$$\frac{(t_4 - t_1) - (t_3 - t_2)}{2} = 16 \text{ ns}$$



And then we're back to where we started, with an offset of 0ns.

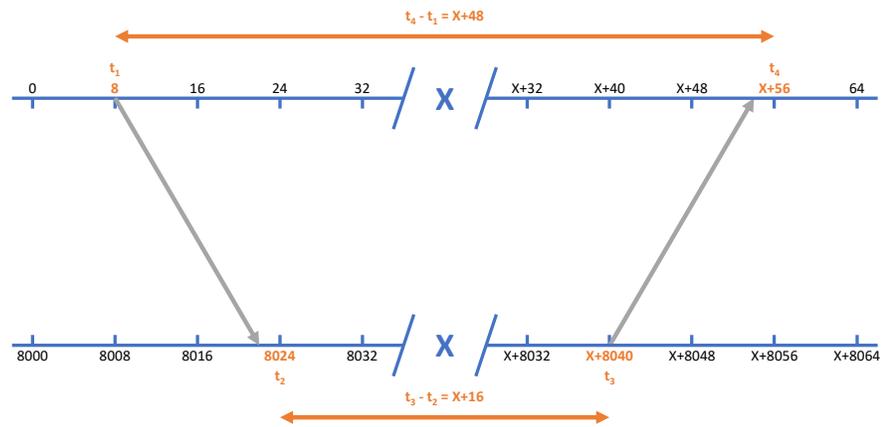
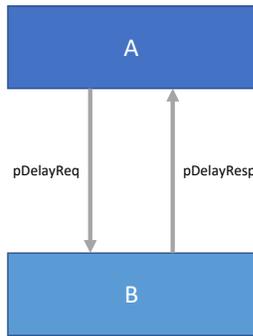
Offset / ns	Measured pDelay / ns
0 – 2	16
>2 – 6	20
>6 – <8	16

Average = 18 ns \neq 14 ns

OK to average as, over multiple measurements, the timelines will move relative to each other.
But on it's own that doesn't help.

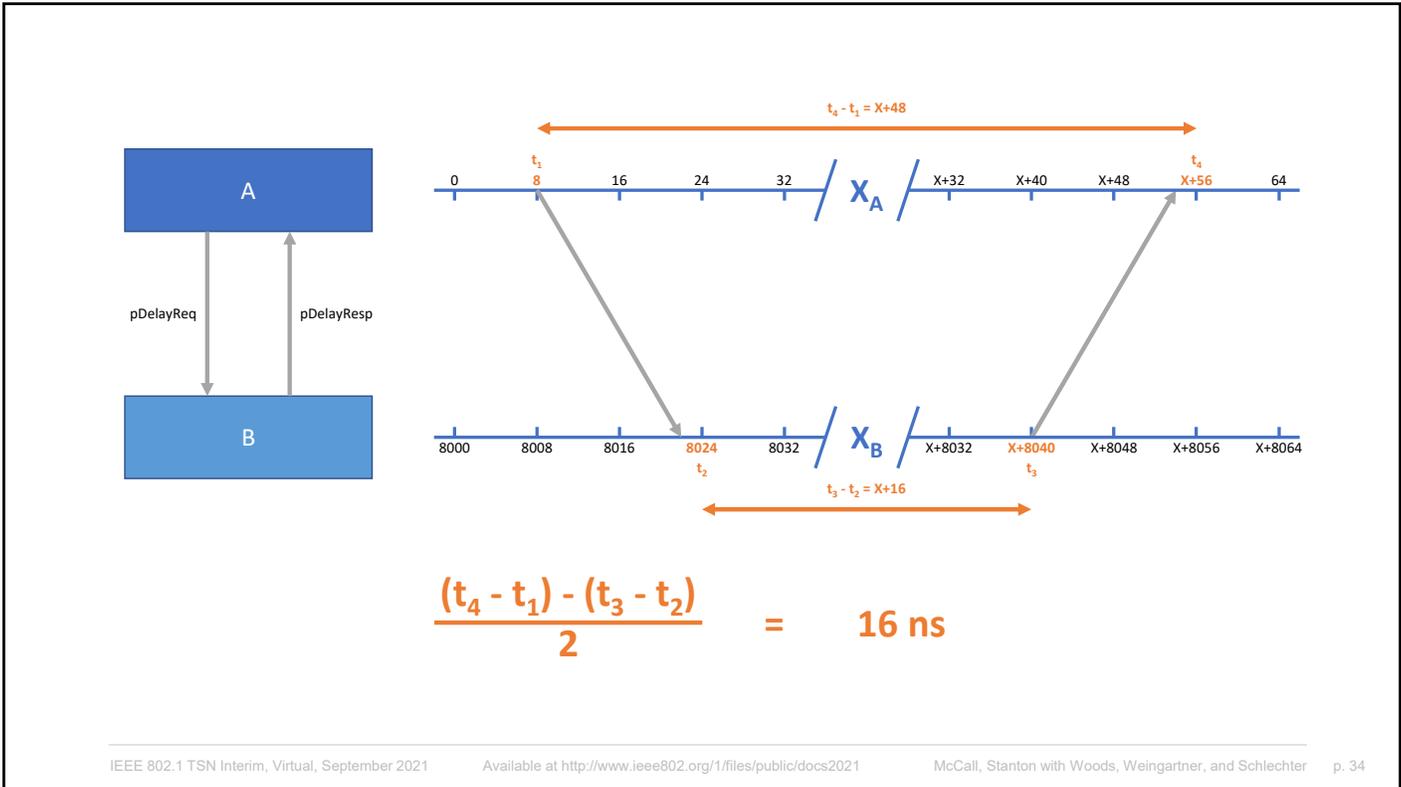
So...does this even work at all?

The good news is that yes, it does. There are two reasons...

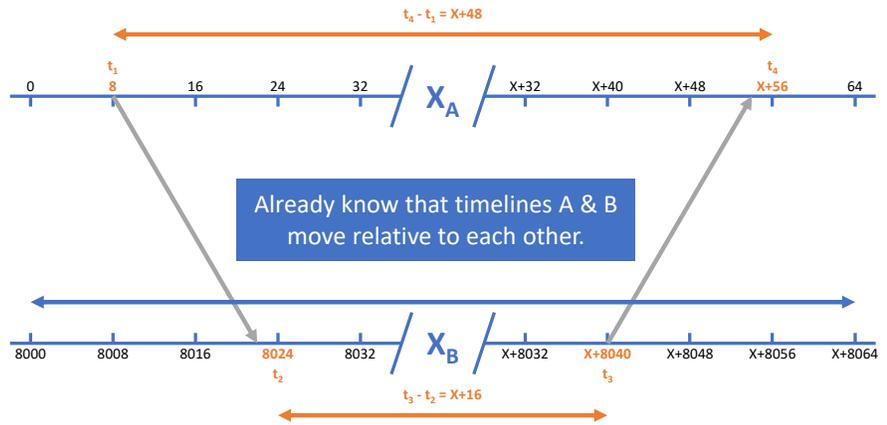
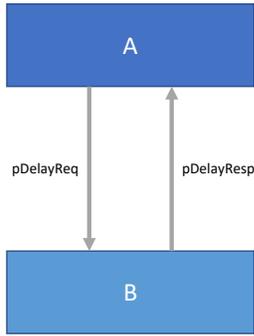


$$\frac{(t_4 - t_1) - (t_3 - t_2)}{2} = 16 \text{ ns}$$

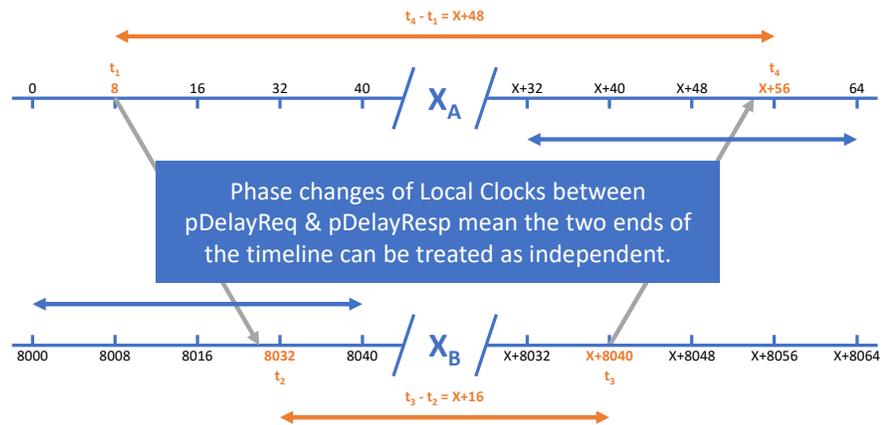
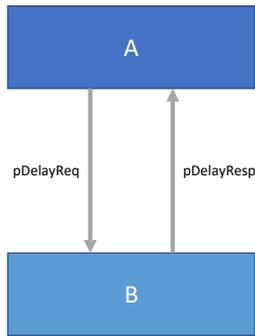
Firstly, there is a relatively large amount of time between the left and right ends of the timeline due to the fact that pDelay Response Time is much longer than pDelay...a fact I previously ignored.



And secondly, that relatively long period of time is measured by two free running local clocks. Running at different speeds...and drifting slightly relative to each other.

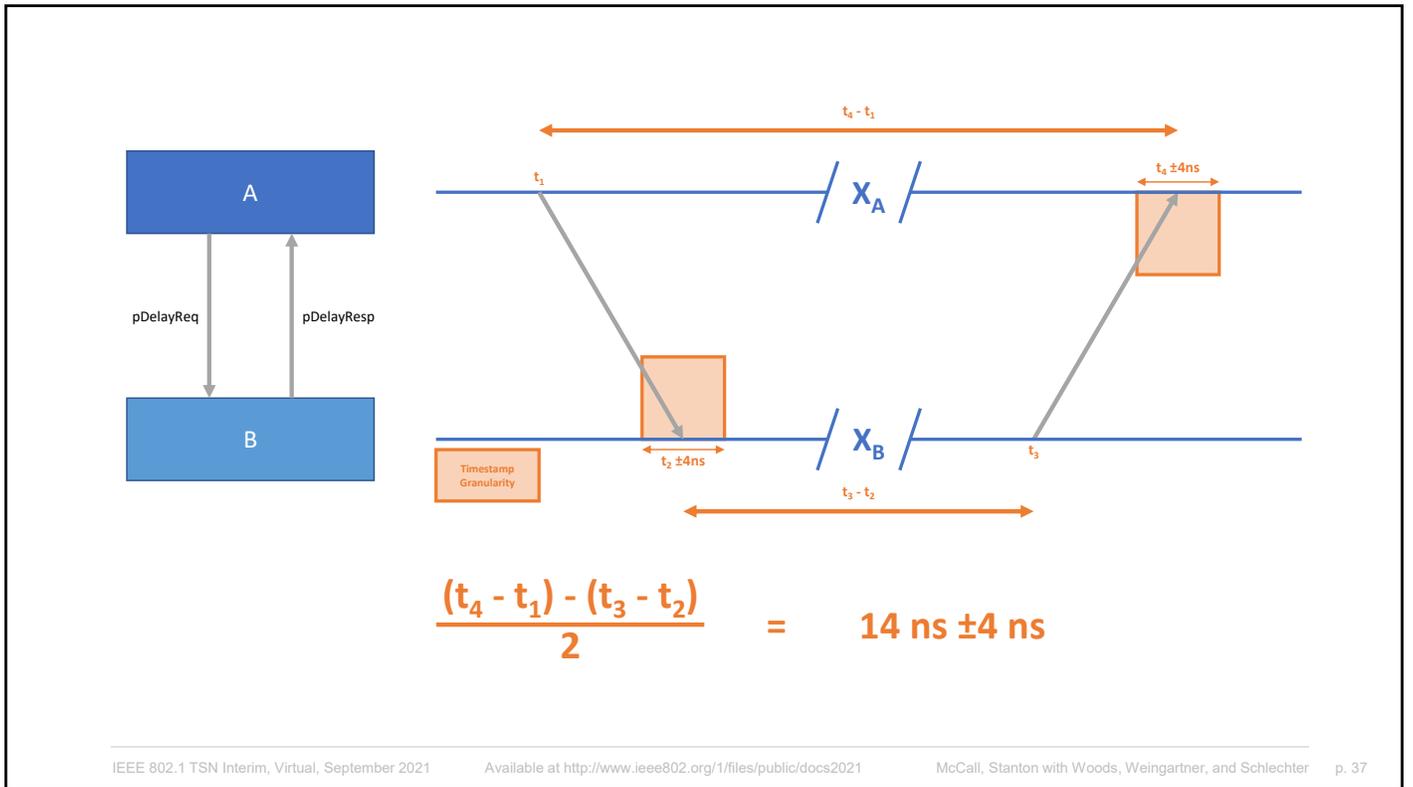


$$\frac{(t_4 - t_1) - (t_3 - t_2)}{2} = 16 \text{ ns}$$

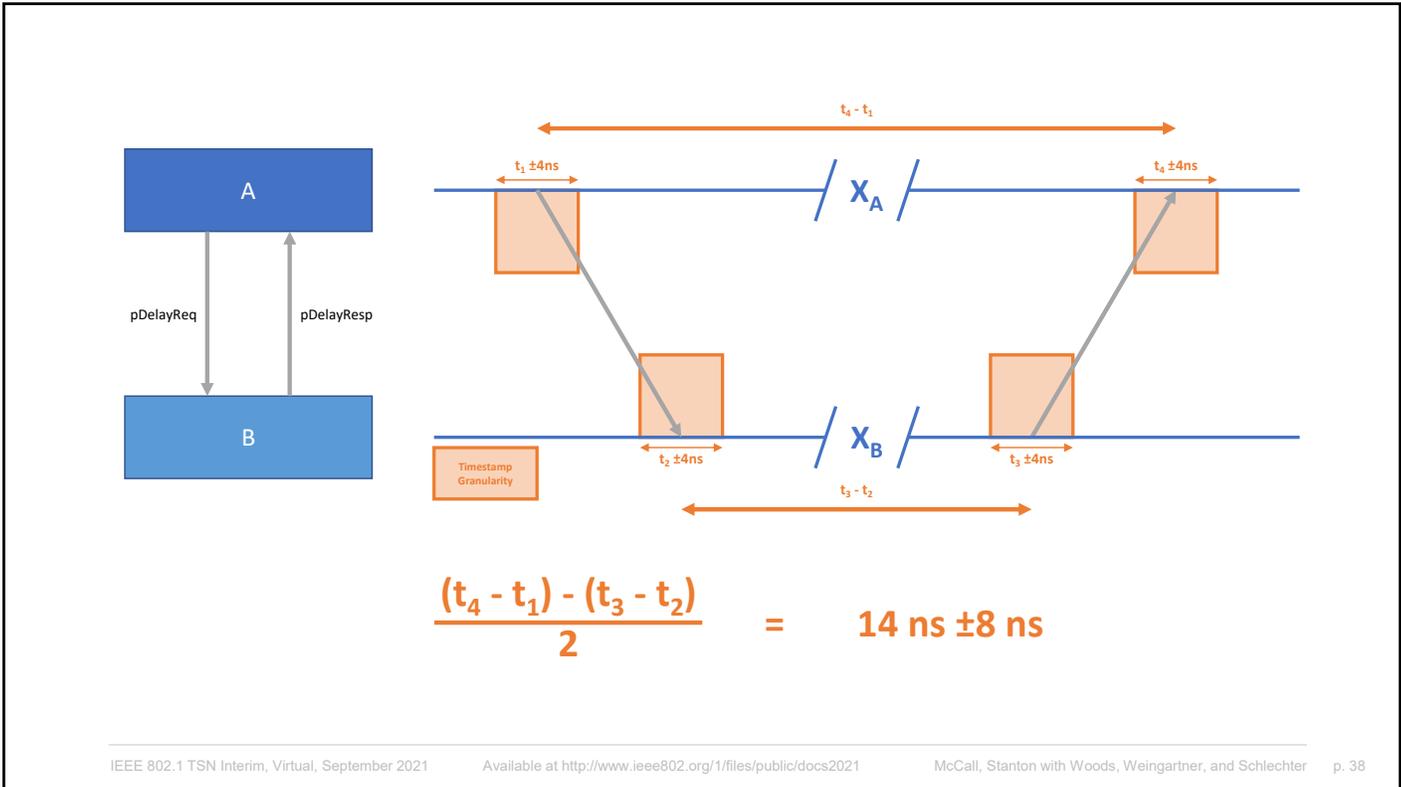


$$\frac{(t_4 - t_1) - (t_3 - t_2)}{2} = 16 \text{ ns}$$

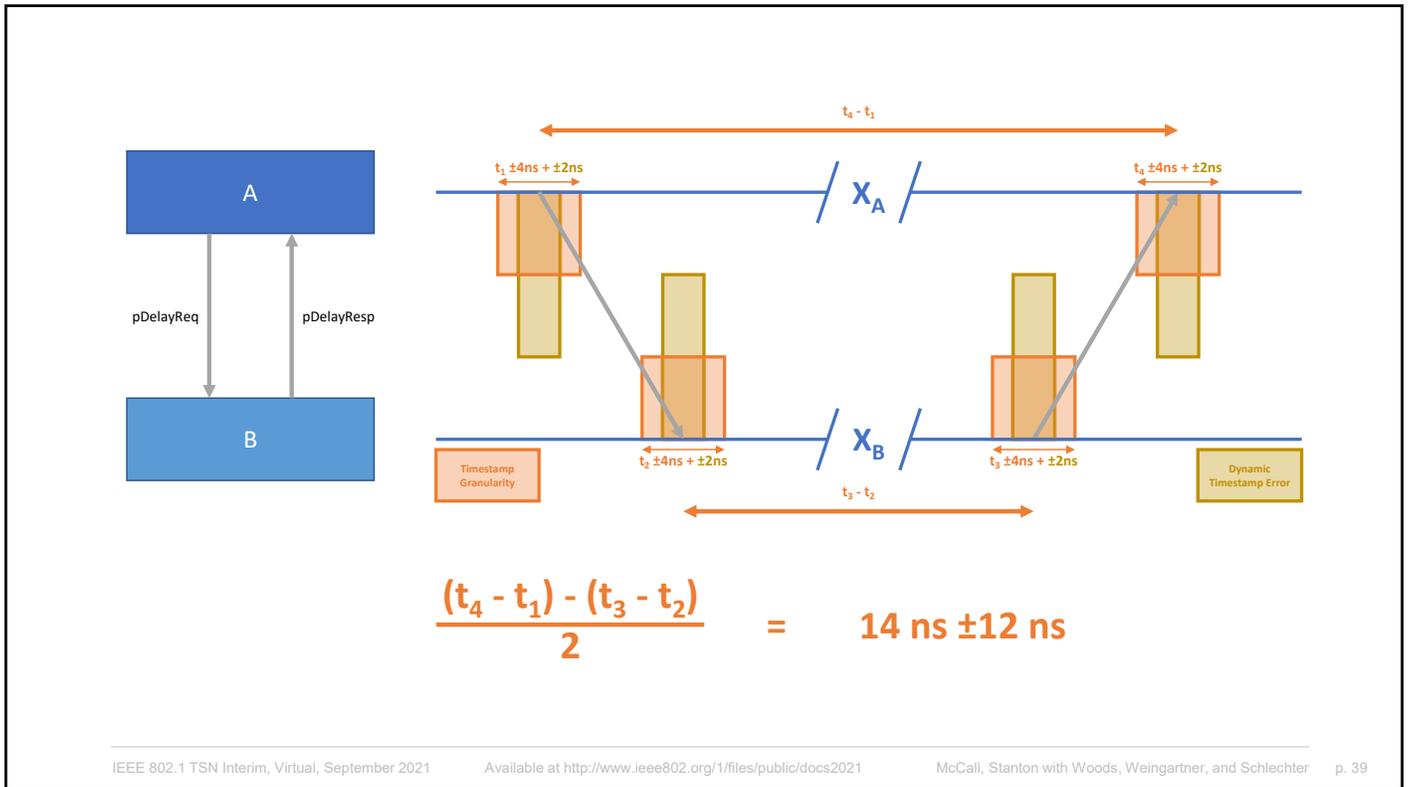
Because the left and right ends of the timeline can effectively be decoupled from each other...because there are different numbers of clock ticks in X_A and X_B (which is eventually compensated for by NRR...which we aren't not currently considering right now)...you can look at the measurement probabilistically...which looks like this...



This shows only error due to Timestamp Granularity...8ns...centred on the 14ns, i.e. ± 4 ns



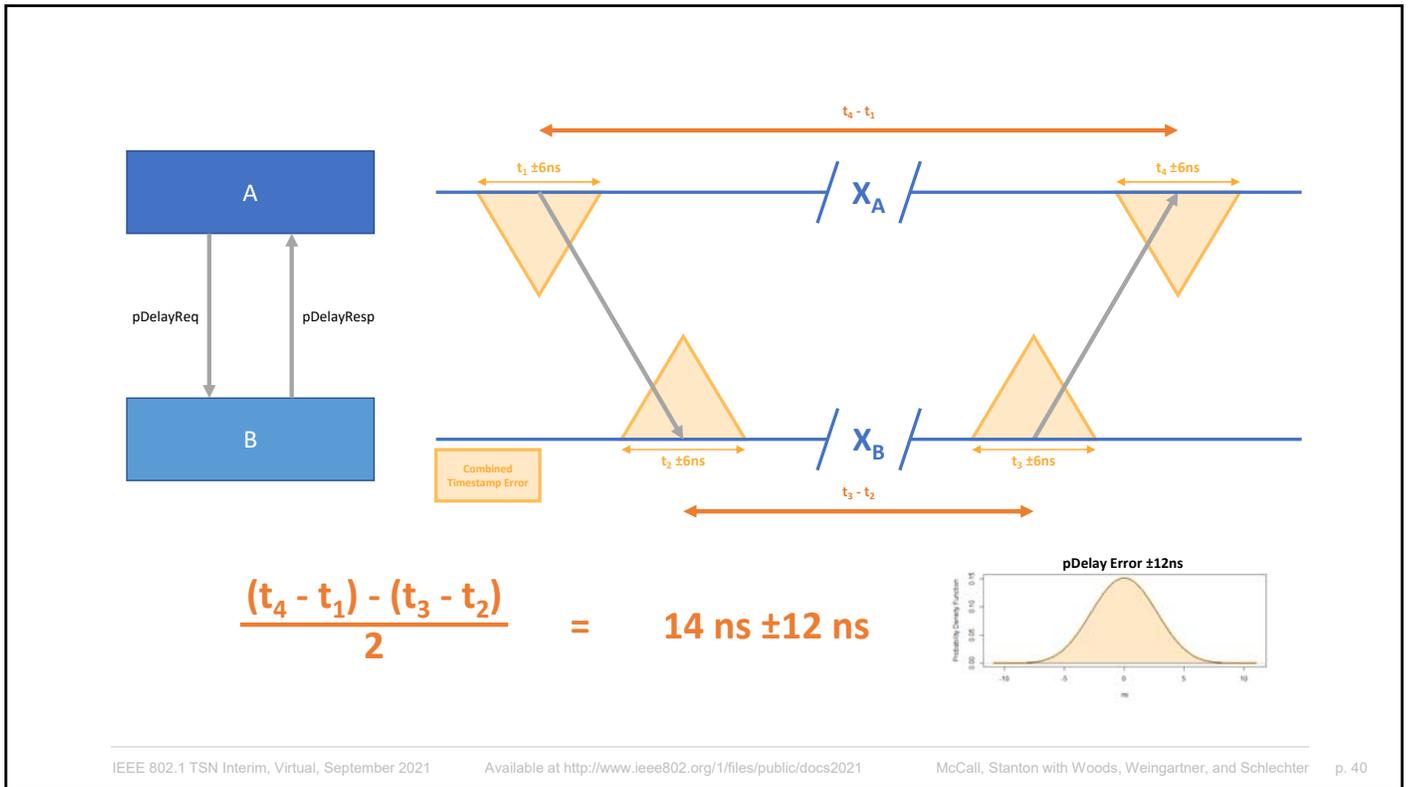
Add TX Timestamp Granularity.



Add Dynamic Time Stamp Error, ±2 ns.

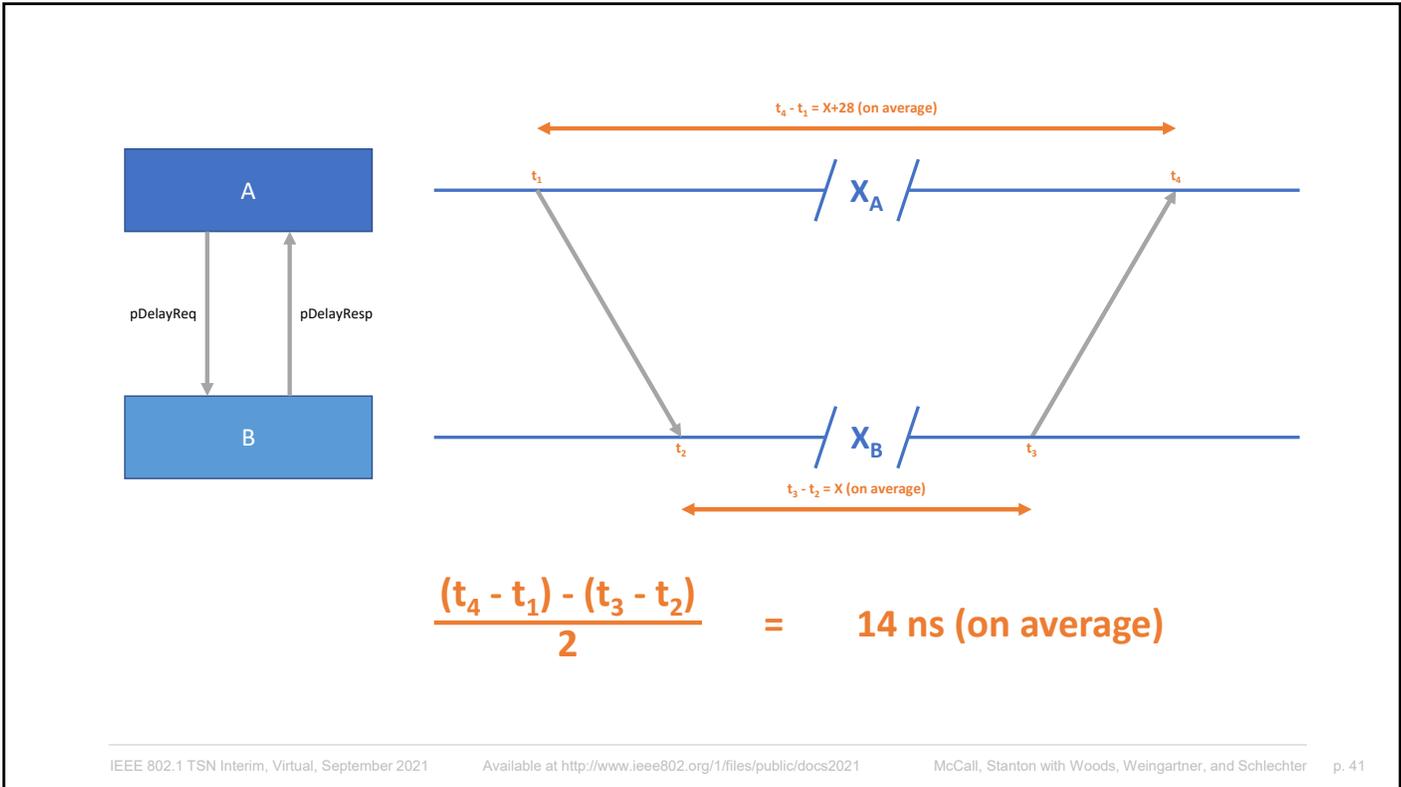
Indications are that this error should be less than Timestamp granularity. It varies with implementation and I'll let Kevin talk about that a bit more later...but for the moment let's just model it as this probability.

If you add together all 8 uniform probabilities...



You get this.

At each TX and RX point, two uniform probabilities combine to give a triangular probability...and those combine to an overall bell curve probability.

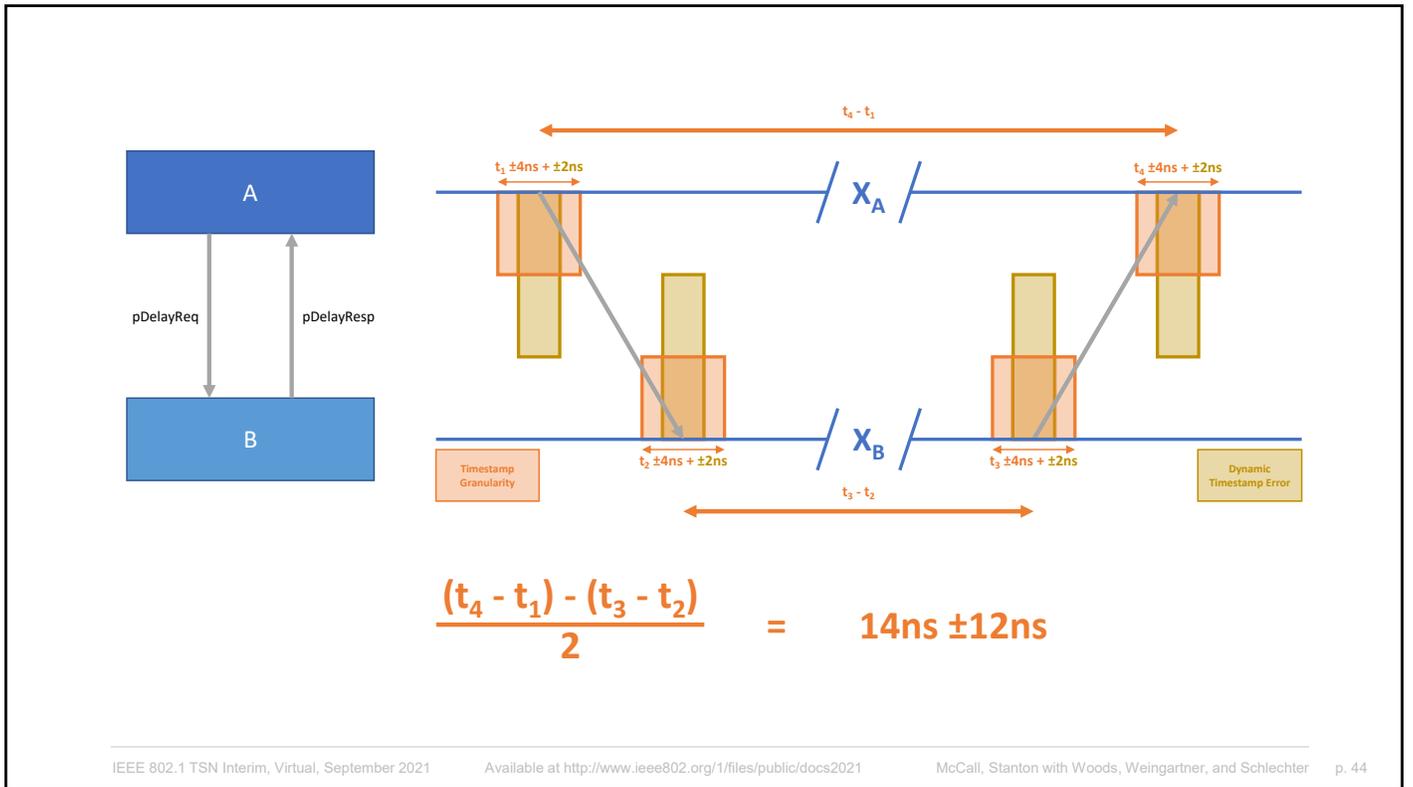


Which can be averaged out.

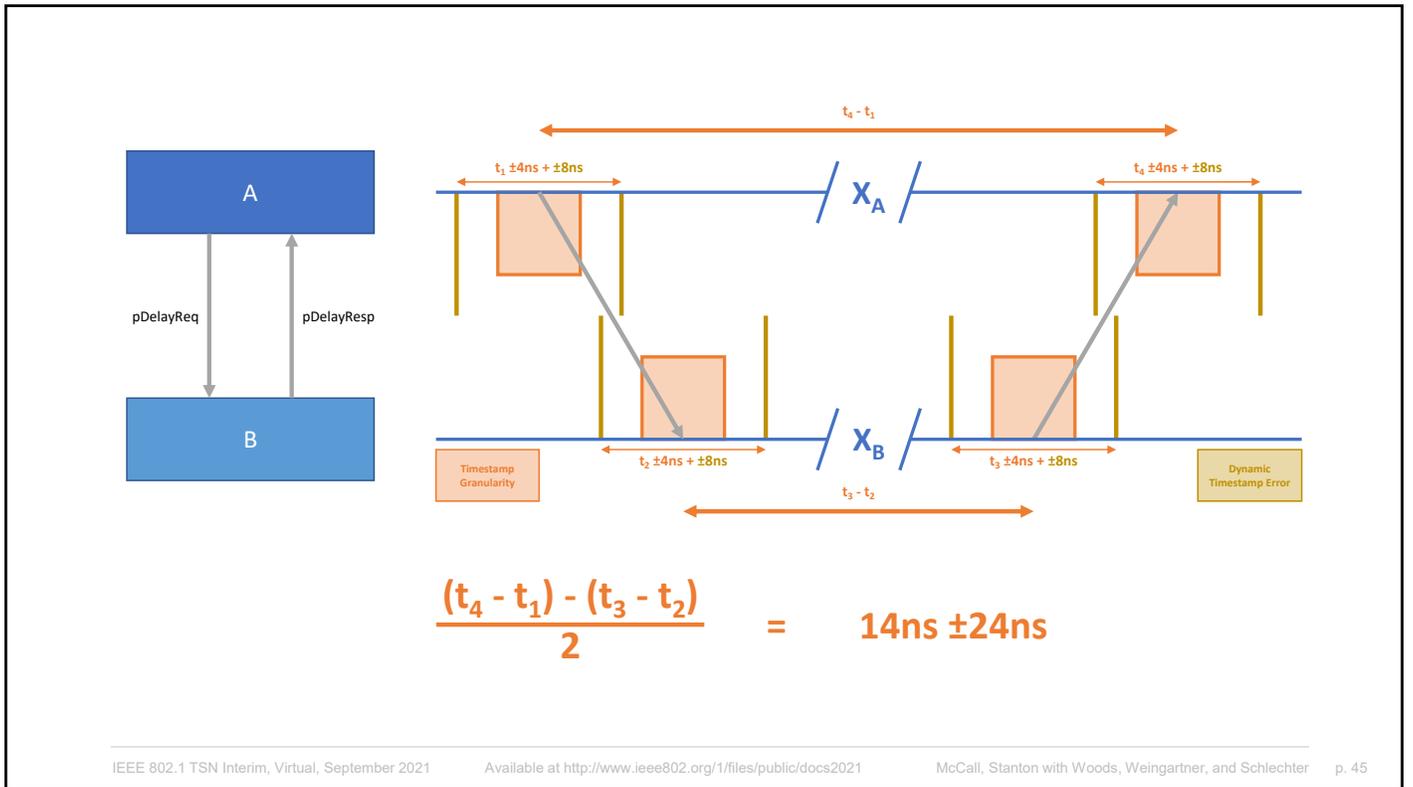
Link Delay Error – Observations

- Averaging of pDelay measurements should be mandatory
 - Delivers meanLinkDelay with much lower error
- Helpful if implementations exhibit error probability distributions that can be averaged out to zero
- For Ethernet...
 - The goal is to estimate a relatively fixed quantity
 - Very low bandwidth low-pass filter is appropriate due to stability of pDelay. (IIR?)
 - Some dedicated startup behaviour may be appropriate to arrive quickly at a useful average
- For wireless...
 - More complex due to multi-path and potential movement of devices, but...
 - Unlikely to have long chains of devices, or at least long chains of mobile devices.
 - Out of scope for the current discussion.

Link Delay & Timestamp Errors – Current Simulation



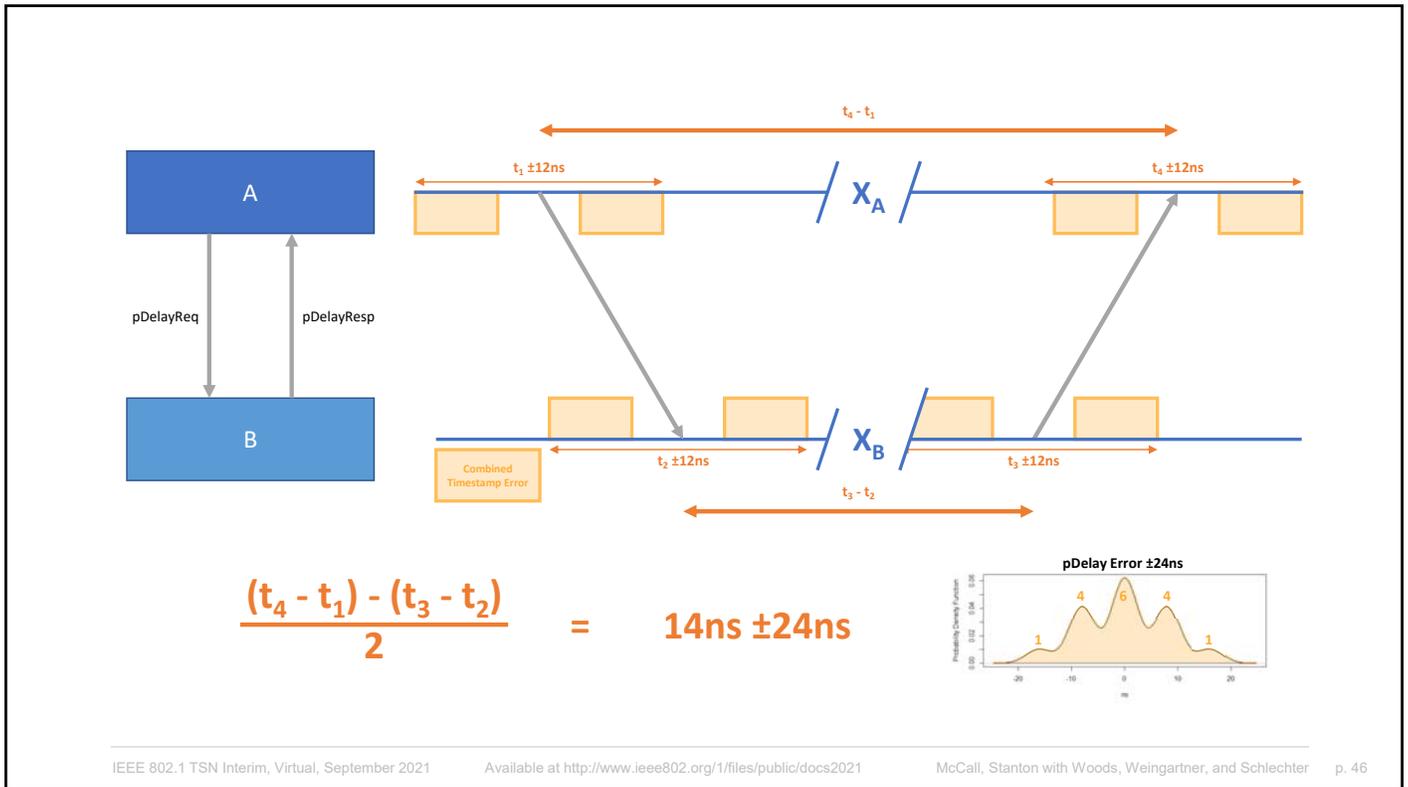
This is the model we just saw. The simulation does not do this. Instead it does...



This.

Dynamic Timestamp Error is $\pm 8\text{ns}$...but it's either $+8$ or -8 with a 50:50 change of either for each and every timestamp. That's represented by the two vertical lines...one at $+8$, one at -8 .

Working out the probabilities...

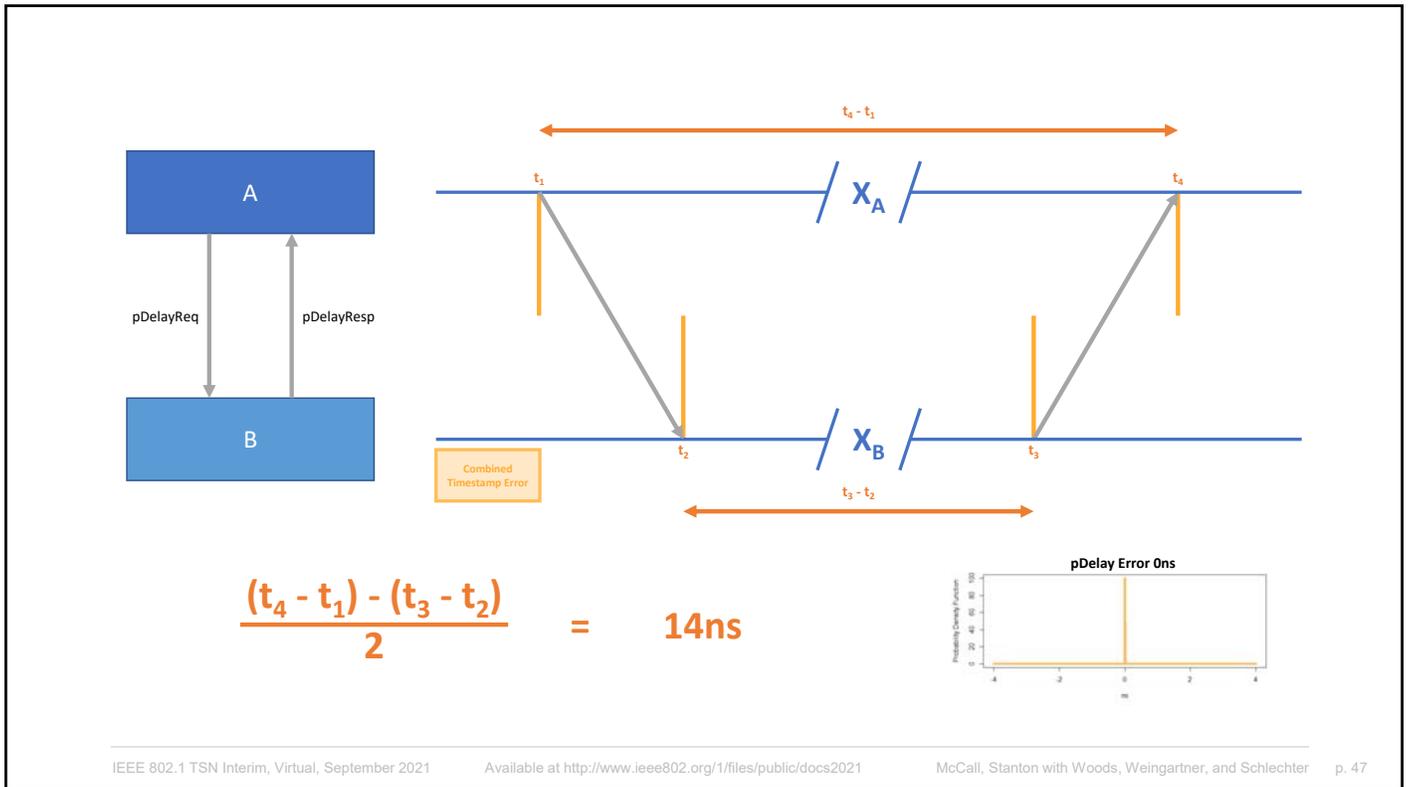


This is Case 1 (and most other cases) in 60802-garner-new-simulation-results-new-freq-stab-model-0421-v02.pdf, i.e. the most recent set of time series simulations.

And there is no averaging of pDelay in the current simulation. The latest pDelay measurement is passed directly into the Correction Field.

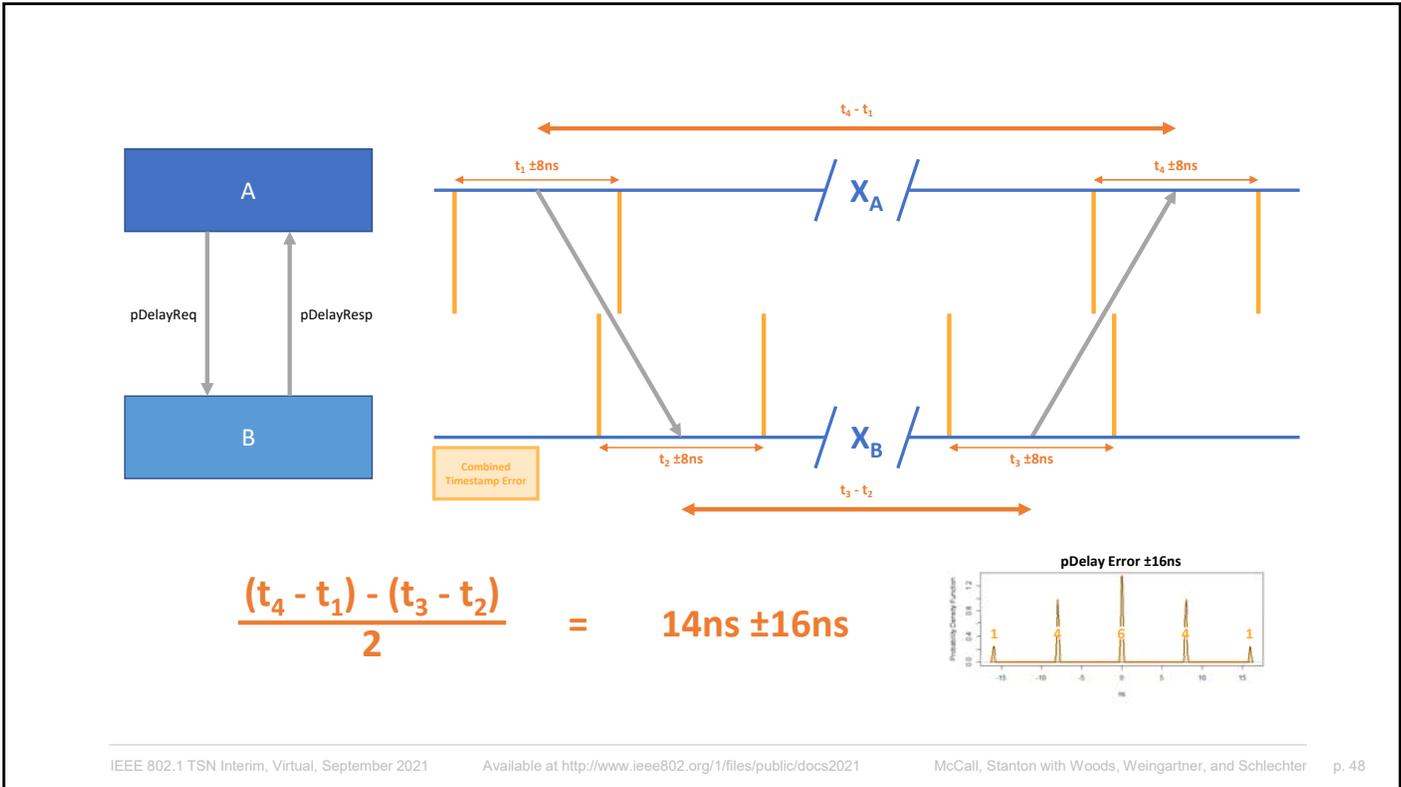
Oh...and if you are wondering if this could result in a negative pDelay, i.e. the simulation of time travel...yes, it could. But not with the current simulation's pDelay value of 150ns.

There were also three other simulations of interest in that same run of simulations.



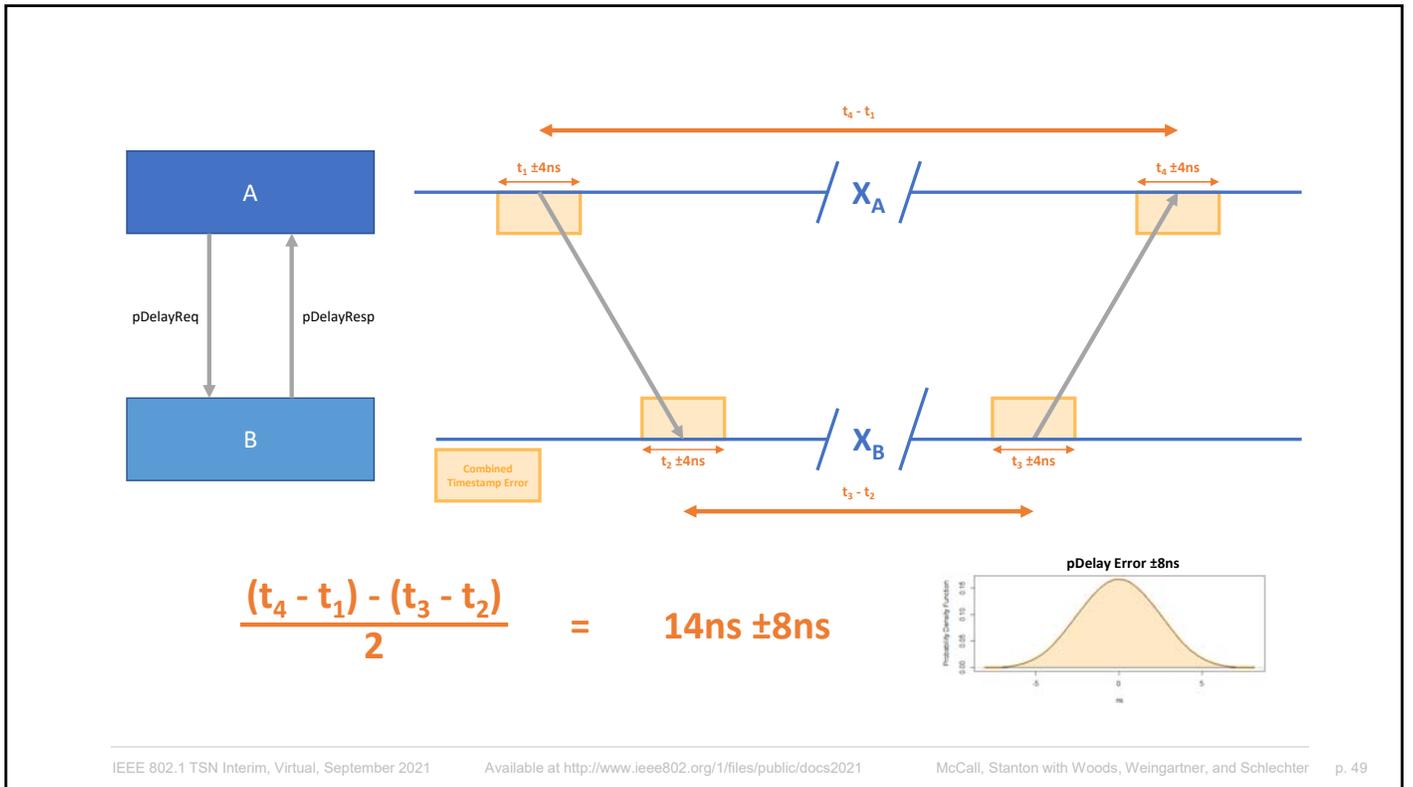
Case 13 in 60802-garner-new-simulation-results-new-freq-stab-model-0421-v02.pdf

This turned both errors to zero. Which, for Link Delay is the equivalent of what averaging pDelay measurements should accomplish.



Case 14 in 60802-garner-new-simulation-results-new-freq-stab-model-0421-v02.pdf

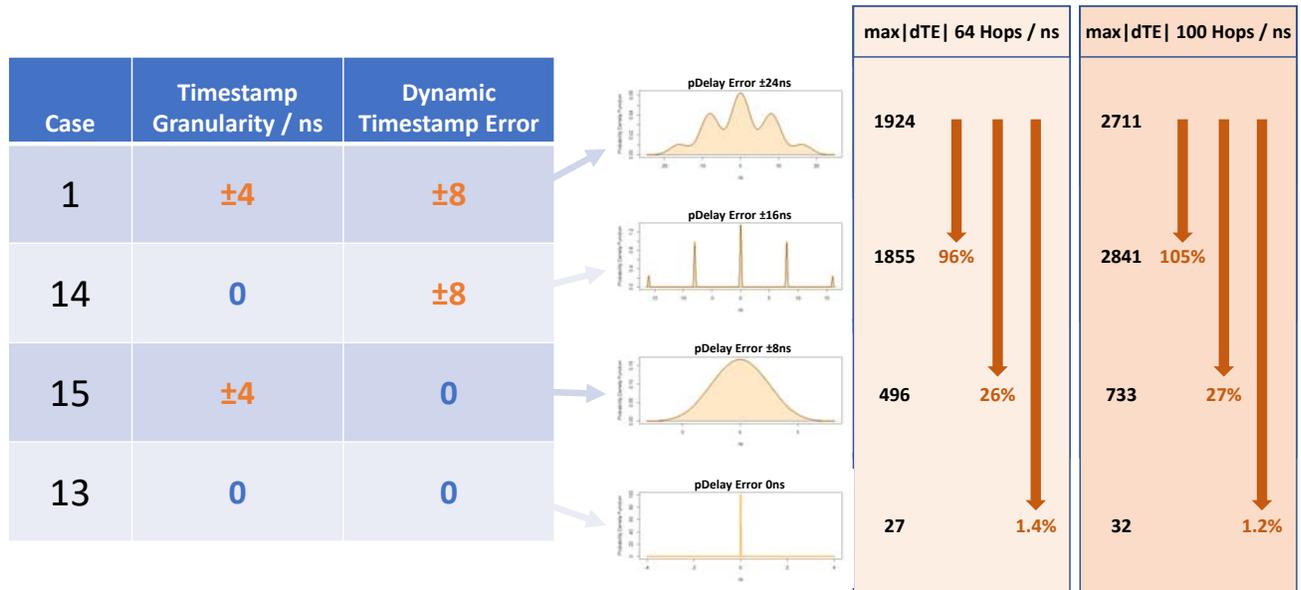
This removed Timestamp Granularity Error...but kept Dynamic Timestamp Error.



Case 15 in 60802-garner-new-simulation-results-new-freq-stab-model-0421-v02.pdf

And this did the opposite...kept Timestamp Granularity Error but removed Dynamic Timestamp Error.

Comparison of Analysis & Latest Simulations



Note that the simulation removed the effects of Timestamp Granularity and Dynamic Timestamp Error from both Link Delay **and** Residence Time.

While Link Delay could get close to this via averaging...you can't do that with Residence Time. So you'll get some big benefits from averaging...but not quite as much as shown here.

Modelling Link Delay Error – Timestamp Granularity & Dynamic Timestamp Error

- Model linear error with uniform distribution for both factors.
- Allow separate RX & TX values for both, i.e. four factors...

	Timestamp Granularity	Dynamic Time Stamp Error
RX	✓	✓
TX	✓	✓

- Add “Error Correction Factor” to account for averaging.
 - One value to cover all errors from all four factors.
 - 25% Effective → Removes 25% of the error.

Formulae – $p\text{Delay}_{\text{errorTS}}$

$$\text{meanLinkDelay}_{\text{error}} = (1 - \text{meanLinkDelay}_{\text{errorCorrection}})(p\text{Delay}_{\text{errorTS}} + p\text{Delay}_{\text{errorNRR}})$$

$$p\text{Delay}_{\text{errorTS}} = \left(\frac{(t_{4\text{TSerror}} - t_{1\text{TSerror}}) - (t_{3\text{TSerror}} - t_{2\text{TSerror}})}{2} \right)$$

$$t_{1\text{TSerror}} = \text{TSGE}_{\text{TX}} + \text{DTSE}_{\text{TX}}$$

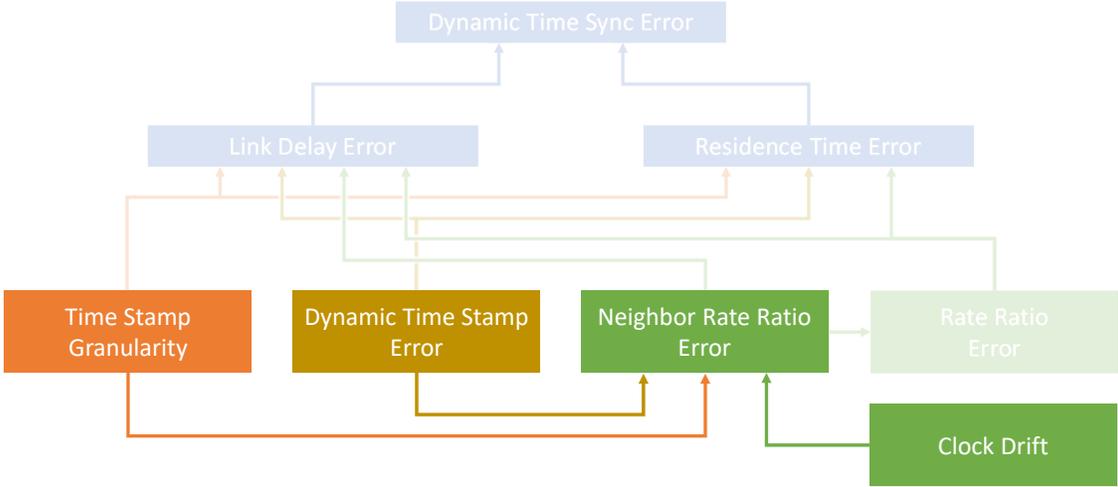
$$t_{2\text{TSerror}} = \text{TSGE}_{\text{RX}} + \text{DTSE}_{\text{RX}}$$

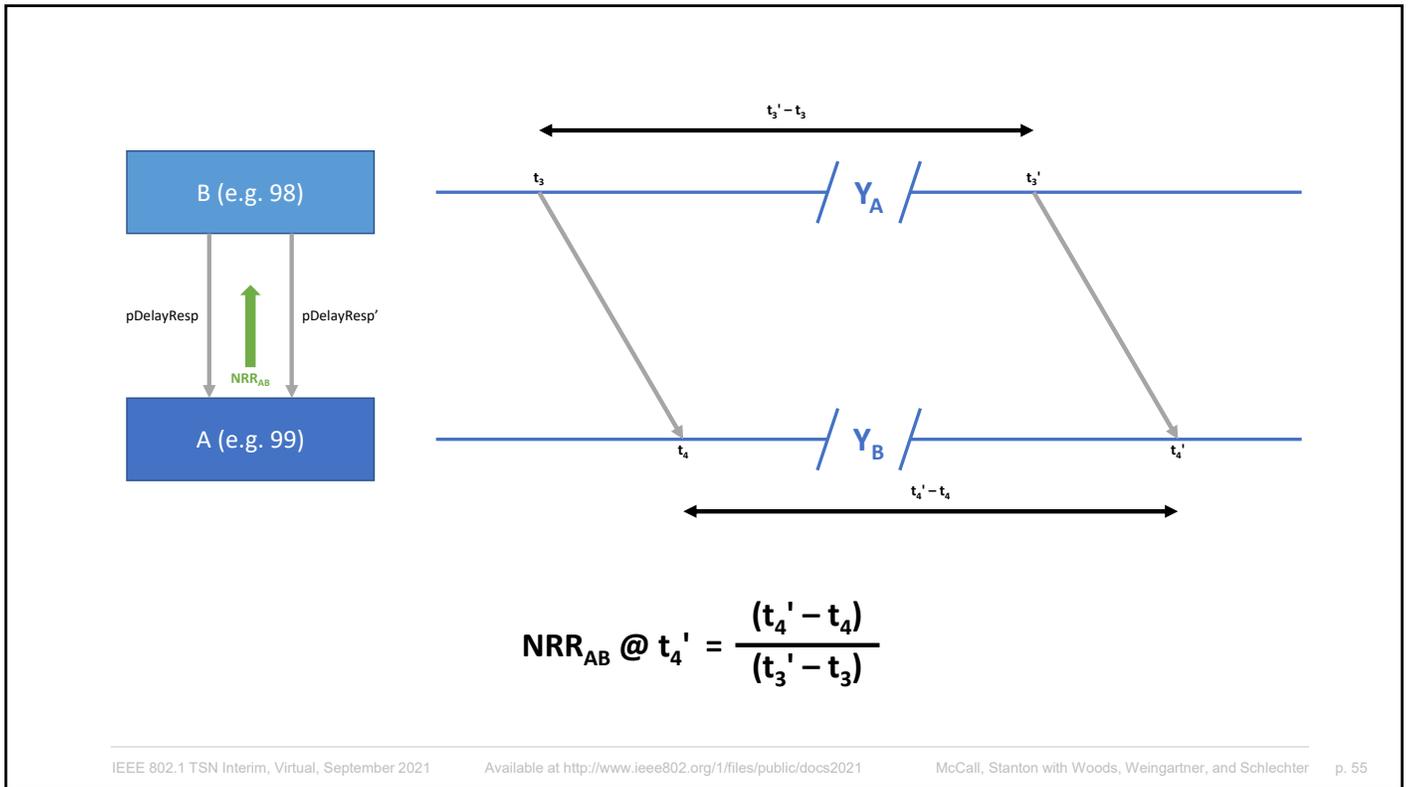
$$t_{3\text{TSerror}} = \text{TSGE}_{\text{TX}} + \text{DTSE}_{\text{TX}}$$

$$t_{4\text{TSerror}} = \text{TSGE}_{\text{RX}} + \text{DTSE}_{\text{RX}}$$

Neighbor Rate Ratio Error

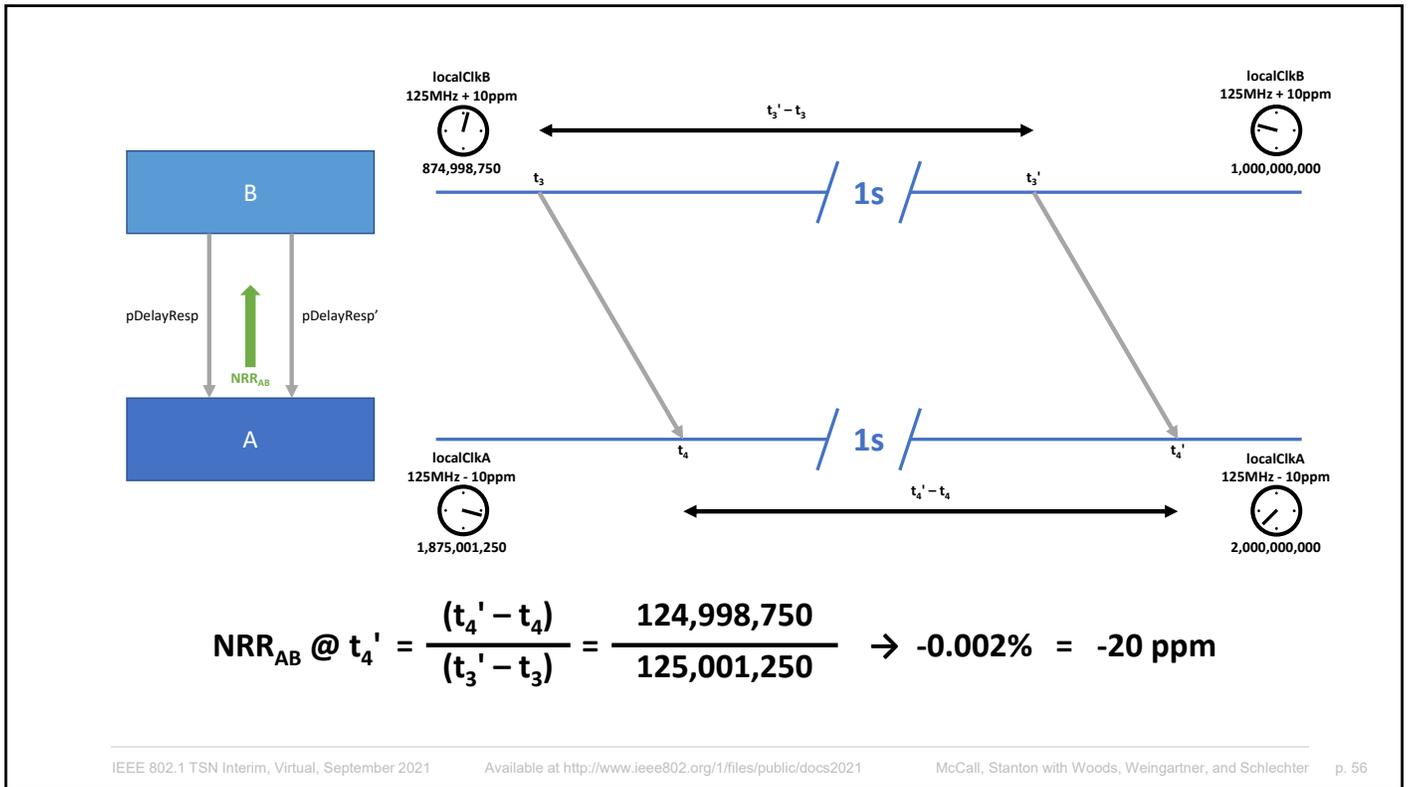
Time Sync – Elements & Relationships



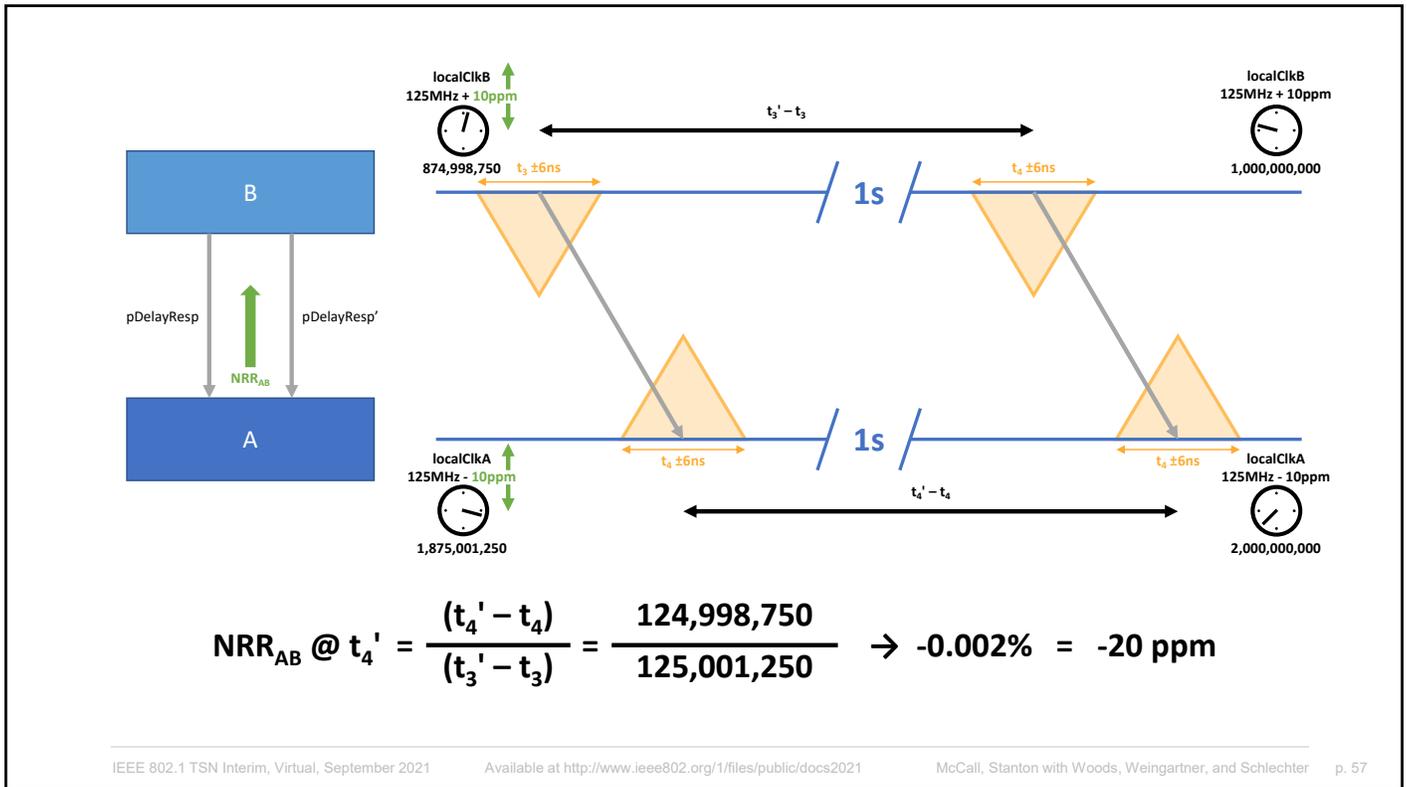


The Platonic ideal...

Note that A uses pDelayResp messages from B to A measure the Rate Ratio between the two (NRR_{AB}). These are the same pDelayResp messages that are used to measure Link Delay between A and B. (I'm just not showing the pDelayReq messages). Also note that the intent here is to measure NRR at t_4 .



Put some realistic number on it...



These are the errors that we could model.

Timestamp Granularity Error & Dynamic Timestamp Error (shown combined)

Clock Drift (not the +10ppm and -10ppm...but the fact that those ppm values are changing over time, i.e. at the left end of the timeline, the points they **aren't** +10ppm and -10ppm)

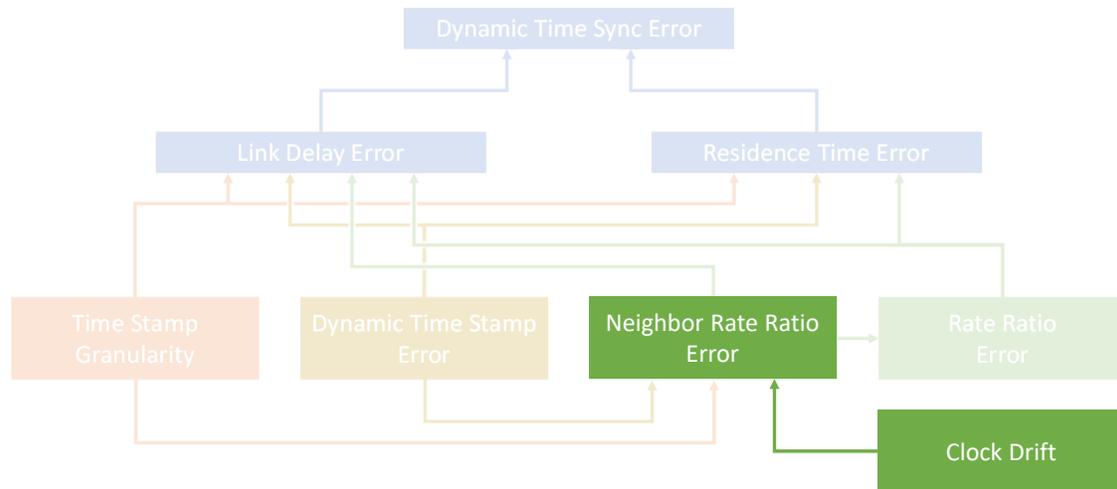
What factors to model?

Factor	YES / NO	
Time Stamp Granularity & Dynamic Time Stamp Error	NO	$\pm 20\text{ns}$ over $1\text{s} = \pm 0.02\text{ppm}$ Will be swamped by other errors.
Drift Rate of Local Clocks A & B (ppm/s) (Difference between Local Clock A ppm/s and Local Clock B ppm/s at start)	YES	For most recent temp profile, max possible difference between two clocks' rate of change of ppm is $\pm 1.2\text{ppm/s}$. $\pm 1.2\text{ppm/s}$ over $1\text{s} = \pm 1.2\text{ppm}$ May be a significant contributor to overall dynamic time sync error.
Rate of change of Drift Rate of Local Clocks A&B (Difference between Local Clocks' ppm/s at start & finish.)	NO	For most recent temp profile, max rate of change of ppm/s (i.e. ppm/s per second) is $\sim \pm 0.006\text{ppm/s}^2$. Rounding up to $\pm 0.01\text{ppm/s}^2$, over 1s (or 5s) the change is not significant. (Note: ignores discontinuities at start and end of temperature ramp, which are not realistic.)

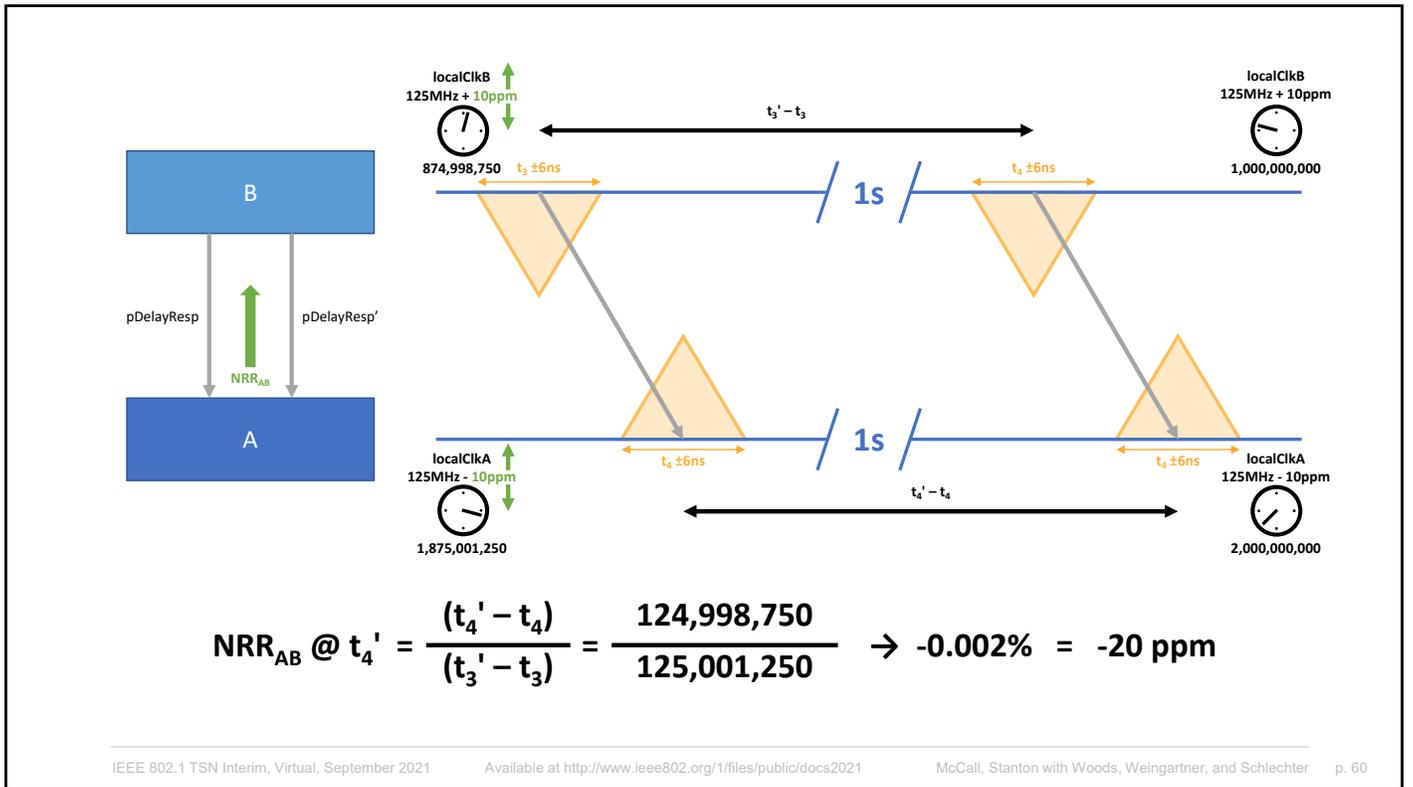
It turns out that the Timestamp errors will be small relative to other errors and can be ignored.

The same applies to the rate of change of clock drift (in ppm/s^2), which is low enough that, over the periods of time we are interested in, clock drift can be treated as constant, i.e. linear.

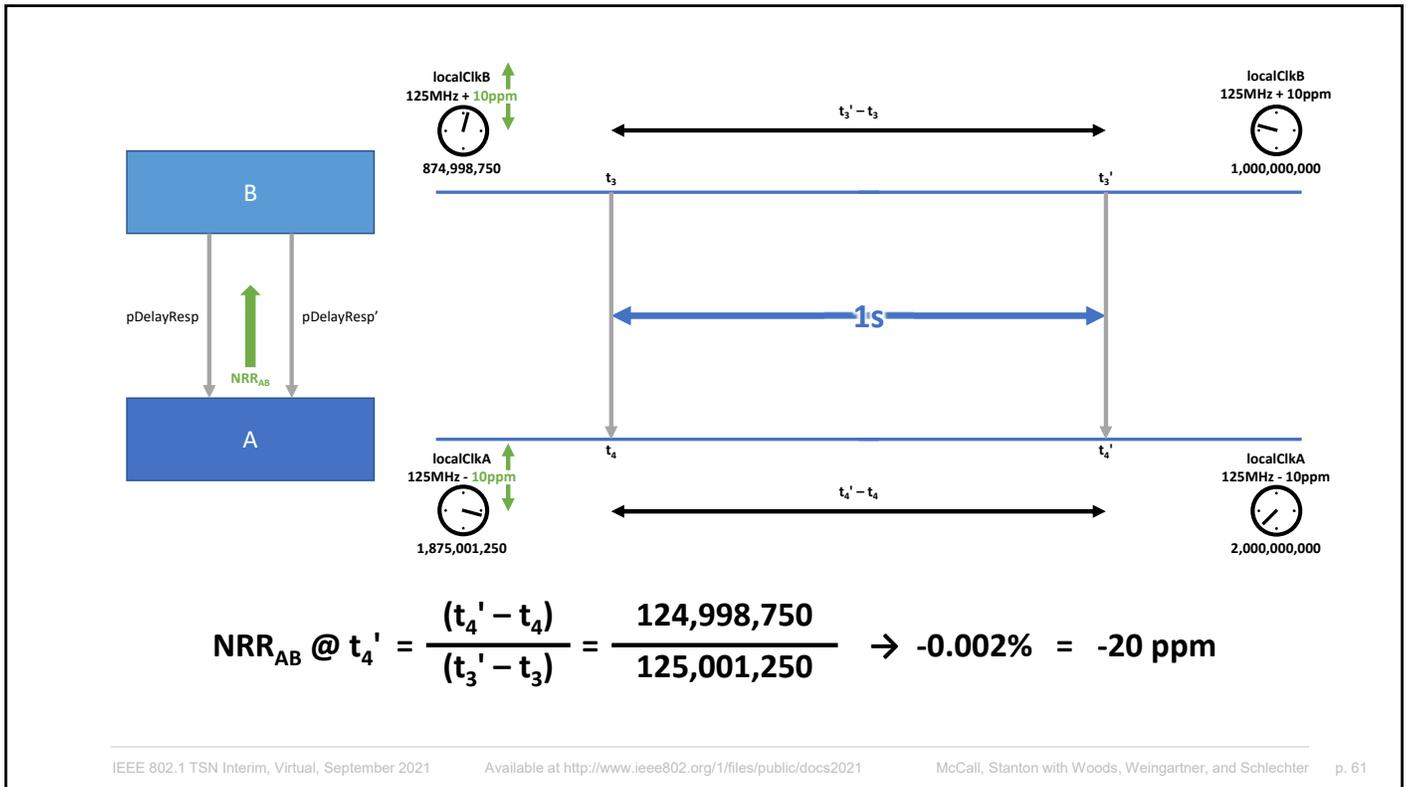
Sources of Dynamic Time Sync Error



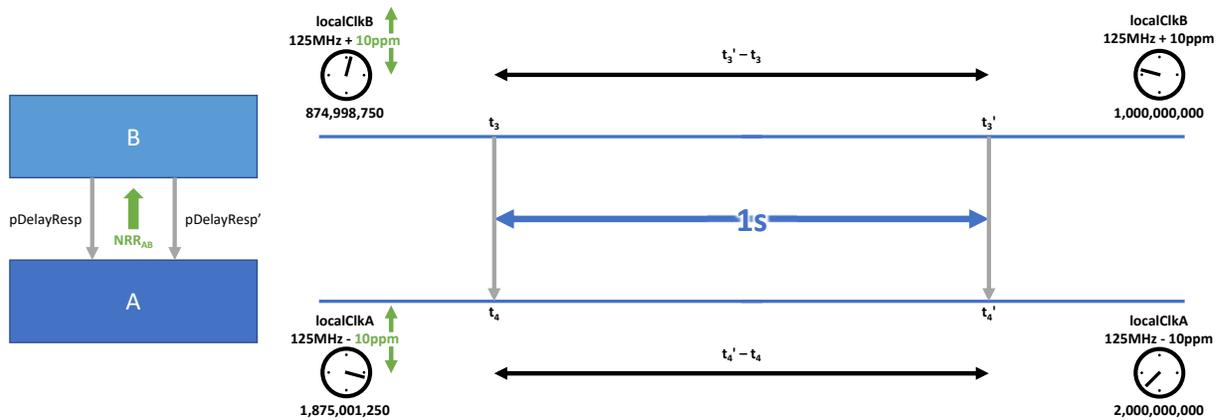
So, we only need to look at Clock Drift...



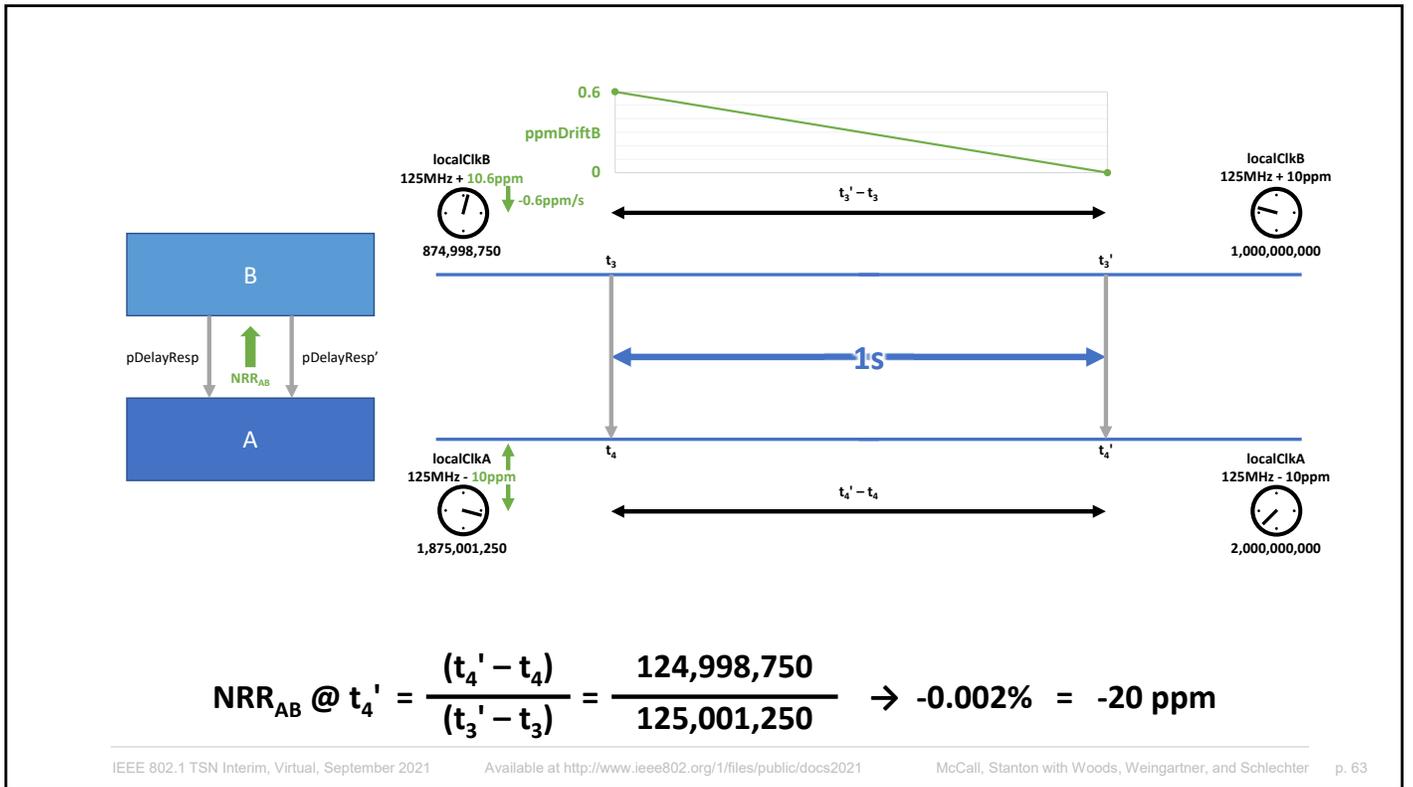
Which means we can remove the Timestamp Errors and also modify the diagram to remove the 1s discontinuity in the middle so that we're looking at the appropriate scale...



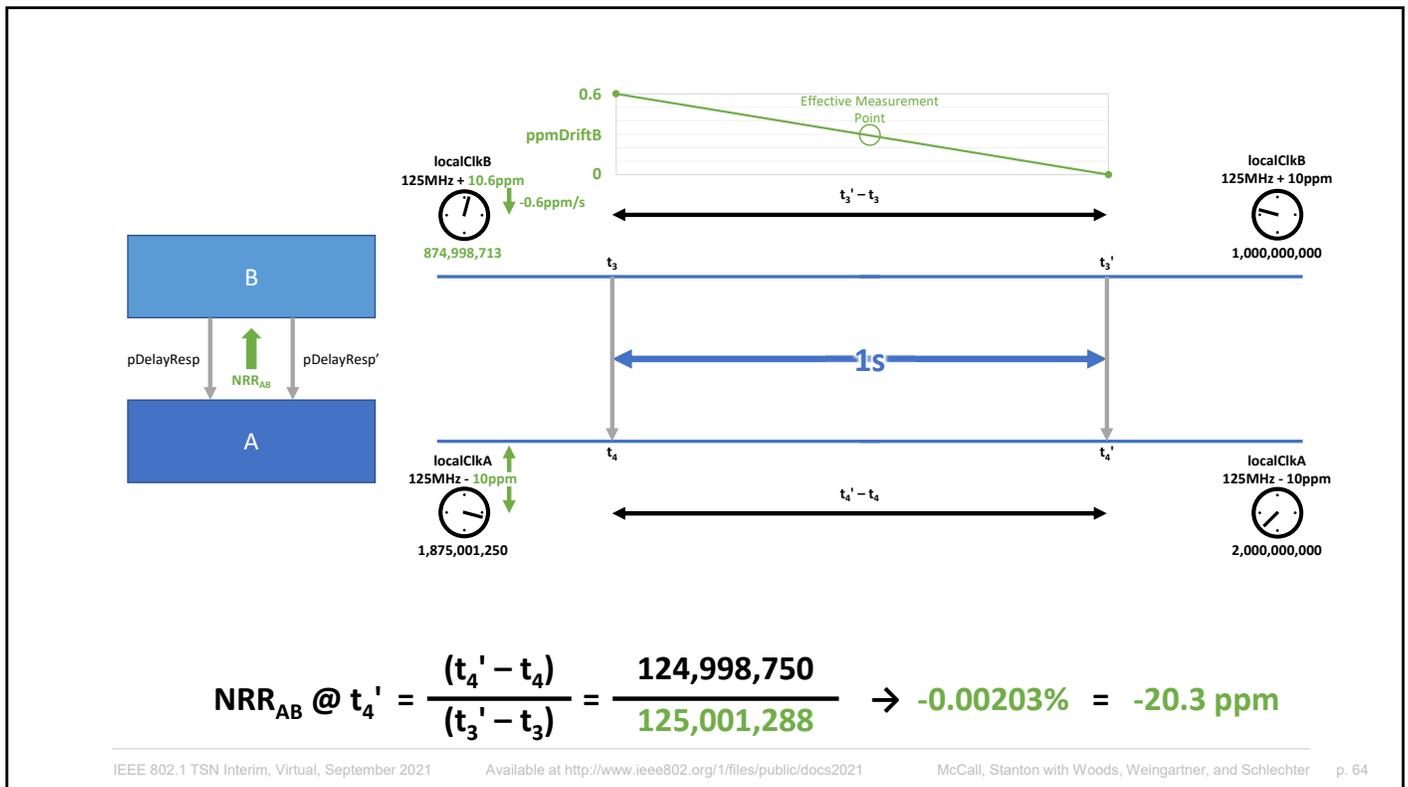
Then squeeze things together to make a bit of room...



$$NRR_{AB} @ t_4' = \frac{(t_4' - t_4)}{(t_3' - t_3)} = \frac{124,998,750}{125,001,250} \rightarrow -0.002\% = -20 \text{ ppm}$$



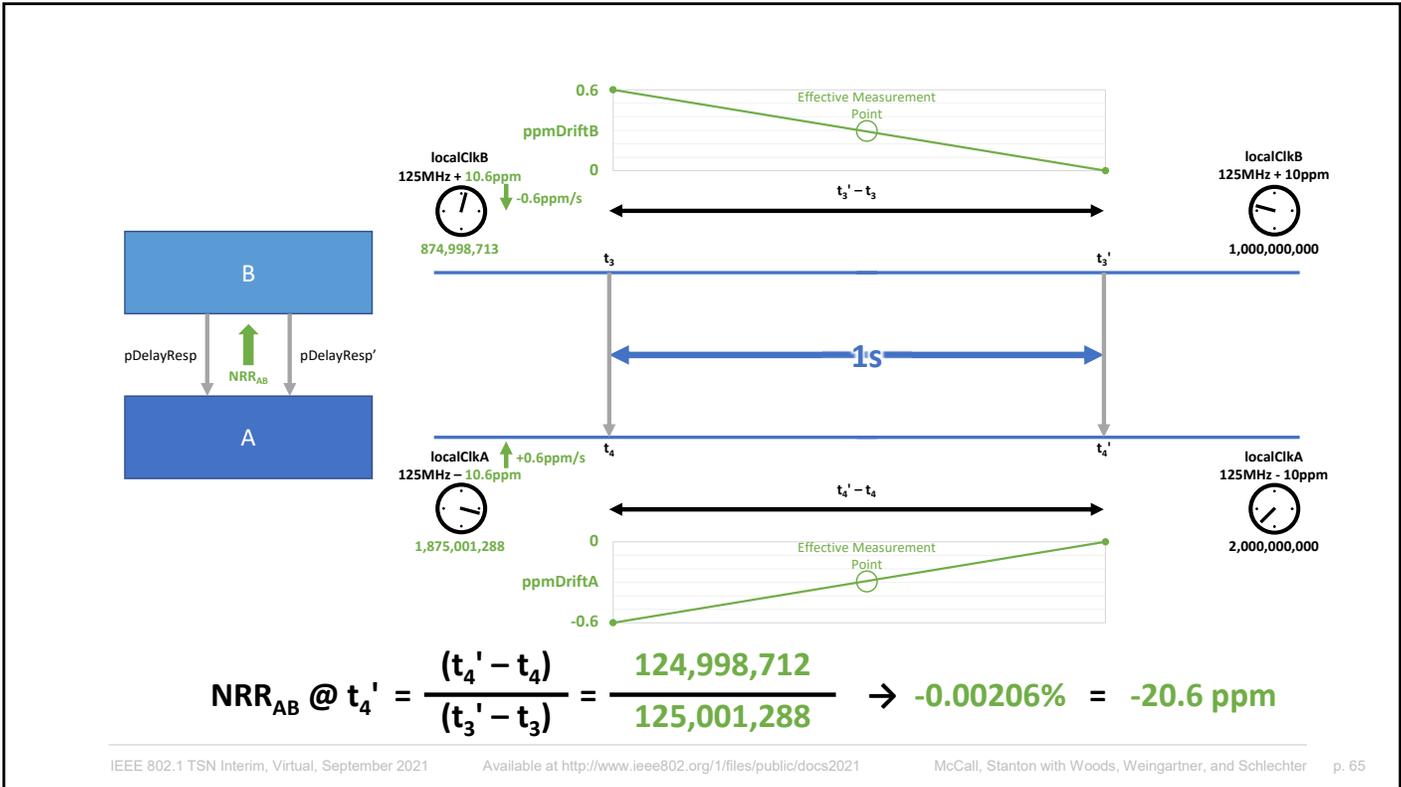
First... add a -0.6ppm drift on B's local clock.



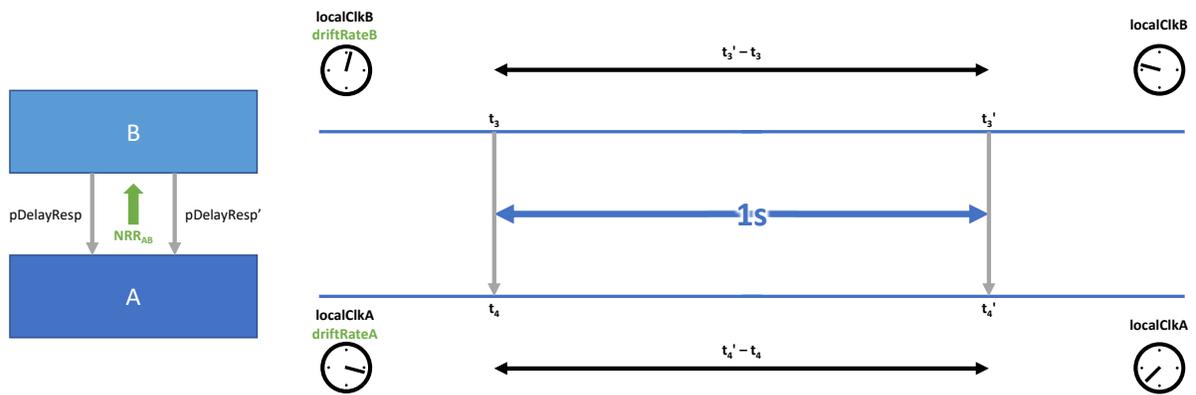
This means that B's local clock was running faster at the left of the timeline than at the right...specifically 0.6ppm faster.

Which means that there were more than 125,001,250 clock ticks between t_3 and t_3' ...which means the fraction is smaller...by 0.3ppm.

Only half of the ppm drift shows up in the NRR because you need to integrate under the curve...or rather line, because we can treat clock drift as linear. Another way of thinking about this is that we're effectively measuring NRR half way between the current and the previous pDelayResp.



Adding in a positive clock drift to A's local clock has a similar effect.



$$\begin{aligned}
 \text{mNRR}_{\text{error}} \text{ (ppm)} &= \frac{(\text{driftRateB} - \text{driftRateA}) * (t_2' - t_2)}{2} \\
 \text{(Difference between measured and actual Neighbor Rate Ratio @ } t_1) & \\
 \text{(For simulation)} &= \frac{(\text{driftRateB} - \text{driftRateA}) * \text{pDelayInterval}}{2}
 \end{aligned}$$

Neighbor Rate Ratio Error - Observations

- There is a difference between actual NRR and measured NRR (mNRR)
- When clock rates are stable mNRR is very accurate
 - Almost no difference between real NRR and mNRR
 - No large impact from timestamp granularity or errors
 - No benefit from averaging mNRR (or taking a median of previous values) if (Timestamp Errors / pDelay Interval) is negligible in terms of ppm (it actually makes things worse)
- When clock rates are drifting mNRR may have significant errors
 - Significant difference between real NRR and mNRR
 - Minimised when clock rates of neighboring devices drift in the same direction
 - Maximised when clock rates of neighboring devices drift in different directions
 - PPM vs. Temp curve for XO means the latter could be more common than expected
- There is the possibility of compensating for Neighbor Rate Ratio Error

Modelling Neighbor Rate Ratio – 1

- Use formula from earlier to model NRR_{error} at $pDelayReq...$

$$mNRR_{error} = \frac{(\text{driftRateB} - \text{driftRateA}) * pDelayInterval}{2}$$

- Model driftRate as stable (0ppm) or varying ± 0.6 ppm (linear uniform distribution)
 - X% stable; 100-X% varying
- Add “Error Correction Factor” to account for modelling of drift (NRR_{drift}), correction of NRR ($cNRR$), and correction of other calculations via use of $aNRR$
 - One value to cover all errors related to NRR_{error} and drift.
 - 25% Effective → Removes 25% of the error.

Modelling Neighbor Rate Ratio – 2

- Actual error when NRR is used (NRR_{error}) will be larger than $mNRR_{error}$ if time has passed since measurement, allowing clocks to drift further apart. Relevant for Residence Time Error.

$$NRR_{error} = mNRR_{error} + (\text{driftRateB} - \text{driftRateA}) * \text{time Since mNRR}$$

- Model time Since mNRR as uniform distribution with min=0, max=pDelay.
- Apply same NRR_{error} correction factor to second term
 - 25% Effective → Removes 25% of the error.

Current Simulation

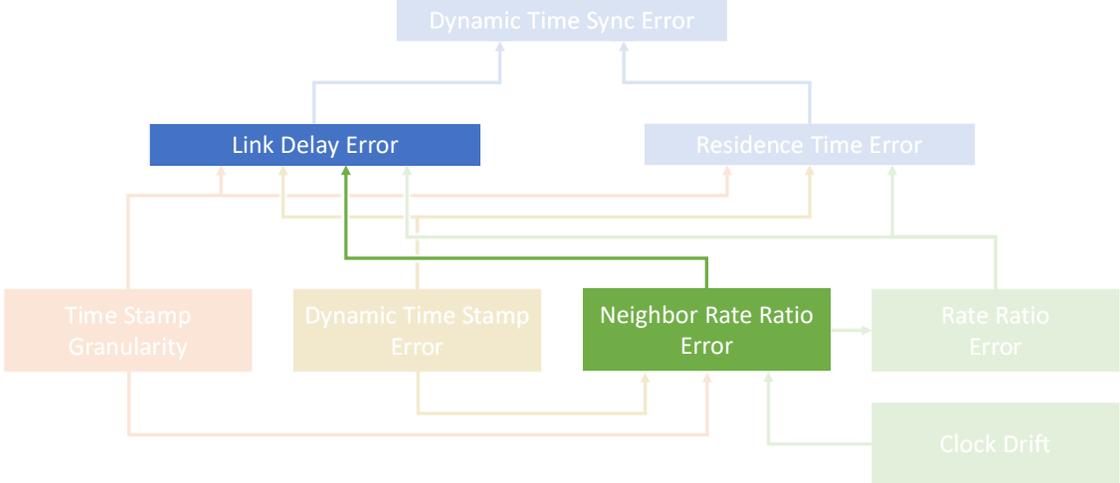
- Mean pDelay Interval: 31.25ms
 - Would mean $mNRR_{error} = \pm 0.009375\text{ppm}$ with no average / median calculations, but...
- NRR is measured based on t_3 & t_4 from most recent pDelayResp message...and t_3 & t_4 from the pDelayResp 7 pDelay Intervals prior. The median of the most recent measurement and the 6 prior measurements is then used as mNRR.
 - Using t_3 & t_4 from the pDelayResp 7 pDelay Intervals prior pushes back the effective measurement point to 7/2, i.e. 3.5 pDelay Intervals prior to mNRR (vs. 0.5 pDelay Intervals prior, i.e. 7x worse).
 - As rate of change of clock drift is (ppm/s²) is low, clock drift can be treated as linear, so the median will be the NRR from 3 pDelay intervals prior, i.e. the effective measurement point is an additional 3 pDelay Intervals prior.
 - Worst case $mNRR_{error}$ due to drift is therefore 13x worse: $\pm 0.12\text{ppm}$
- pDelay Interval is short enough that Timestamp errors have an appreciable effect
 - $\pm 24\text{ns}$ over 31.25ms is $\pm 0.77\text{ppm}$
- Worst Case $mNRR_{error}$: $\pm 0.89\text{ppm}$
- Effect of Clock drift on calculations made between NRR measurements is low due to low pDelay Interval
 - $\pm 1.2\text{ppm}$ over 31.25ms is $\pm 0.0375\text{ns}$

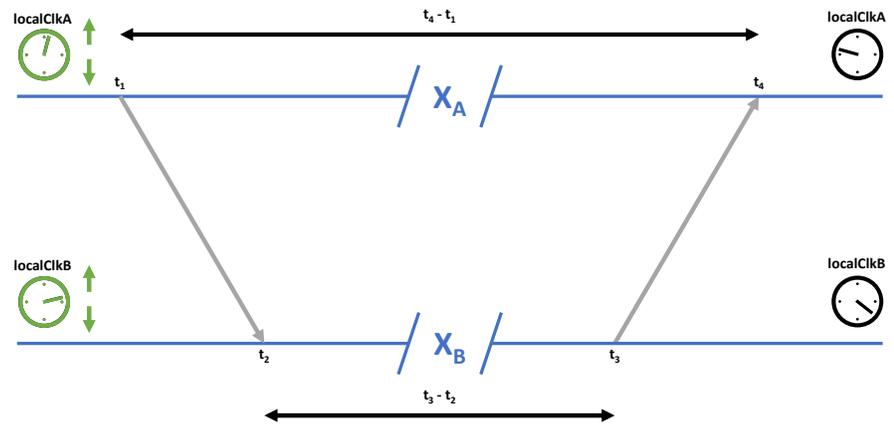
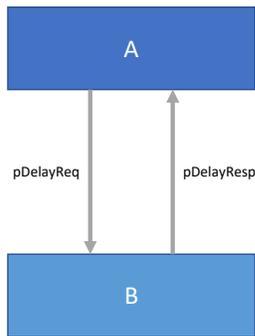
Formulae –NRR_{error}

$$mNRR_{error}(n) = \text{driftRateError}_{correction} \times \left(\frac{pDelayInterval}{2} \right) (\text{clockDrift}(n-1) - \text{clockDrift}(n))$$

Link Delay Error – Neighbor Rate Ratio Error

Time Sync – Elements & Relationships





$$pDelay = \frac{(t_4 - t_1) - mNRR(t_3 - t_2)}{2}$$

$$pDelay_{errorNRR} (ns) = \frac{mNRR_{error}(t_3 - t_2)}{2}$$

$$p\text{Delay}_{\text{errorNRR}} (\text{ns}) = \frac{\text{NRR}_{\text{error}}(t_3 - t_2)}{2}$$

WORST CASE:

$$p\text{Delay}_{\text{errorNRR}} (\text{ns}) = \frac{10\text{ms} * (0.6\text{ppm})}{2}$$
$$= 3\text{ns}$$

Conclusion: Model $p\text{DelayError}_{\text{NRR}}$

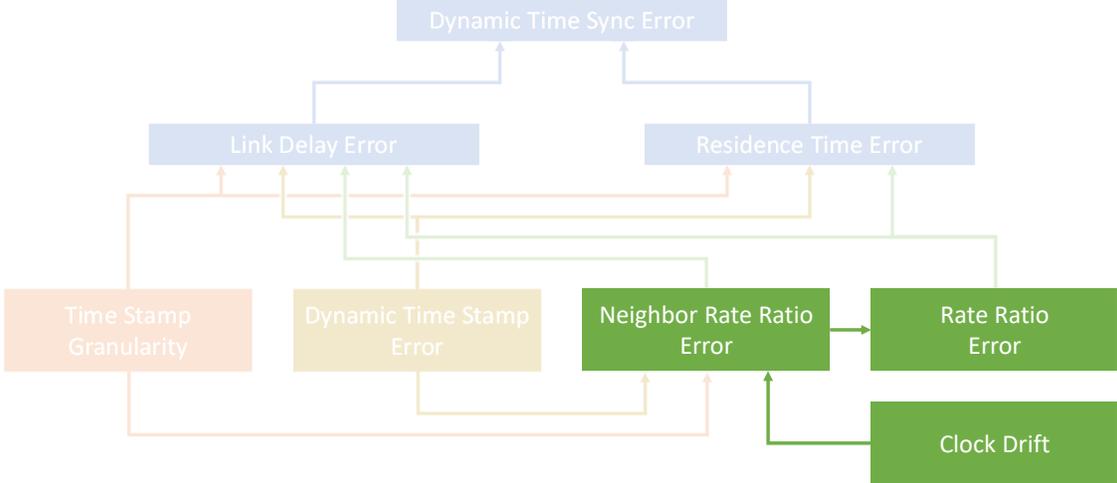
Formulae – pDelay_{errorNRR}

$$\text{meanLinkDelay}_{\text{error}} = (1 - \text{meanLinkDelay}_{\text{errorCorrection}})(\text{pDelay}_{\text{errorTS}} + \text{pDelay}_{\text{errorNRR}})$$

$$\text{pDelay}_{\text{errorNRR}} = \text{mNRR}_{\text{error}} \left(\frac{\text{pDelayResponse}}{2} \right)$$

Rate Ratio Error

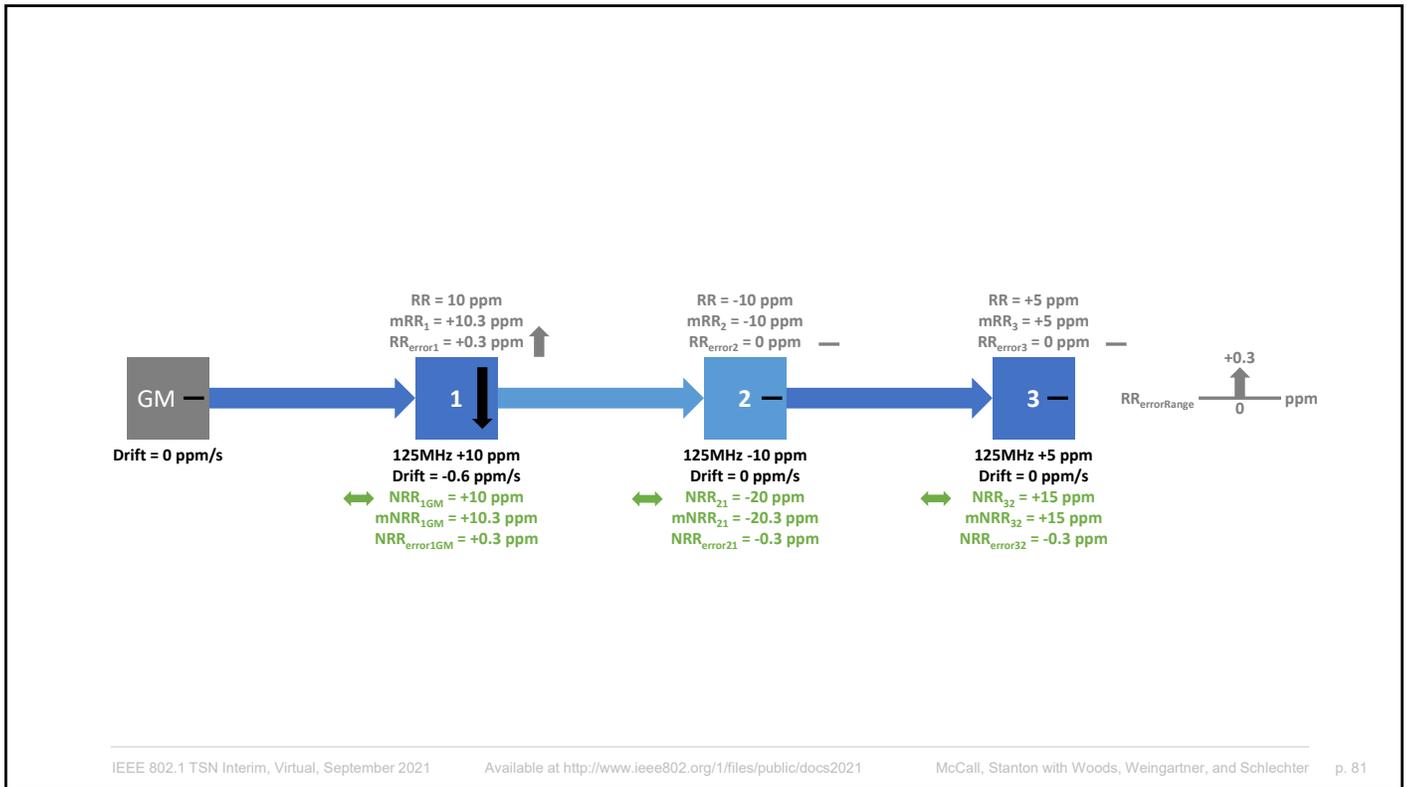
Time Sync – Elements & Relationships



$$\text{RateRatio}_n = \text{RateRatio}_{n-1} * \text{NRR}$$



Start with a “simple” analysis.



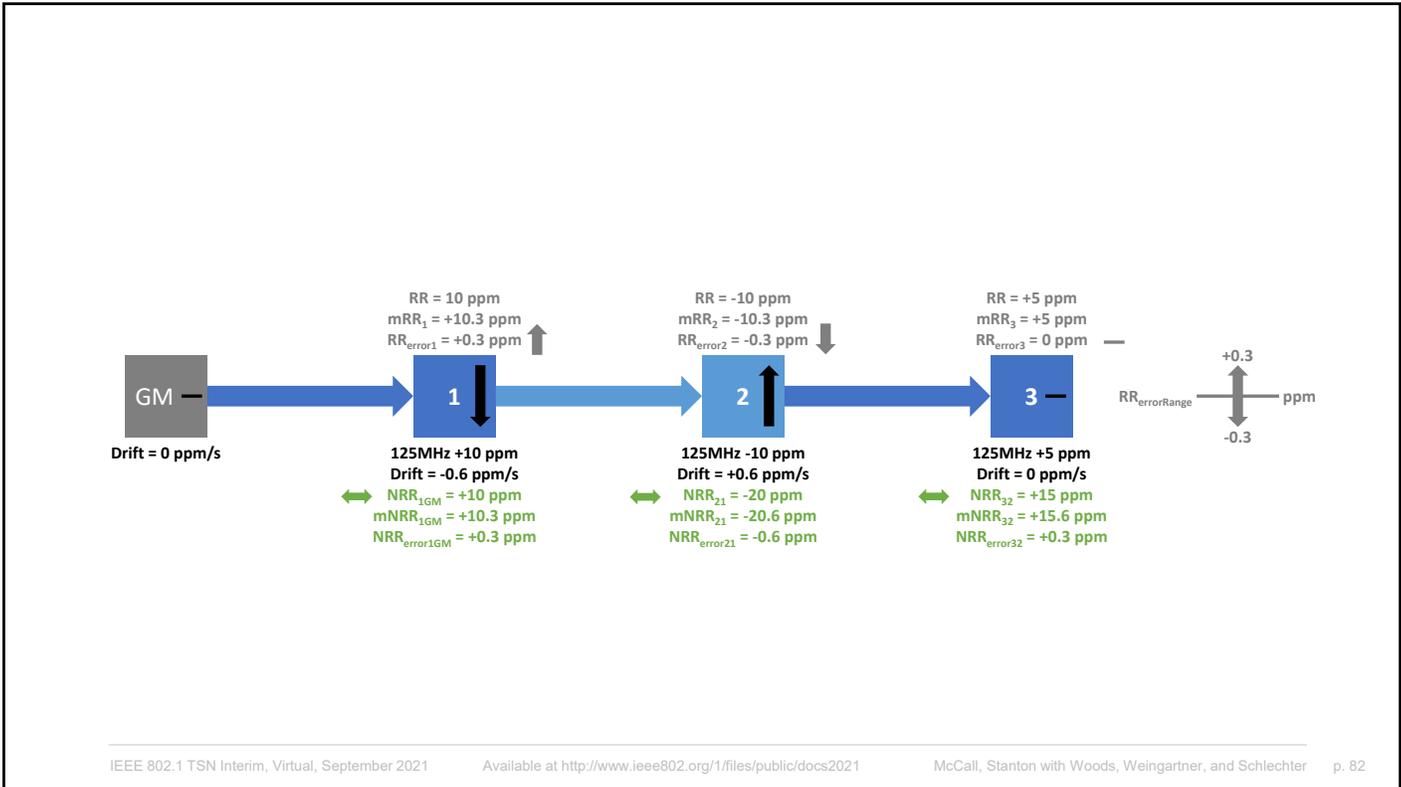
Simple in that...if you feed NRR directly into RR, and don't account for time between measuring mNRR and using mNRR, i.e. applying it to Rate Ratio.

For pDelay Interval = 1s

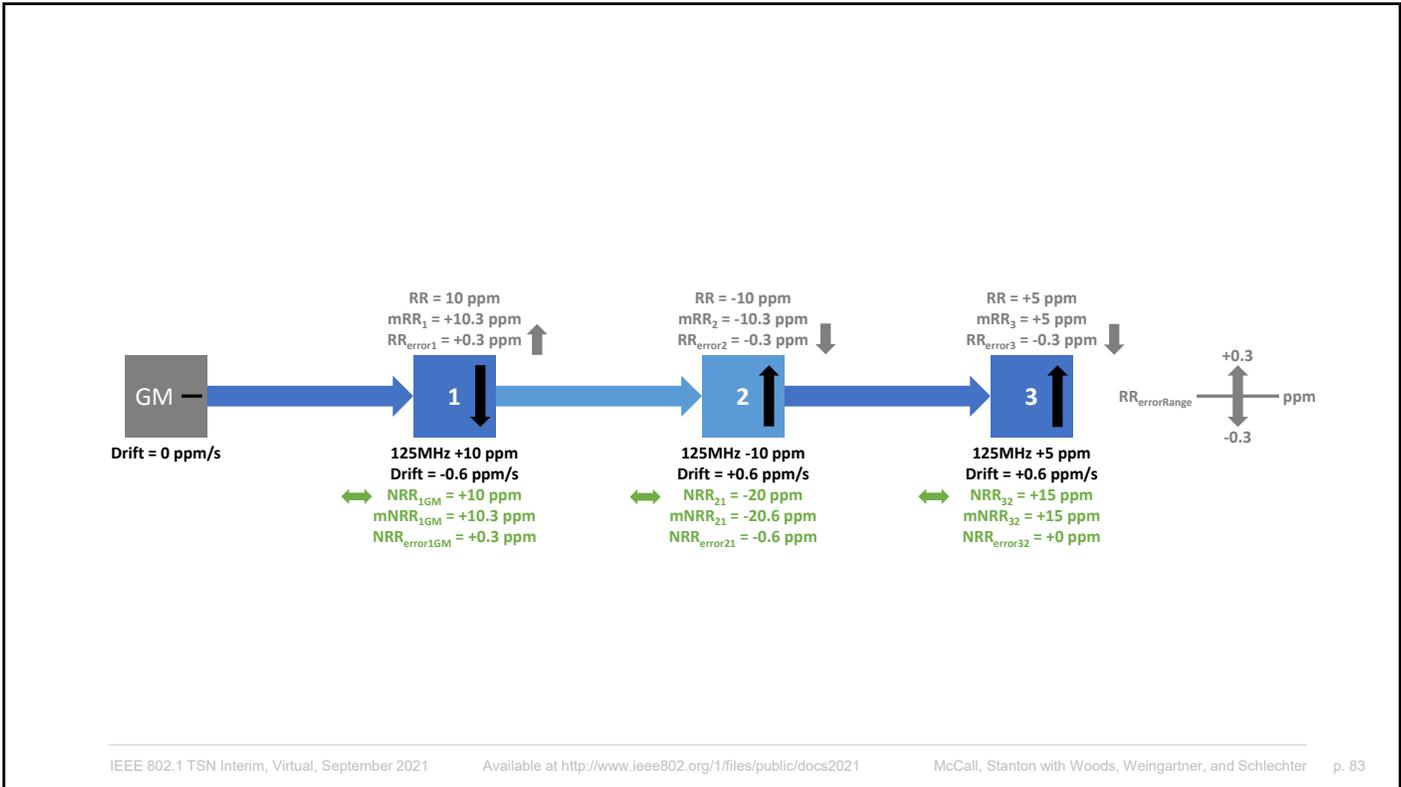
...which is the 802.1AS-2020 default value, and the value used in this presentation when looking at NRR error.

See earlier in the presentation for full analysis.

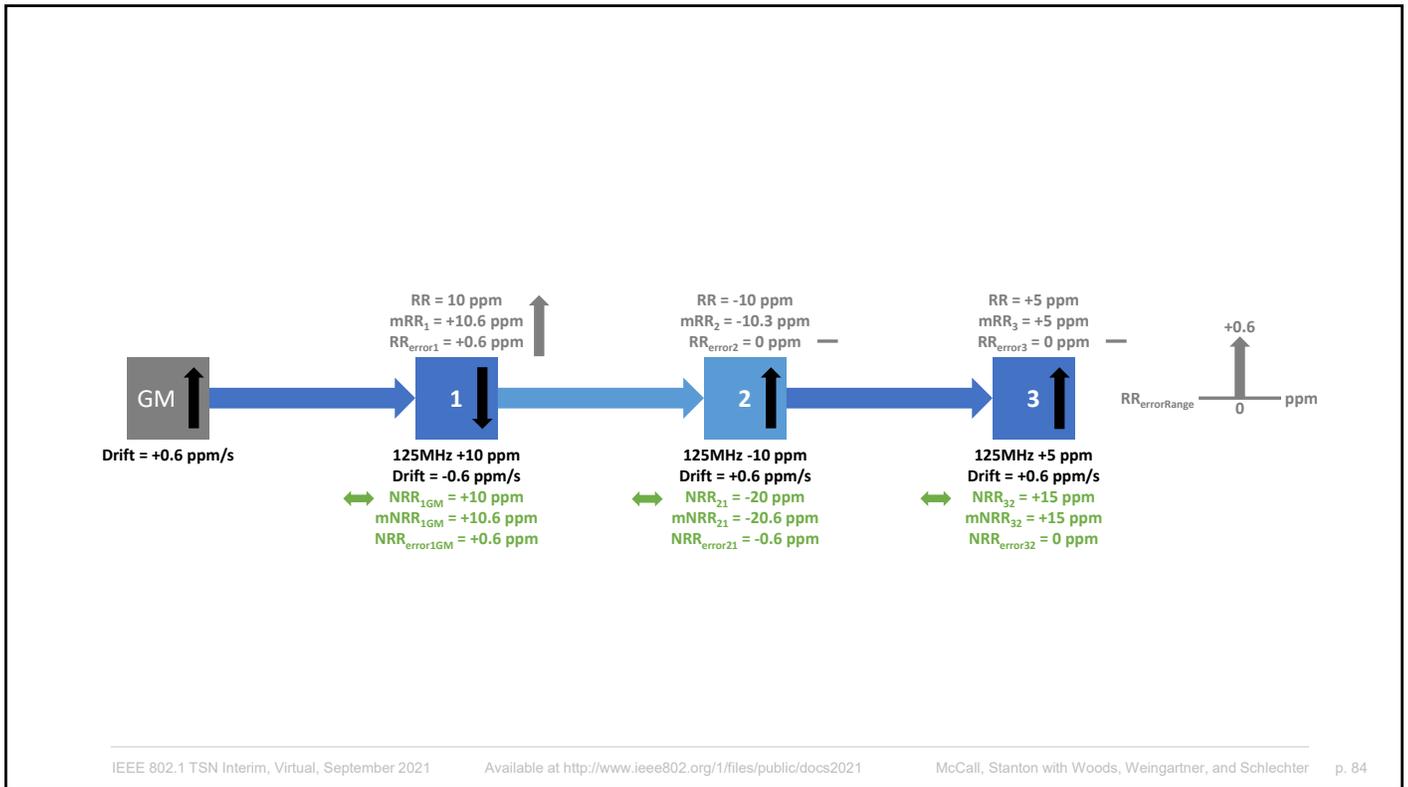
The RR_{error} at 1 is fully removed at 2. This makes sense as RR is trying to calculate the ratio of Local Clock to GM clock. If everything is working correctly, the error will depend on the drift between the two...not the drift of clocks in the transmission path.



This continues to line up as we add Clock Drift to 2...



...and 3.



Adding Clock Drift to the GM is very different because it affects the Rate Ratio Error at every single downstream device.

Worse, it pushes them all in one direction, which could have further implications...when you start adding in the next layer of Rate Ratio Error...

$$RR_{\text{error}} \text{ (ppm)} = \frac{(\text{driftRate}_{\text{GM}} - \text{driftRate}_{\text{local}}) * \text{pDelayInterval}}{2} \quad ?$$

Which looks a lot like the NRRerror equation, but with $\text{driftRate}_{\text{GM}}$ instead of driftRate_{n-1}

Note: Effect of GM clock drift skews RR_{error} for all downstream instances in the same direction.

This will have implications for how errors accumulate.

Therefore, model GM clock separately, i.e. separate ppm/s range vs. other clocks in the system.

(A more accurate GM clock could address up to 50% of basic Rate Ratio related errors in the system.)

Note that the latest time series simulations do not include any GM clock error.)

However, there is an additional source of error.

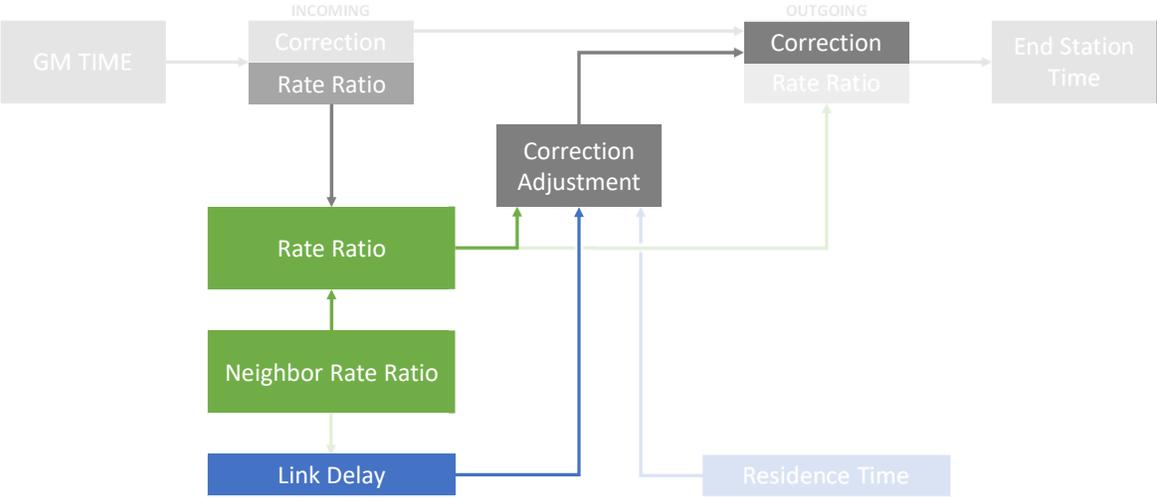
RateRatio is used during Time Sync processing.

RateRatio is effectively measured at the same time as NRR. There is a time delay between measurement and use...which is different at each node.

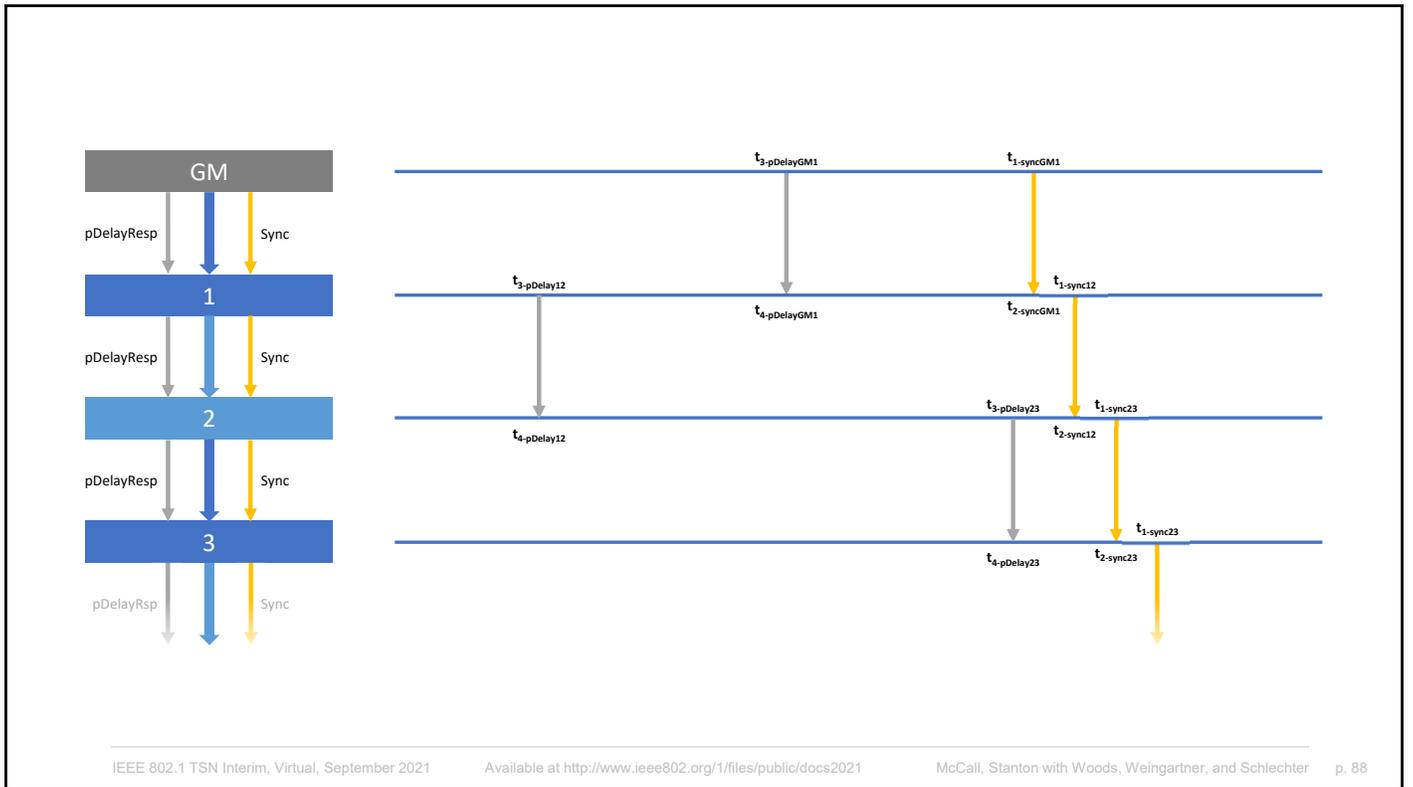
(NRR is measured as part of pDelay exchange.)

The “simple” analysis suggest that the equation at the top of the screen is sufficient.
But the time delay means we need to dig deeper...

Time Sync – Elements & Relationships

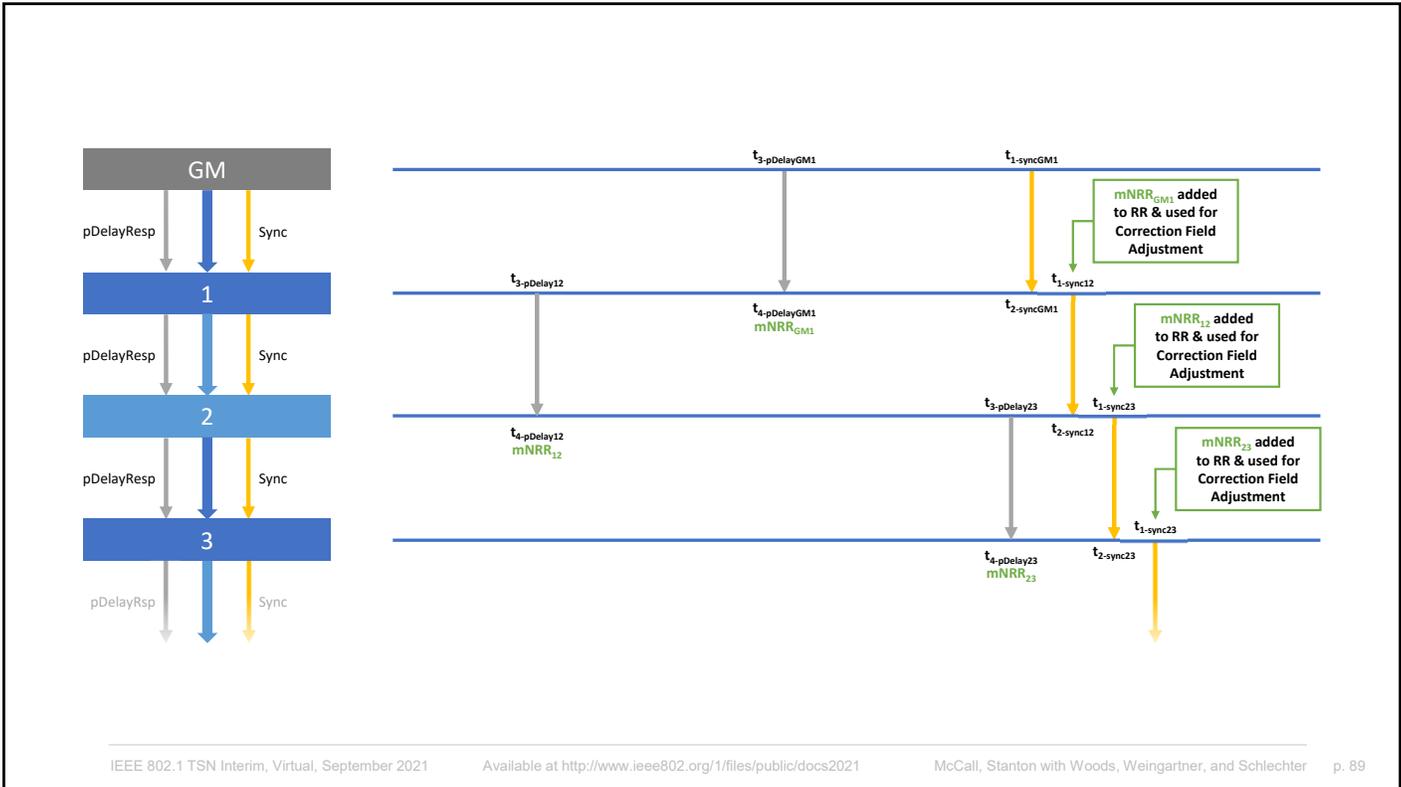




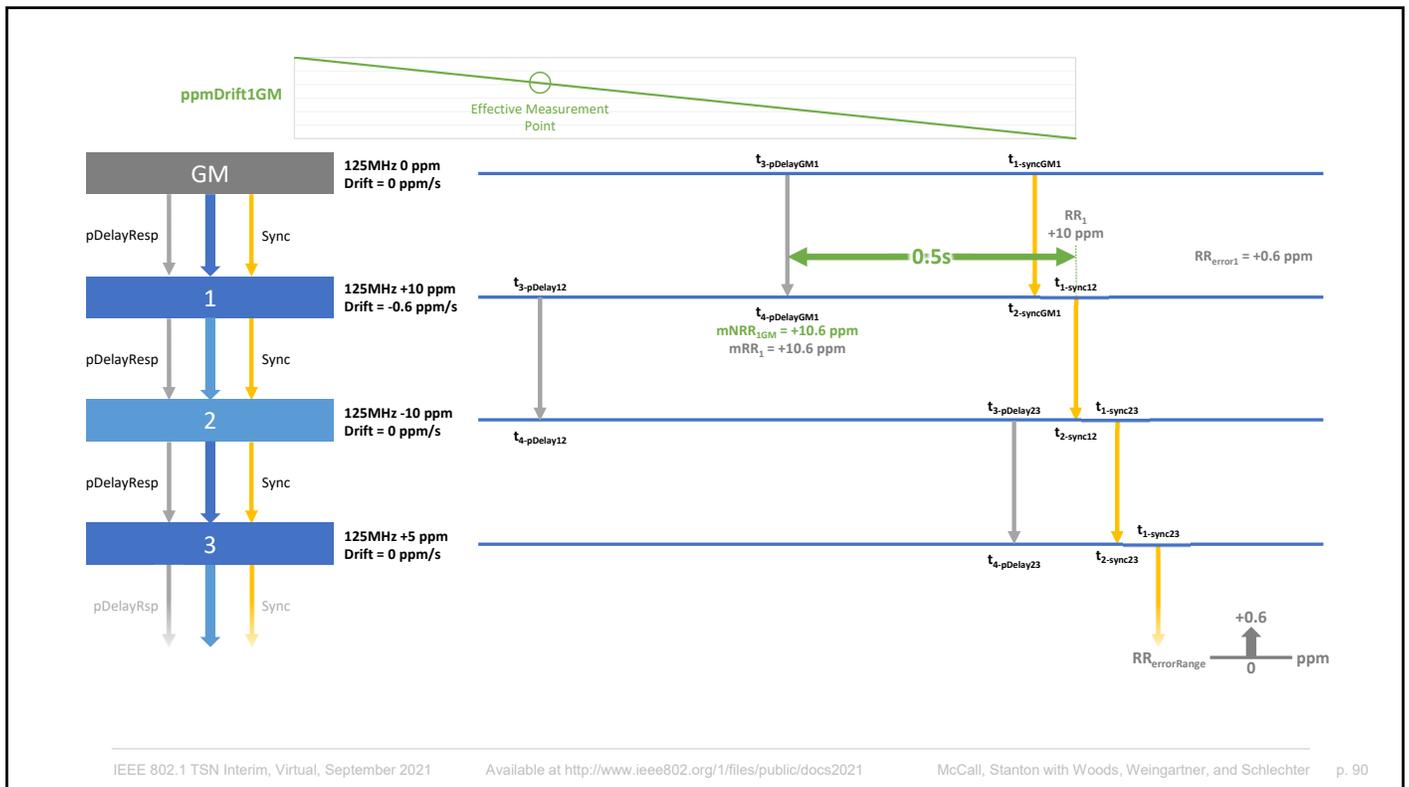


A more complex analysis includes different periods between pDelay and Sync messages. The minimum is close to zero (at scales of tens of hundreds of milliseconds) while the maximum is the pDelay Interval.

Note that in this graphic I've used an unrealistically long Residence Time to make it visible at this scale. The error analysis is unaffected.



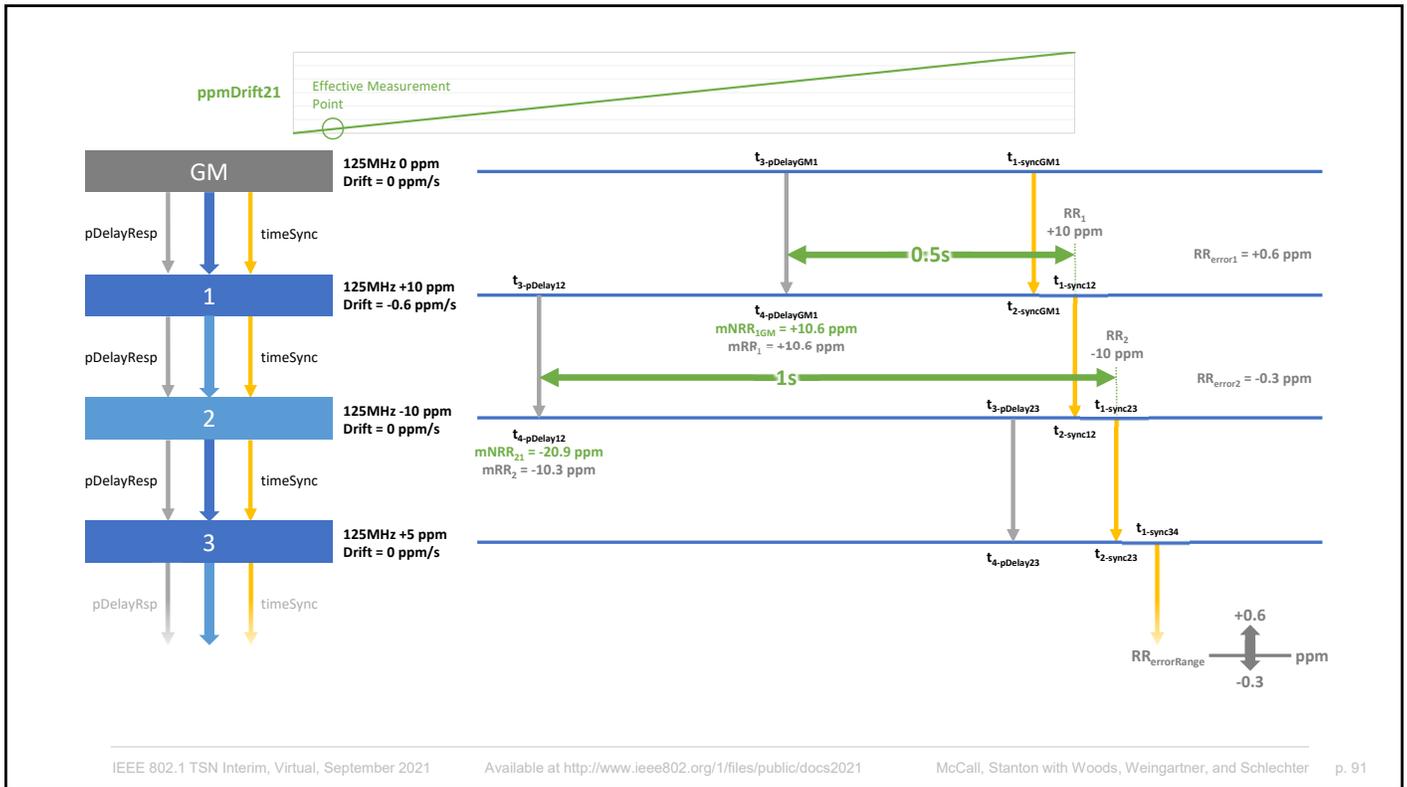
NRR is measured at last pDelayResp message (mNRR)...but added to Rate Ratio at end of Residence Time (Sync message TX). This means that the amount of adjustment to Rate Ratio (i.e. NRR) is effectively measured at mNRR.



Looking at the first step...

mNRR is effectively measured pDelay Interval / 2 prior to the incoming pDelay message. When pDelay Interval is 1 second, that's 0.5s prior to the mNRR measurement, i.e. it has an error of +0.3ppm.

But that is 0.5s prior to when the measurement is actually used...by which time there is an additional 0.5s of drift leading to an additional +0.3ppm of error, for a total of +0.6ppm.

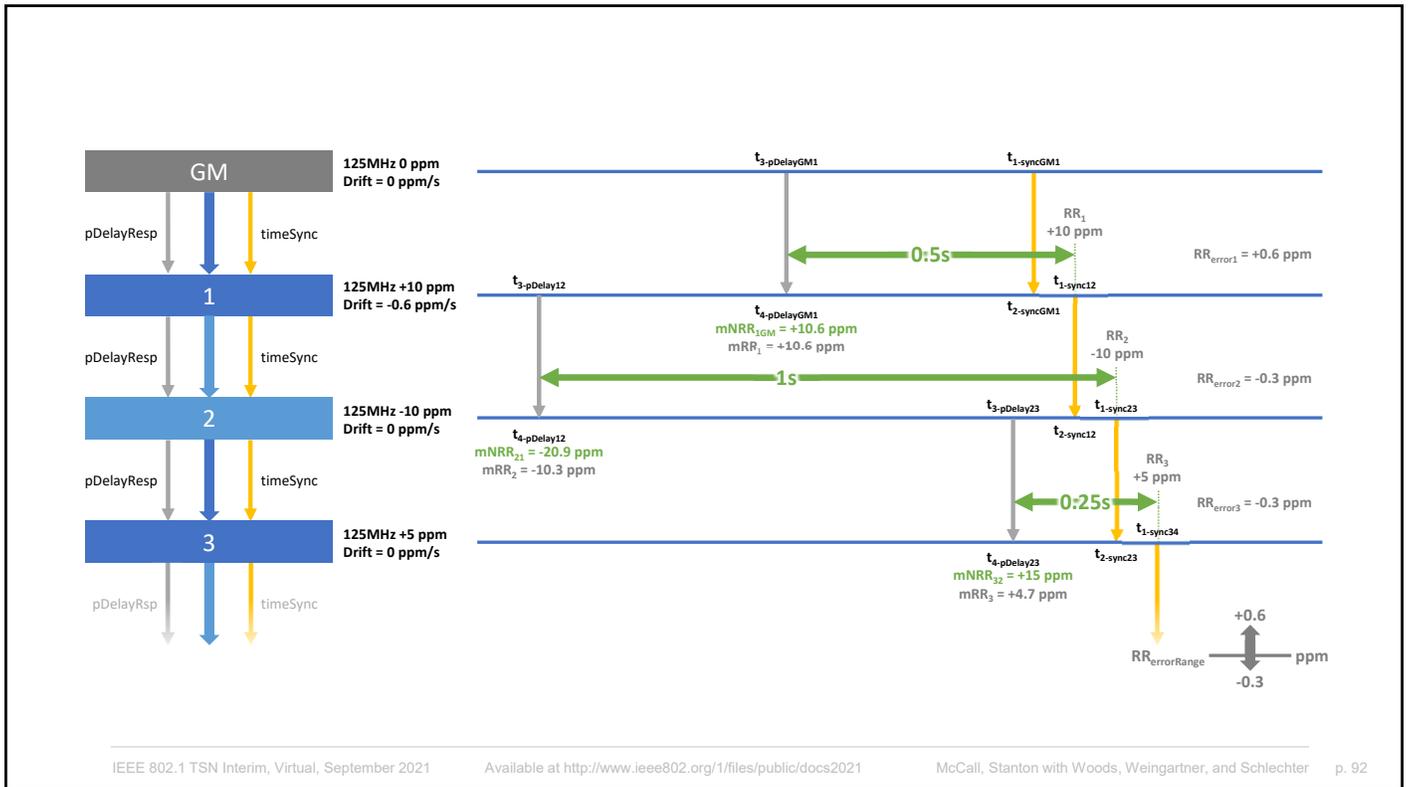


In the second step, the delay between measuring NRR and using the measurement is at it's maximum, i.e. a pDelay Interval of 1 second...which means the effective measurement point is 1.5 seconds prior to when the value is used.

The NRR error at the point when it's added to the Rate Ratio is therefore off by - 0.9ppm...which doesn't just cancel out the +0.6ppm error from the previous step, it overshoots by 0.3ppm.

Note that the calculations shown above do not explicitly include the effect of clock drift during Residence Time. However, this does not affect the error analysis.

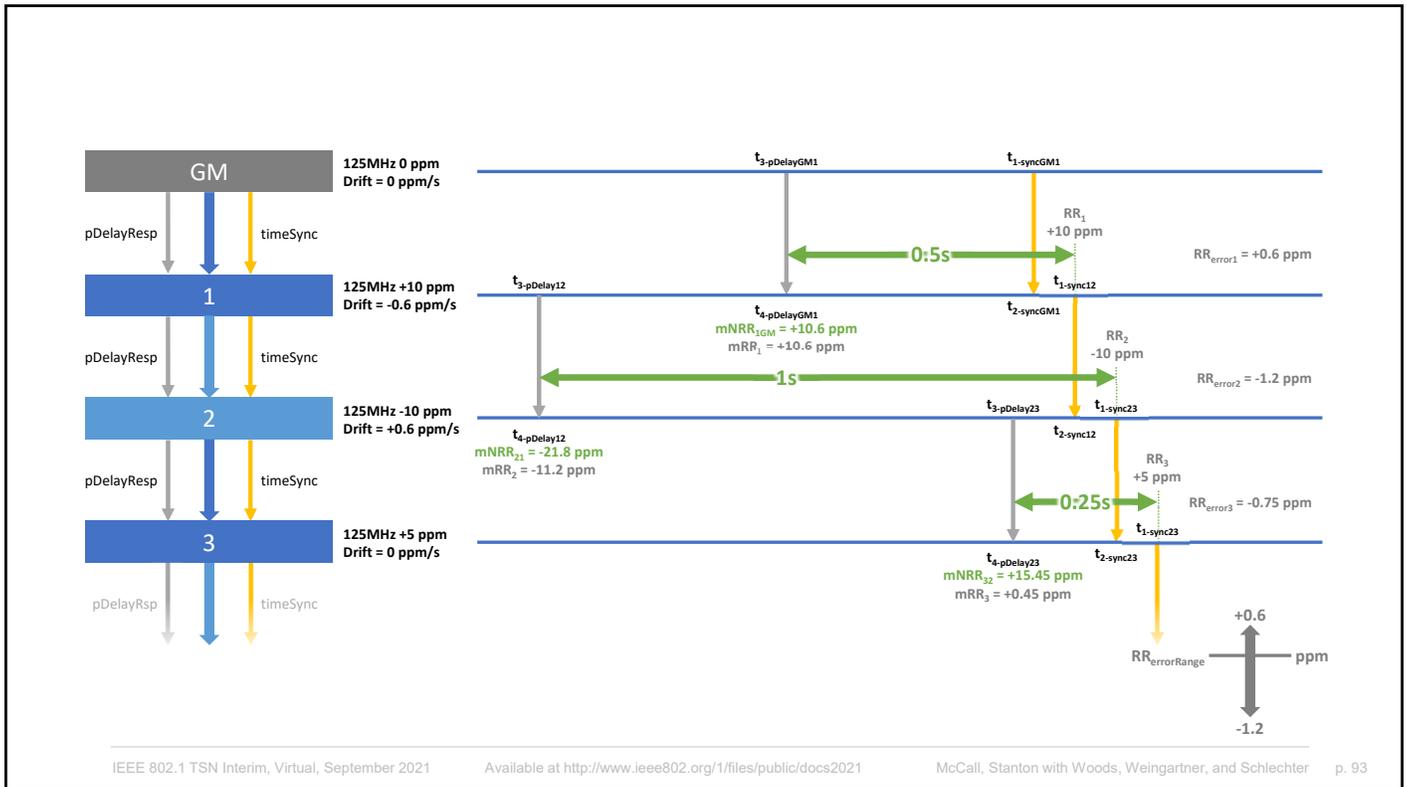
- Clock₁ will drift slightly (slower, since its drift rate is -0.6ppm) between t_{2-sync12} (where it's marked as +10ppm) and t_{1-sync23} where it's implicitly used in the calculation of RR₂ but...
- Clock drift over a typical Residence Time (default 802.1AS maximum is 10ms) is not significant, e.g. ±1.2ppm over 10ms is 0.012ppm. And even this can be ignored because...
- NRR_{error} and RR_{error} are driven by a clock's drift rate, not it's absolute value, and the rate of change of clock drift (ppm/s²) is low enough that, over the time periods of interest, clock drift can be considered linear. (See meanLinkDelay_{error} for details.)



There is no relative drift between 2 and 3, so there is no error in the NRR measurement...which means there is still nothing to cancel out the 0.3ppm error in Rate Ratio.

The basic mNRR errors from the previous section (the simple analysis) cancel out, because pDelay interval is constant (or close enough; if you want to be really accurate, you'd need to model variation in pDelay interval, but for the purpose of this error model, we'll assume there is no variation).

But the errors due to drift can remain in the system. This is the same inputs as for the simple analysis...but with added delay. In the simple analysis the $RR_{\text{errorRange}}$ was ± 0.3 ppm. In this case, it's already higher. And it can get worse.

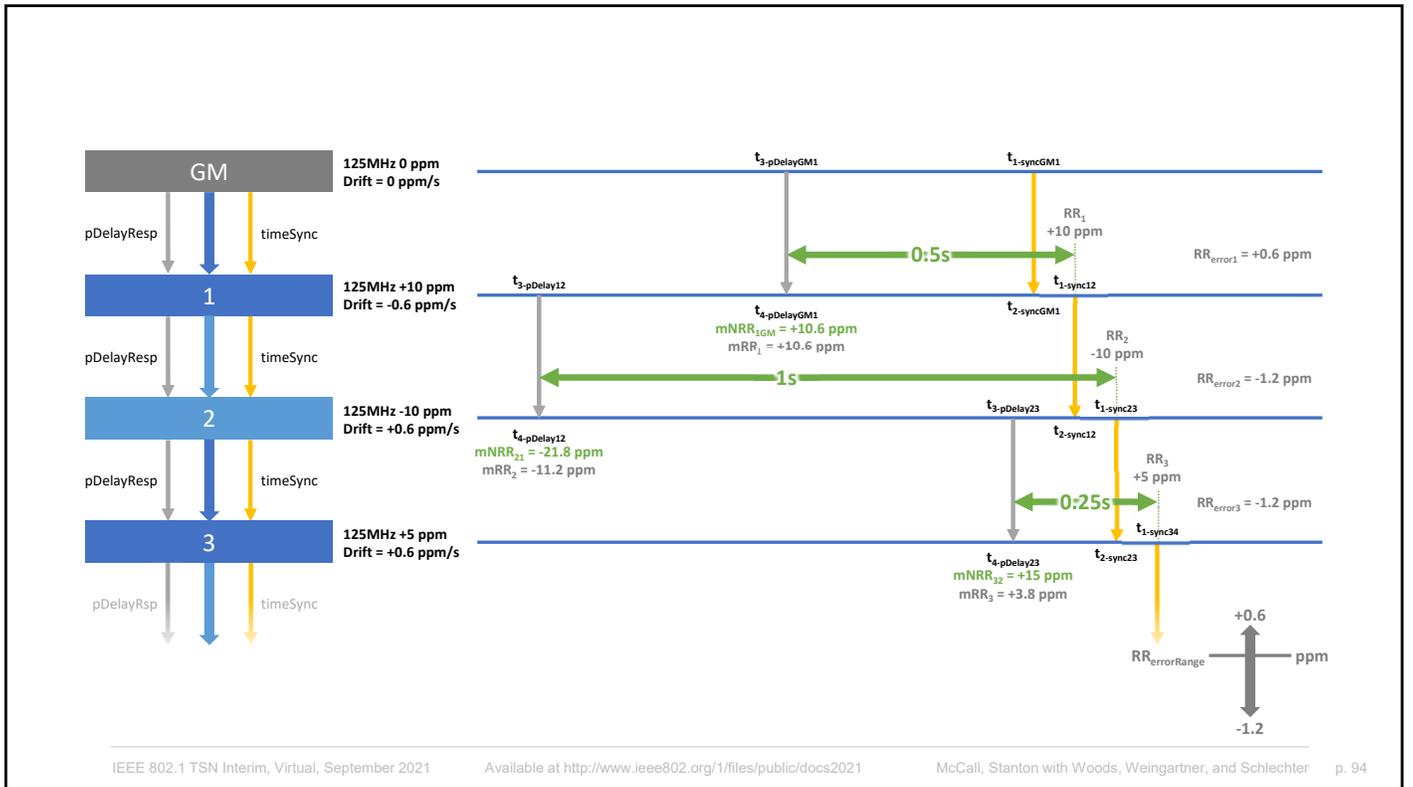


Adding a +0.6ppm/s drift to 2's clock means the mNRR is off by -1.2ppm (the difference between 2 and 1's clocks) multiplied by 1.5s, which is the time from the effective measurement point (pDelay Interval / 2, i.e. 0.5s) plus the delay between mNRR and when it's used (a full pDelay Interval in this case, i.e. 1s).

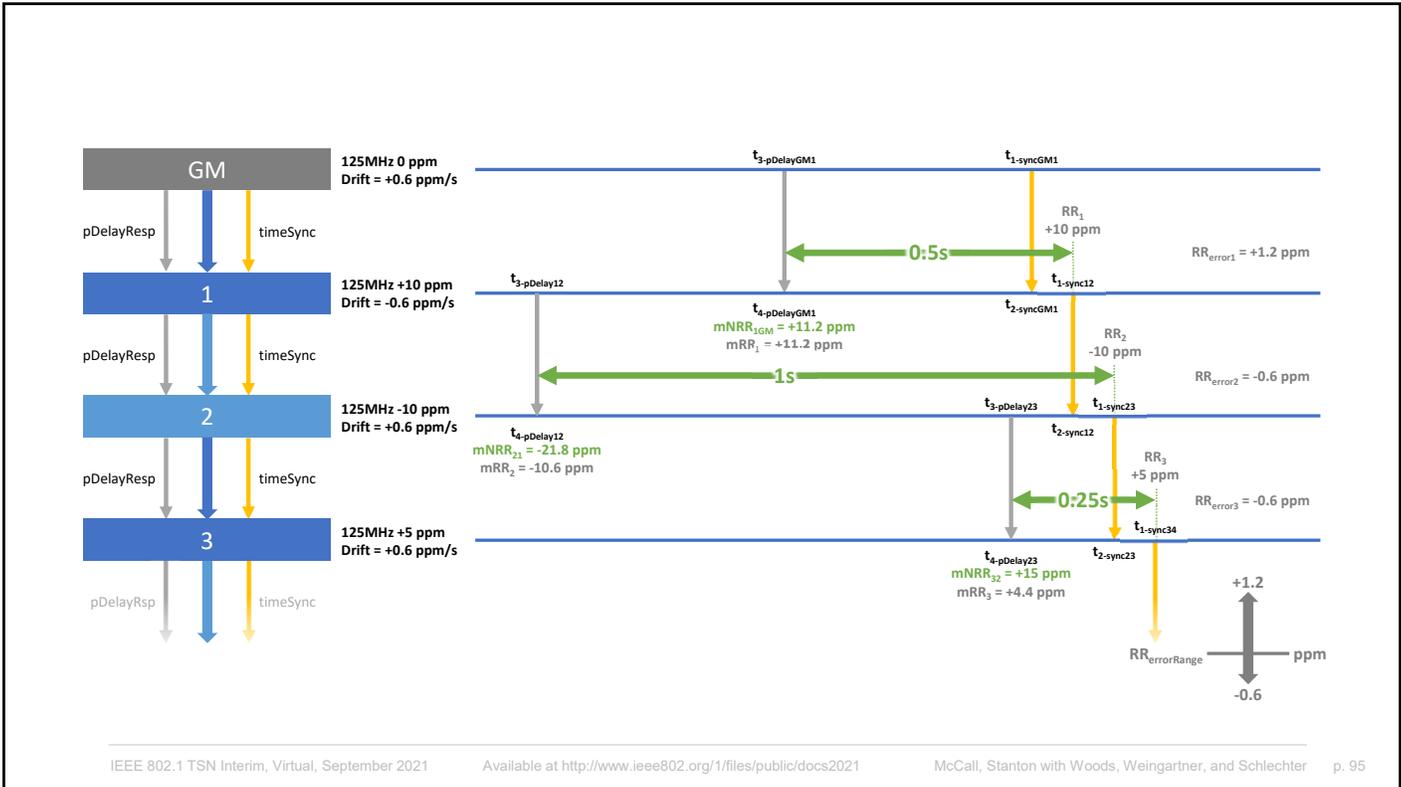
$-1.2\text{ppm} \times 1.5\text{s} = -1.8\text{ppm}$, so mNRR_{21} isn't -20ppm, or -20.6ppm, it's -21.8ppm.

On the previous slide mRR2 corrected mRR1 (of +0.3ppm) and overshoot by 0.3ppm to end up at -0.3ppm. In this case, with 2's clock also drifting...in the opposite direction to 1...and such a long delay between mNRR21 and the Sync message...the overshoot is now an additional 0.9ppm to end up at -1.2ppm.

Some of that is offset over the next hop by the difference in drift rate between 3 and 2. In the simple analysis of this scenario, $\text{RR}_{\text{error}2}$ was -0.3ppm and all of the error (i.e. 0.3ppm) was offset when moving on to $\text{RR}_{\text{error}3}$. In this scenario 0.45ppm is offset...which makes sense since the effective measurement point is 0.75s prior to when the measurement is used. But that still leaves -0.75ppm remaining, which will be passed down the chain of devices and, if every subsequent device has zero drift, will remain at the End Station.



Adding the +0.6pp drift to device 3 (as was done in the simple analysis) has the expected effect. There is no relative drift between 3 and 2, so NRR_{error} is zero and RR_{error3} remains the same as RR_{error2} .



Note that a GM Clock Drift has a similar effect as in the simple analysis: it pushes every subsequent Rate Ratio measurement off in one direction, but in this case the effect is amplified by the initial delay between $mNRR_{1GM}$.

Observations

- Effect of RR_{error} can be significant
 - See section on effect of RR_{error} Residence Time Error for details
- RR_{error} due to $mNRR_{error}$ cancels out at the next node
- RR_{error} due to additional Clock Drift between when NRR is measured ($mNRR$) and then used as part of Time Sync does not necessarily cancel out.
 - Can accumulate to some degree. Should average to zero, unless...
- GM clock drift affects every device in the chain, and in the same direction, i.e. affects the average
- Less time between pDelay message (used to measure NRR) and TimeSync is better
 - Lower pDelay Interval is one method
 - If pDelay Interval is $<$ Sync Interval many pDelay messages are “wasted”
 - Setting pDelay to a multiple of Sync Interval and synchronising to send pDelay just before Sync would be more efficient
- Anything that can be done to compensate for NRR_{error} will flow through to RR_{error}

Modelling Rate Ratio Error

- Model two aspects...
 - Error from $mNRR_{error}$
 - Errors out along the chain, node-to-node-to-node.
 - $RR_{errorNRR}$ never deviates by more than $mNRR_{errorRange}$
 - See earlier for equation
 - Error from drift between time of mNRR measurement and Time Sync
 - Should average out eventually...but with potentially large variances as the errors won't cancel out node-to-node-to-node.
 - $RR_{errorDrift}$ may vary by a lot more than $mNRR_{errorRange}$
 - Need more work to model and quantify this.

Formulae – RR_{error}

$$RR_{error} = RR_{errorNRR} + RR_{errorDrift}$$

$$RR_{errorNRR} = \text{driftRateError}_{correction} \times \frac{p\text{DelayInterval}}{2} (\text{clockDrift}_{GM} - \text{clockDrift})$$

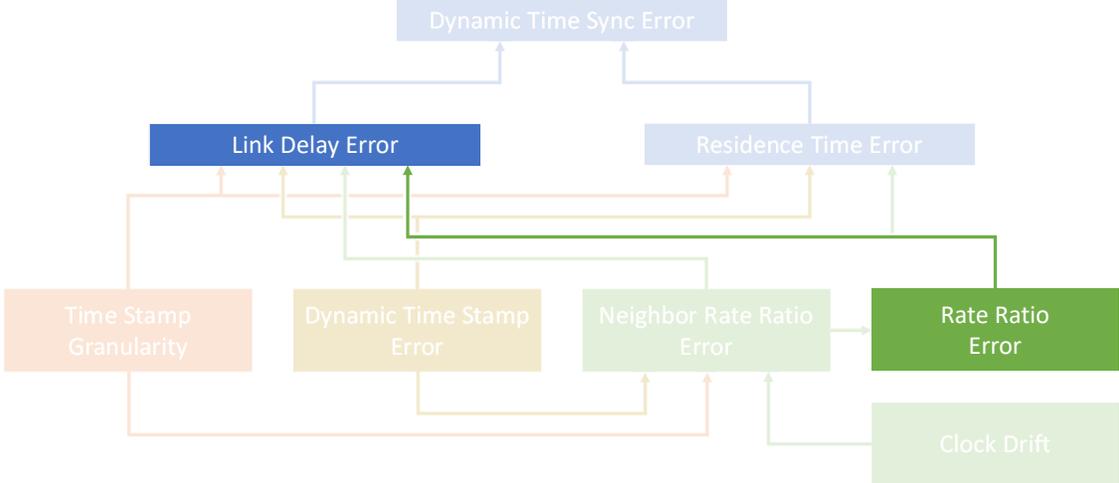
$$RR_{errorDrift}(n) = RR_{errorDrift}(n-1) + \text{driftRateError}_{correction} \times \text{delay}_{mNRR_Sync} (\text{clockDrift}_{GM} - \text{clockDrift}(n))$$

$$\text{delay}_{mNRR_Sync} \sim U(\min = 0, \max = p\text{DelayInterval}(1 - p\text{DelayRespSync}_{correction}))$$

delay_{mNRR_Sync} is modelled as a uniform continuous distribution with minimum = 0 and a maximum = $p\text{DelayInterval}(1 - p\text{DelayRespSync}_{correction})$.

pDelay Error – Rate Ratio Error

Time Sync – Elements & Relationships



$$pDelayError_{RR} \text{ (ns)} = pDelay * RR_{Error}$$

WORST CASE:

$$RR_{Error} \text{ (ppm)} = \pm 10$$

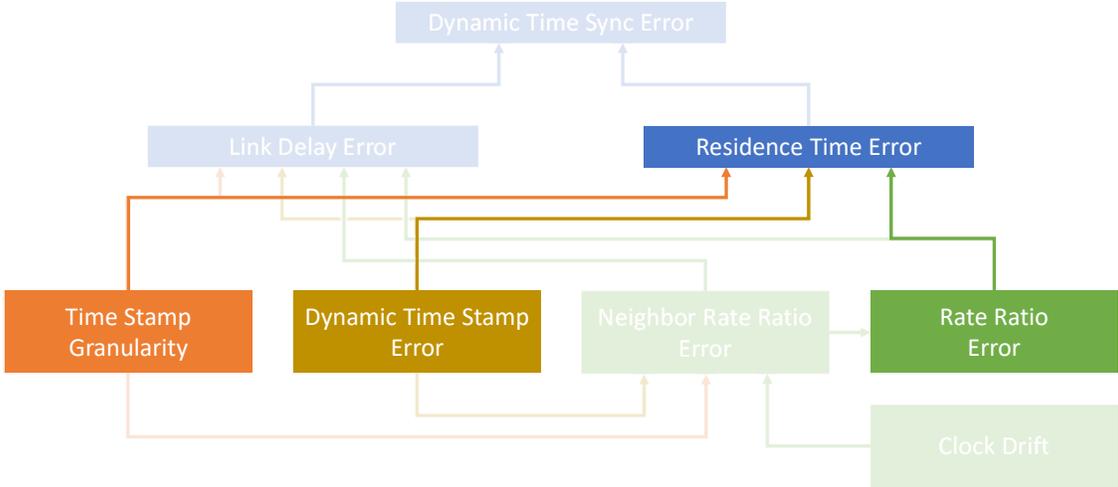
$$pDelayError_{RR} \text{ (ns)} = 150\text{ns} * (\pm 10 \text{ ppm})$$
$$= 0.0015\text{ns (i.e. 1.5ps)}$$

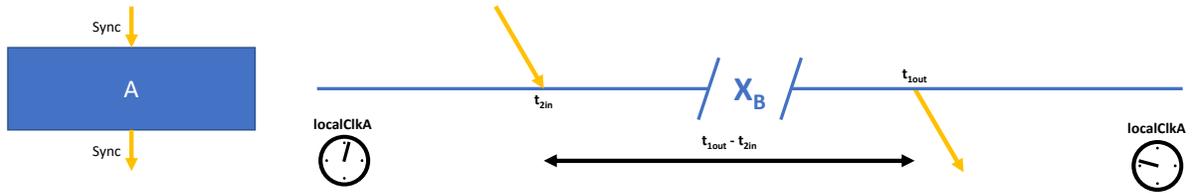
150ns pDelay → 45m

Conclusion: Do Not Model $pDelayError_{RR}$

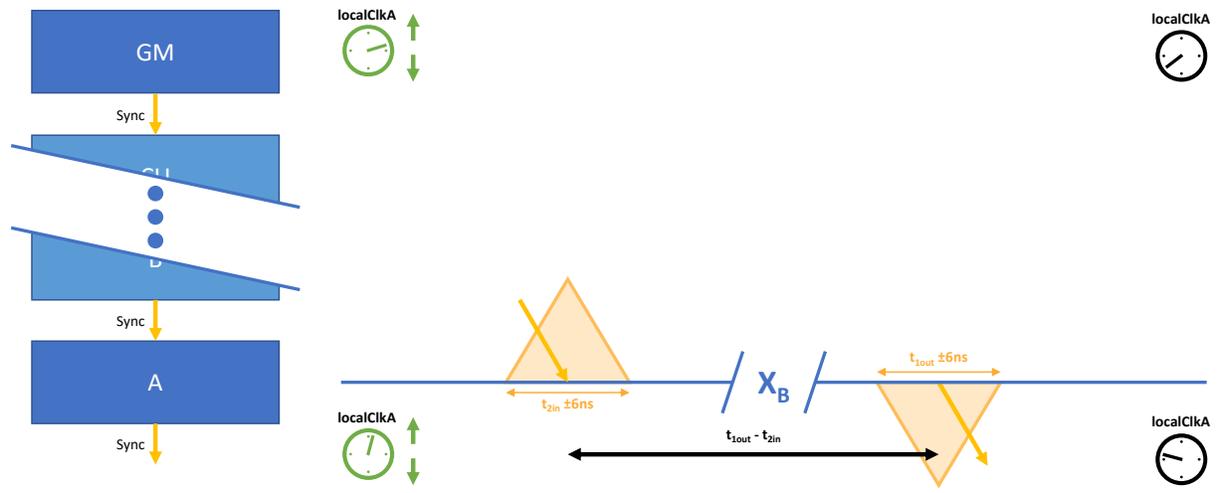
Residence Time Error – Timestamp Errors & Rate Ratio Error

Time Sync – Elements & Relationships





$$\text{residenceTime} = \text{RateRatio} (t_{1\text{out}} - t_{2\text{in}})$$



$$\text{residenceTime} = \text{RateRatio} (t_{1out} - t_{2in})$$

$$\text{residenceTime}_{\text{error}} = t_{2in\text{error}} + t_{1out\text{error}} + \text{RR}_{\text{error}} (t_{1out} - t_{2in})$$

Conclusion: Model residenceTime_{error2in} and residenceTime_{error2out}

$$\text{residenceTimeError}_{RR} \text{ (ns)} = (t_{1\text{out}} - t_{2\text{in}}) * RR_{\text{Error}}$$

WORST CASE:

$$RR_{\text{Error}} \text{ (ppm)} = \pm 10$$

$$\text{residenceTimeError}_{RR} \text{ (ns)} = (t_{1\text{out}} - t_{2\text{in}}) * RR_{\text{Error}}$$

$$= (10\text{ms}) * 10\text{ppm}$$

$$= 100\text{ns}$$

150ns pDelay → 45m

Conclusion: Model residenceTimeError_{RR}

Formulae – residenceTime_{error}

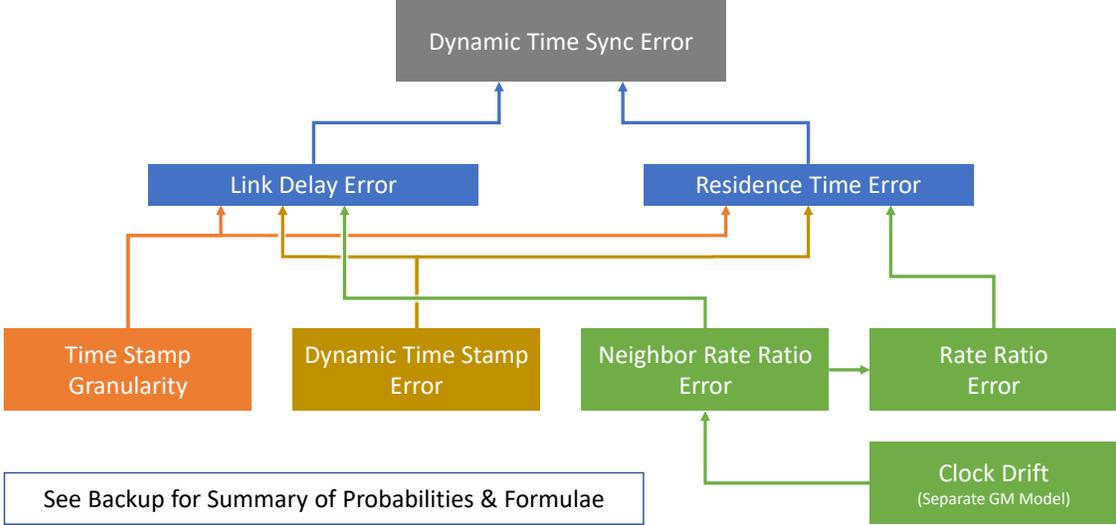
$$residenceTime_{error} = residenceTime_{errorTS} + residenceTime_{errorRR}$$

$$residenceTime_{errorTS} = t_{1Error} - t_{2Error}$$

$$t_{2Error} = TSGE_{RX} + DTSE_{RX} \quad t_{1Error} = TSGE_{TX} + DTSE_{TX}$$

$$residenceTime_{errorRR} = RRe_{error}(residenceTime)$$

Time Sync – Errors to Model



Summary of Main Insights

- meanLinkDelay should be an average of pDelay measurements.
 - Should nearly eliminate meanLinkDelay_{error}
- NRR should not be averaged.
 - Provided (Timestamp Errors / pDelay Interval) is negligible in terms of ppm
- There are significant benefits to the GM having greater ppm stability (lower ppm/s).
- It is hard (impossible?) to eliminate Timestamp granularity and Dynamic Timestamp error effects from Residence Time
 - Need more work to see how far these elements can realistically be reduced
- Synchronising pDelayReq to arrive just before Sync would minimise the effect of Clock Drift.
- There may be opportunities to compensate for NRR_{error}
 - And therefore RR_{error}

Recommended Next Steps

- **Complete the Model**
- **Validate the Model using existing simulation(s) from Geoff**
- **Report back with a specific proposal containing 802.1AS parameters (e.g. rates etc.) and other recommendations that would meet the time accuracy goals for JP 60802 review**

Backup

Input Errors

Error	Distribution	Default Min	Default Max	Equation	Unit
<i>clockDrift_{GM}</i> <small>(<i>clockDrift_{GMmin}</i> & <i>clockDrift_{GMmax}</i>)</small>	X% Stable (zero) (100-X)% Uniform	-0.6	+0.6	$clockDrift_{GM} \sim U(clockDrift_{GMmin}, clockDrift_{GMmax})$	ppm/s
<i>clockDrift</i>	X% Stable (zero) (100-X)% Uniform	-0.6	+0.6	$clockDrift \sim U(-clockDrift, +clockDrift)$	ppm/s
<i>TSGE_{TX}</i>	Uniform	-4	+4	$TSGE_{TX} \sim U(-TSGETX, +TSGETX)$	ns
<i>TSGE_{RX}</i>	Uniform	-4	+4	$TSGE_{RX} \sim U(-TSGERX, +TSGERX)$	ns
<i>DTSE_{TX}</i>	Uniform	-2	+2	$DTSE_{TX} \sim U(-DTSETX, +DTSETX)$	ns
<i>DTSE_{RX}</i>	Uniform	-1	+1	$DTSE_{RX} \sim U(-DTSERX, +DTSERX)$	ns

**Formulae below take into account the unit of the input value.
Formulae in the main presentation do not.**

Other Variable Inputs

Error	Default Value	Unit
<i>pDelay</i>	150	ns
<i>pDelayInterval</i>	1,000	ms
<i>pDelayResponse</i>	10	ms
<i>residenceTime</i>	10	ms

**Formulae below take into account the unit of the input value.
Formulae in the main presentation do not.**

Correction Factor Inputs

Error	Default Value	Unit
<i>meanLinkDelay_{errorTS}correction</i>	0	%
<i>driftRateError_{correction}</i>	0	%
<i>pDelayRespSync_{correction}</i>	0	%

Formulae – Top Level

$$DTE(x) = \sum_{n=1}^x (\text{meanLinkDelay}_{\text{error}}(n) + \text{residenceTime}_{\text{error}}(n)) \quad \text{ns}$$

$$\text{meanLinkDelay}_{\text{error}} = (1 - \text{meanLinkDelay}_{\text{errorCorrection}})(p\text{Delay}_{\text{errorTS}} + p\text{Delay}_{\text{errorNRR}} + p\text{Delay}_{\text{errorRR}}) \quad \text{ns}$$

$$\text{residenceTime}_{\text{error}} = \text{residenceTime}_{\text{errorTS}} + \text{residenceTime}_{\text{errorRR}} \quad \text{ns}$$

**The initial model will not include factors shown in grey.
x is number of hops.**

Formulae – meanLinkDelay_{error}

$$\text{meanLinkDelay}_{\text{error}} = (1 - \text{meanLinkDelay}_{\text{errorCorrection}})(p\text{Delay}_{\text{errorTS}} + p\text{Delay}_{\text{errorNRR}} + p\text{Delay}_{\text{errorRR}}) \quad \text{ns}$$

$$p\text{Delay}_{\text{errorTS}} = \left(\frac{(t_{4PD\text{Error}} - t_{1PD\text{Error}}) - (t_{3PD\text{Error}} - t_{2PD\text{Error}})}{2} \right) \quad \text{ns}$$

$$t_{1PD\text{Error}} = \text{TSGE}_{\text{TX}} + \text{DTSE}_{\text{TX}} \quad t_{2PD\text{Error}} = \text{TSGE}_{\text{RX}} + \text{DTSE}_{\text{RX}} \quad \text{ns}$$

$$t_{3PD\text{Error}} = \text{TSGE}_{\text{TX}} + \text{DTSE}_{\text{TX}} \quad t_{4PD\text{Error}} = \text{TSGE}_{\text{RX}} + \text{DTSE}_{\text{RX}} \quad \text{ns}$$

$$p\text{Delay}_{\text{errorNRR}} = \frac{m\text{NRR}_{\text{error}}}{10^6} \left(\frac{p\text{DelayResponse} \times 10^6}{2} + p\text{Delay}_{\text{errorTS}} \right) \quad \text{ns}$$

$$= m\text{NRR}_{\text{error}} \left(\frac{p\text{DelayResponse}}{2} + \frac{p\text{Delay}_{\text{errorTS}}}{10^6} \right) \quad \text{ns}$$

$$p\text{Delay}_{\text{errorRR}} = \frac{RR_{\text{error}}}{10^6} (p\text{Delay} + (1 - p\text{Delay}_{\text{errorCorrection}})(p\text{Delay}_{\text{errorTS}} + p\text{Delay}_{\text{errorNRR}})) \quad \text{ns}$$

Formulae – residenceTime_{error}

$$residenceTime_{error} = residenceTime_{errorTS} + residenceTime_{errorRR} \quad \text{ns}$$

$$residenceTime_{errorTS} = t_{1Error} - t_{2Error} \quad \text{ns}$$

$$t_{2Error} = TSGE_{RX} + DTSE_{RX} \quad t_{1Error} = TSGE_{TX} + DTSE_{TX} \quad \text{ns}$$

$$residenceTime_{errorRR} = \frac{RR_{error}}{10^6} (residenceTime \times 10^6 + residenceTime_{errorTS}) \quad \text{ns}$$

$$= RRe_{error} \left(residenceTime + \frac{residenceTime_{errorTS}}{10^6} \right) \quad \text{ns}$$

Formulae –NRR_{error}

$$mNRR_{error}(n) = \text{driftRateError}_{correction} \times \left(\frac{pDelayInterval}{2 \times 10^3} + \frac{NRRInterval_{error}}{10^9} \right) (\text{clockDrift}(n-1) - \text{clockDrift}(n)) \quad \text{ppm}$$

$$NRRInterval_{error} = (t_{4PDerror} - t_{4previousPDerror}) - (t_{3PDerror} - t_{3previousPDerror}) \quad \text{ns}$$

$$t_{3previousPDerror} = TSGE_{TX} + DTSE_{TX} \quad t_{4previousPDerror} = TSGE_{RX} + DTSE_{RX} \quad \text{ns}$$

The factors $t_{3PDerror}$ and $t_{4PDerror}$ are the same as in the $pDelay_{error}$ calculation.
(Same value. No need for new random numbers, as the same $pDelayResp$ message is used to measure NRR.)

Formulae – RR_{error}

$$RR_{error} = RR_{errorNRR} + RR_{errorDrift} \quad \text{ppm}$$

$$RR_{errorNRR} = \text{driftRateError}_{correction} \times \left(\frac{pDelayInterval}{2 \times 10^3} + \frac{NRRinterval_{error}}{10^9} \right) (\text{clockDrift}_{GM} - \text{clockDrift}) \quad \text{ppm}$$

The factor $NRRinterval_{error}$ is the same as in the NRR_{error} calculation as it is the same error.

$$RR_{errorDrift}(n) = RR_{errorDrift}(n-1) + \text{driftRateError}_{correction} \times \frac{delay_{mNRR_Sync}}{10^3} (\text{clockDrift}_{GM} - \text{clockDrift}(n)) \quad \text{ppm}$$

$$delay_{mNRR_Sync} \sim U(0, (1 - pDelayRespSync_{correction})pDelayInterval) \quad \text{ms}$$

Does not include any effect of changing clock drift (ppm/s²) during Residence Time. See speaker notes in RR_{error} section for details.

Special Cases

$$\mathit{clockDrift}(0) = \mathit{clockDrift}_{GM}$$

For the first hop ($n = 1$), $n - 1 = 0$, i.e. the first device in chain, which is the GM.

$$\mathit{residenceTime}_{\mathit{error}}(x) = 0$$

For the last hop ($n = x$), the Sync message is not passed on so there is no Residence Time Error.

Of Interest...

- What happens when there is GM Clock Drift?
 - The ability to control clockDrift_{GM} separately allows this to be investigated.
- How much do the various factors influence DTE?
 - Also: how do they accumulate along the chain?
 - The separation of various factors until late in the process allows detailed analysis
- How effective are the various correction fields?

Careful construction of the model will help when it comes to extracting the values of interest (minimums, maximums, means, etc...) as the model is run, without taking excessive time or using excessive memory.