Transmitter Pre-Compensation for 320 MB/sec SCSI

Andrew Bishop
Quantum Corporation

Objectives

- Look at two forms of transmitter pre-compensation that might be used for 320 MB/sec SCSI
  - Mathematically model timing pre-compensation
    - How does timing pre-compensation help:
      - Isolated pulses
      - High frequency patterns (0, 1, 0, 1, etc.)
  - Use experimental data to examine amplitude pre-compensation using real cables with real loads.
    - How much does amplitude pre-comp improve signal integrity?
    - How much amplitude pre-comp is required for:
      - short cables?
      - long cables?
    - Determine the optimum value for amplitude pre-comp

Transmitter Pre-Compensation for 320 MB/sec SCSI

Part I

Timing Pre-Compensation Study

Disclaimer

- All data collected in this section of the presentation is based on a simple *optimistic* model for SCSI signals.
- Real factors such as DC attenuation, reflections, offsets, frequency dependent skew, differential skew, etc. will degrade performance even more than suggested.
- This simple model is used to facilitate the analysis and to gain some insight into what can be expected in actual practice.

Timing Pre-Compensation Basics

- Adjust the timing of pulse edges to improve receiver amplitude noise margin and set-up time.
- If the pulse is sent earlier, then response will rise further past the threshold.
- If the pulse is earlier, there is more set-up time, and less hold time for the previous bit.
- Timing pre-comp has largest effect for an isolated pulse.
- Timing pre-comp has ***no effect*** for high-frequency patterns (0, 1, 0, 1, etc.).

Effect on Received Amplitude

Simplified Transmitter Waveform:

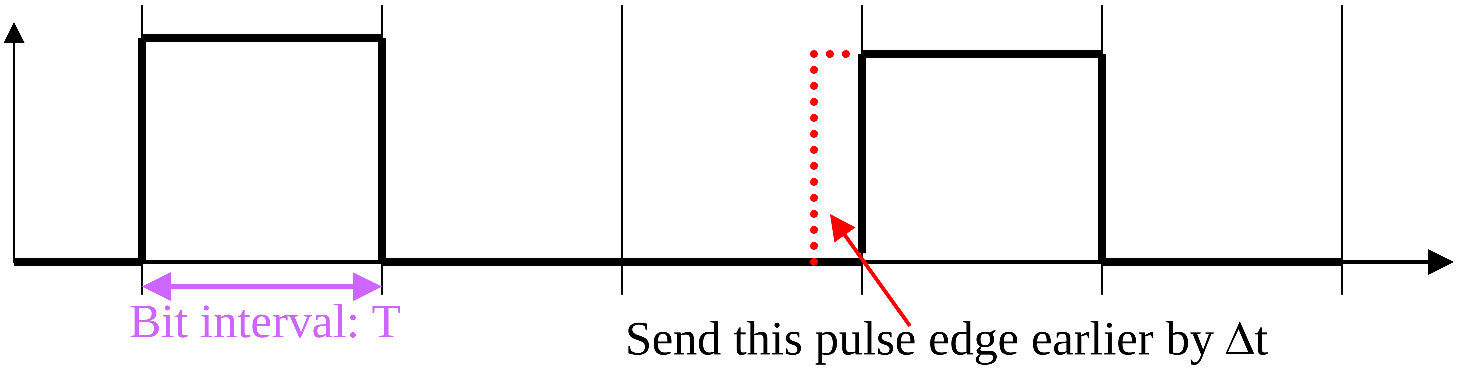

Receiver Waveform:

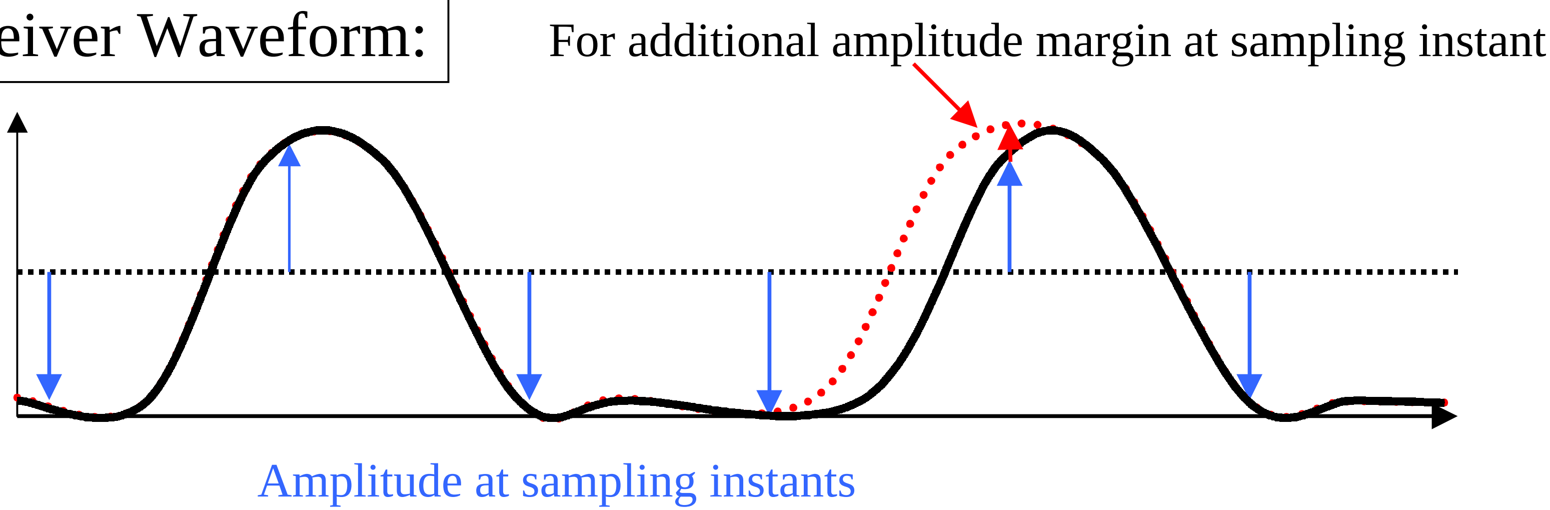

Receiver Waveform Model

- Assuming a first order model for cable the received waveform during bit interval (k+1) can be expressed as:

$$v_{k+1}(t) + 500mV = \begin{cases} X_{k+1} = 0 & V_k \exp(-t/\tau) \\ X_{k+1} = 1 & 1000mV - (1000mV - V_k)\exp(-t/\tau) \end{cases}$$

- Where:
  - X_{k+1}: is the bit sent during interval $(k+1)$
  - V_k: is the amplitude at the end of bit interval (k)
  - τ: is the dominant first-order time constant of the cable
  - $v_{k+1}(t)$: is receiver waveform during bit interval $(k+1)$

- This model assumes a zero-delay, just a first-order amplitude roll-off with frequency and normalized peak-to-peak amplitude of 1000mV

Receiver Waveform Model

$$v_{k+1}(t)+500mV=\begin{cases} X_{k+1}=0 & V_k\exp(-t/\tau) \\ X_{k+1}=1 & 1000mV-(1000mV-V_k)\exp(-t/\tau) \end{cases}$$

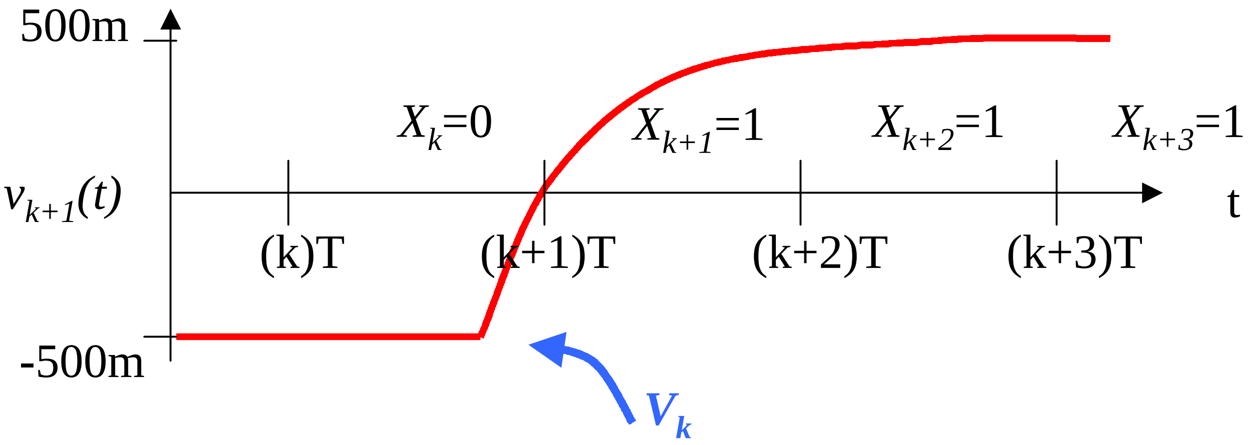

- This assumes several optimistic scenarios:
  - input is an ideal square wave
  - no frequency dependent time skew in cable
  - Output has a voltage swing of 1V (no attenuation in the cable)

Receiver Waveform Model

- First order cable model:
  - attenuation versus frequency
  - characterized by a single pole roll-off

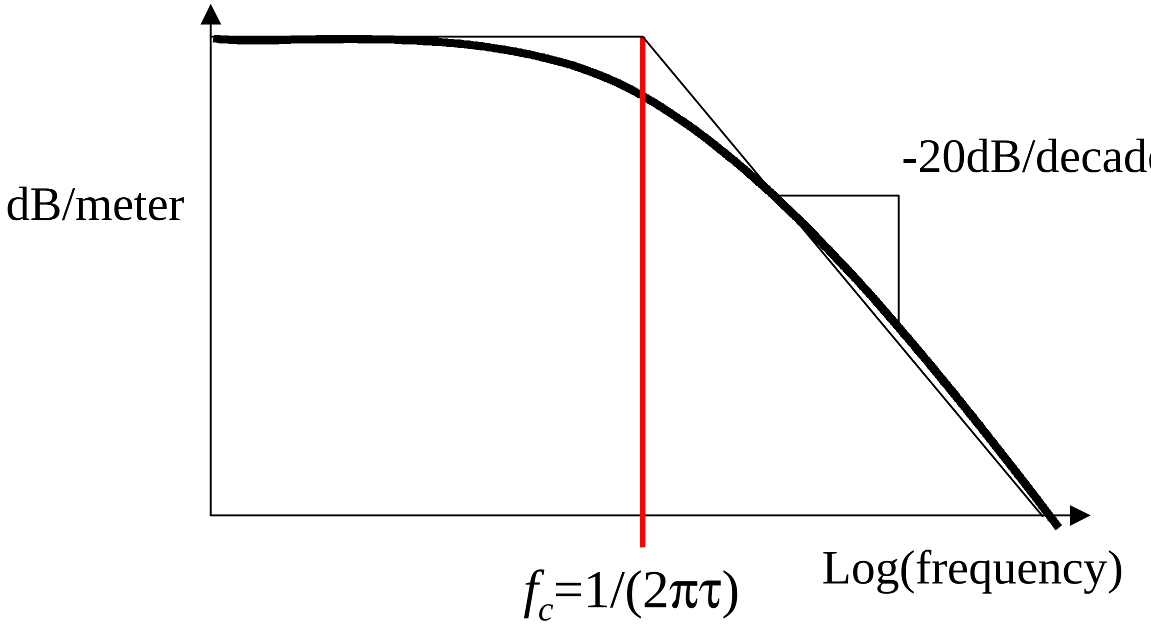

Pre-comp Amplitude Margin

- Receiver input is sample of waveform at the sampling instant (T_{sample})
- Sampling instant defined by position of zero crossings for high frequency pattern (de-skew).

$$v_{k+1}(T_{sample}) + 500mV = \begin{cases} X_{k+1} = 0 & V_k \exp(-T_{sample}/2\tau) \\ X_{k+1} = 1 & 1000mV - (1000mV - V_k)\exp(-T_{sample}/2\tau) \end{cases}$$

- For worst case (isolated bit) we have:

$$X_{k-N}, .. X_k = 0, X_{k+1} = 1, X_{k+2} = 0 \therefore V_k = 0.$$

- Without pre-compensation the amplitude is:

$$v_{k+1}(T_{sample}) = \left(1000mV - 1000mV \exp(-T_{sample}/2\tau)\right) - 500mV$$

- With timing pre-comp shift (Δt) the amplitude is:

$$v_{k+1}(T_{sample}) = \left(1000mV - 1000mV \exp(-(T_{sample} + \Delta t)/\tau)\right) - 500mV$$

Receiver Amplitude

- Numerical calculations for isolated pulse:
  - $\tau \approx 1/(2\pi\ 40\ \text{MHz})$, Δt = 0ps → 2ns by 500ps

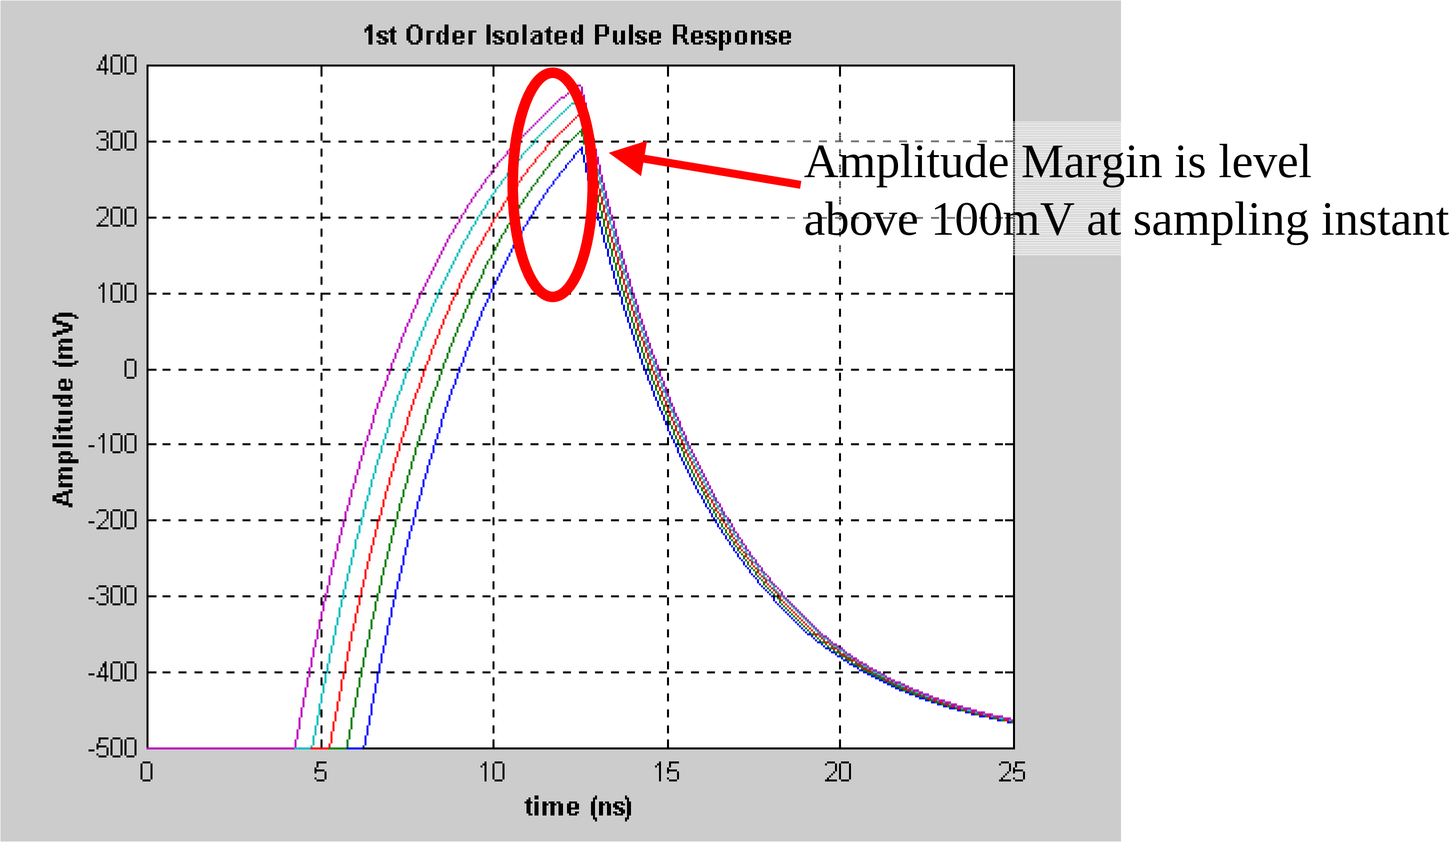

Amplitude Margin

- Amplitude margin *vs.* timing pre-comp shift (Δt) for $f_c = 40\text{MHz}$:

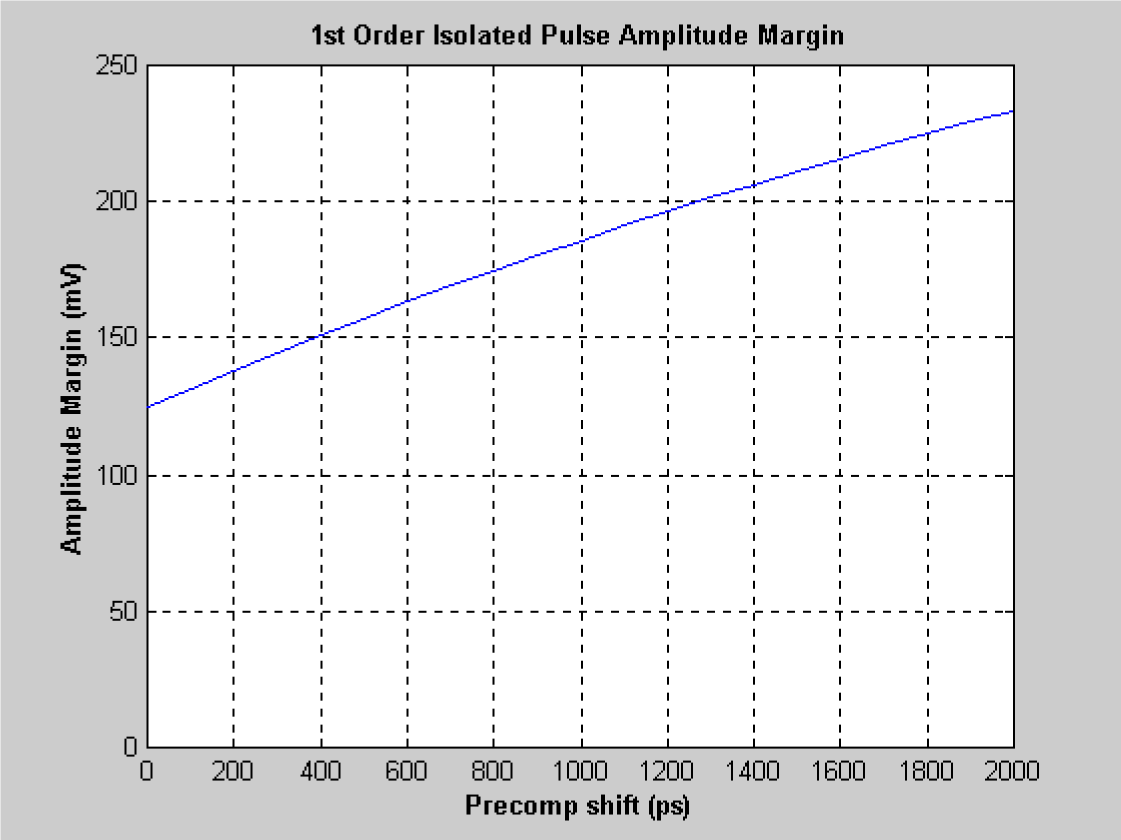

Amplitude Margin Improvement with Pre-Comp

- Improvement in amplitude margin with timing pre-comp shift Δt for f_c = 40 MHz:

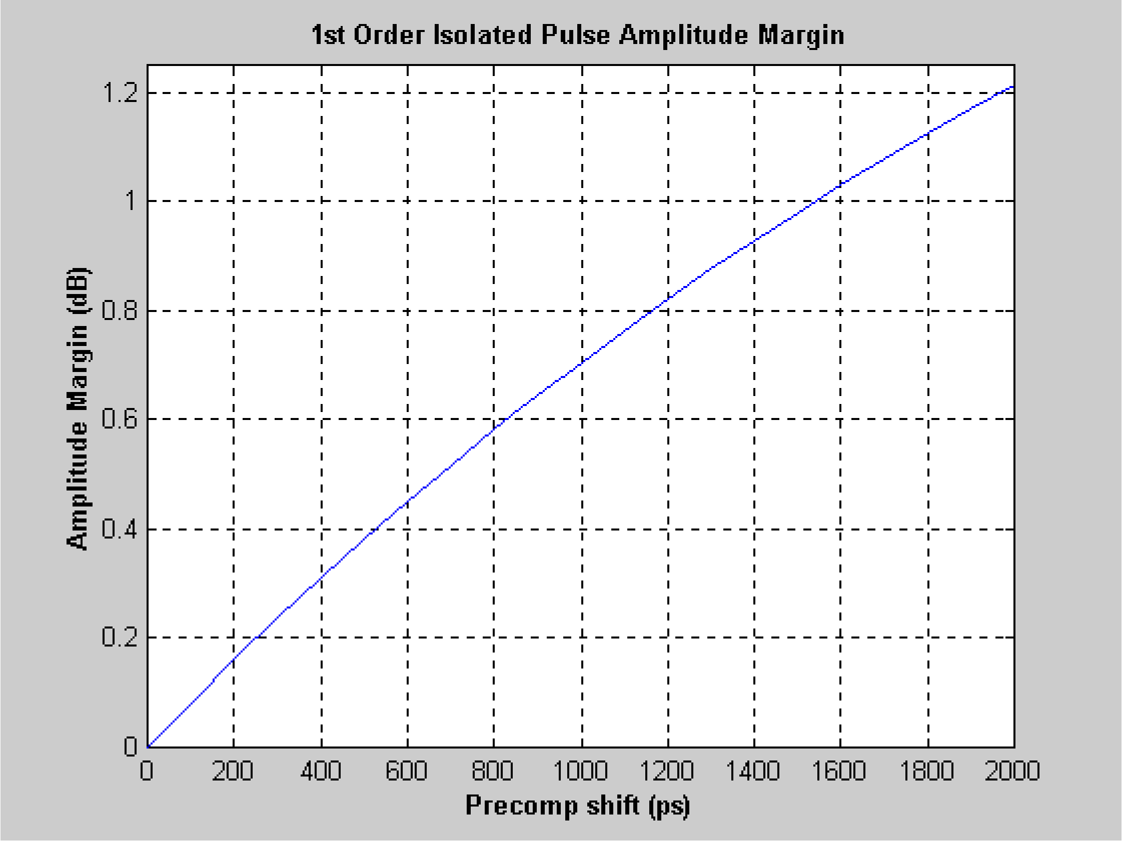

Timing Pre-comp and Set-up/Hold Time

Transmitter Waveform:

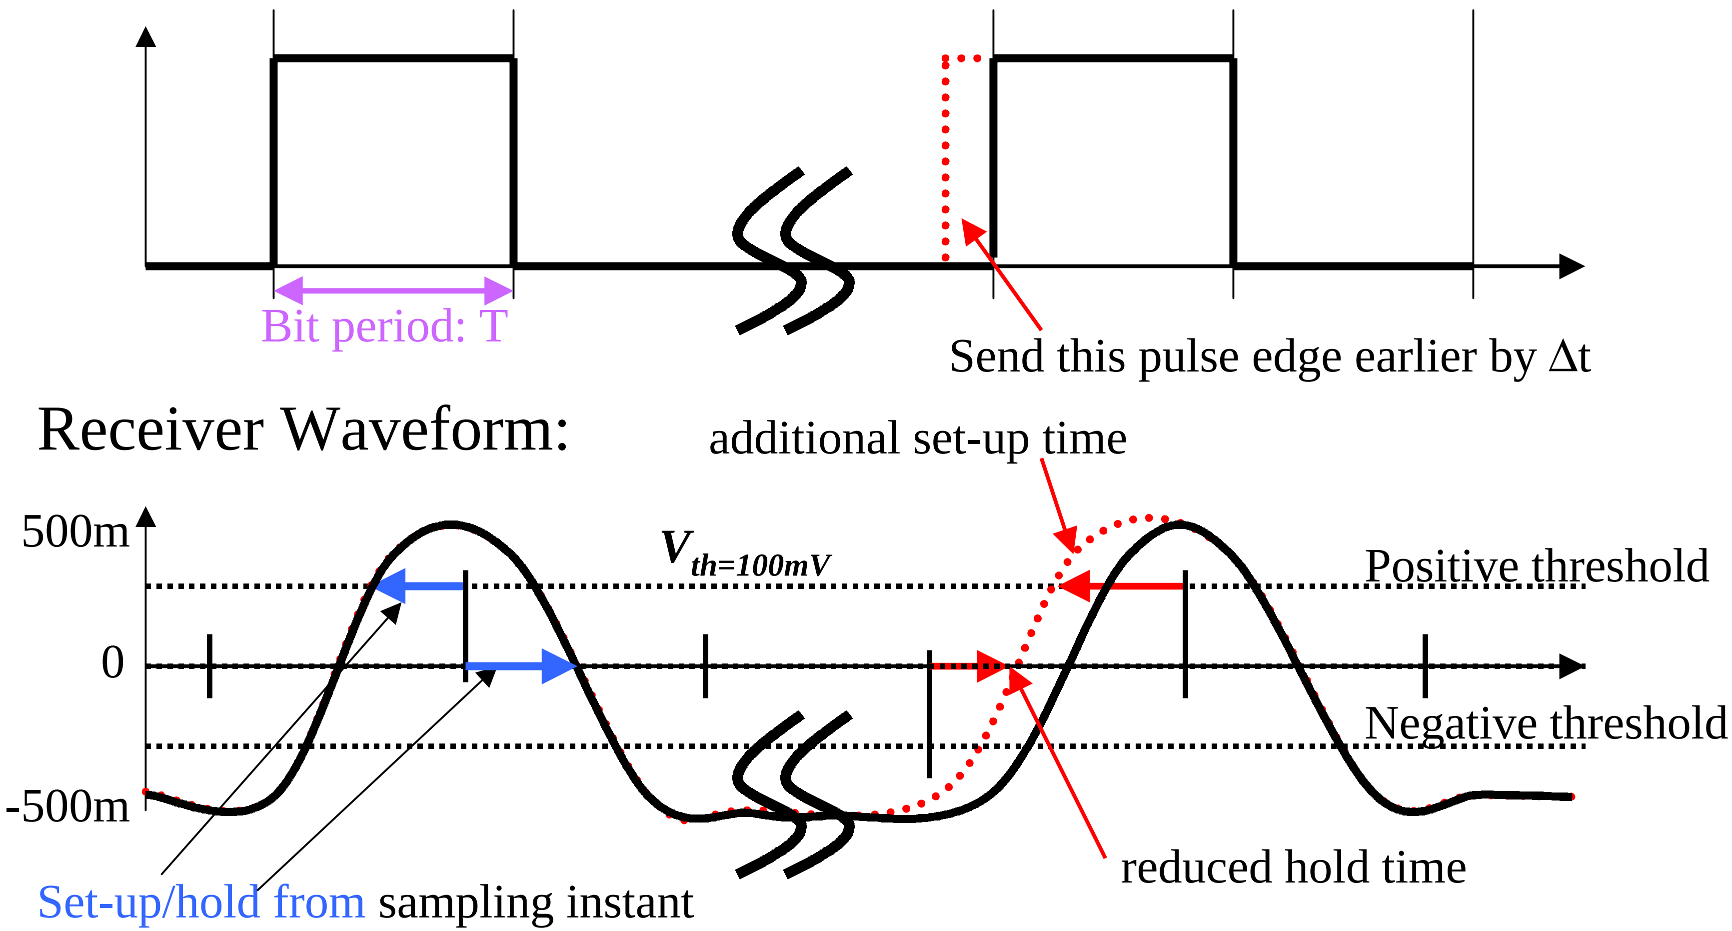

Set-up/Hold Time

- Time before sampling instant (T_{sample}) that waveform passes positive threshold (V_{setup})

$$t_{su}(V_{setup}) = T_{sample} + \Delta t + \tau \ln\left(\frac{500mV - V_{setup}}{1000mV}\right)$$

- Set-up time for current bit linearly increases with pre-comp shift (Δt)
- Hold time for previous bit linearly decreases with pre-comp shift (Δt)

$$t_h(V_{hold}) = T_{sample} - \Delta t - \tau \ln\left(\frac{500mV - V_{hold}}{1000mV}\right)$$

- Isolated bit set-up/hold times calculated assuming that the starting voltage V_k = 0 (ISI voltage)

Set-up/Hold Time

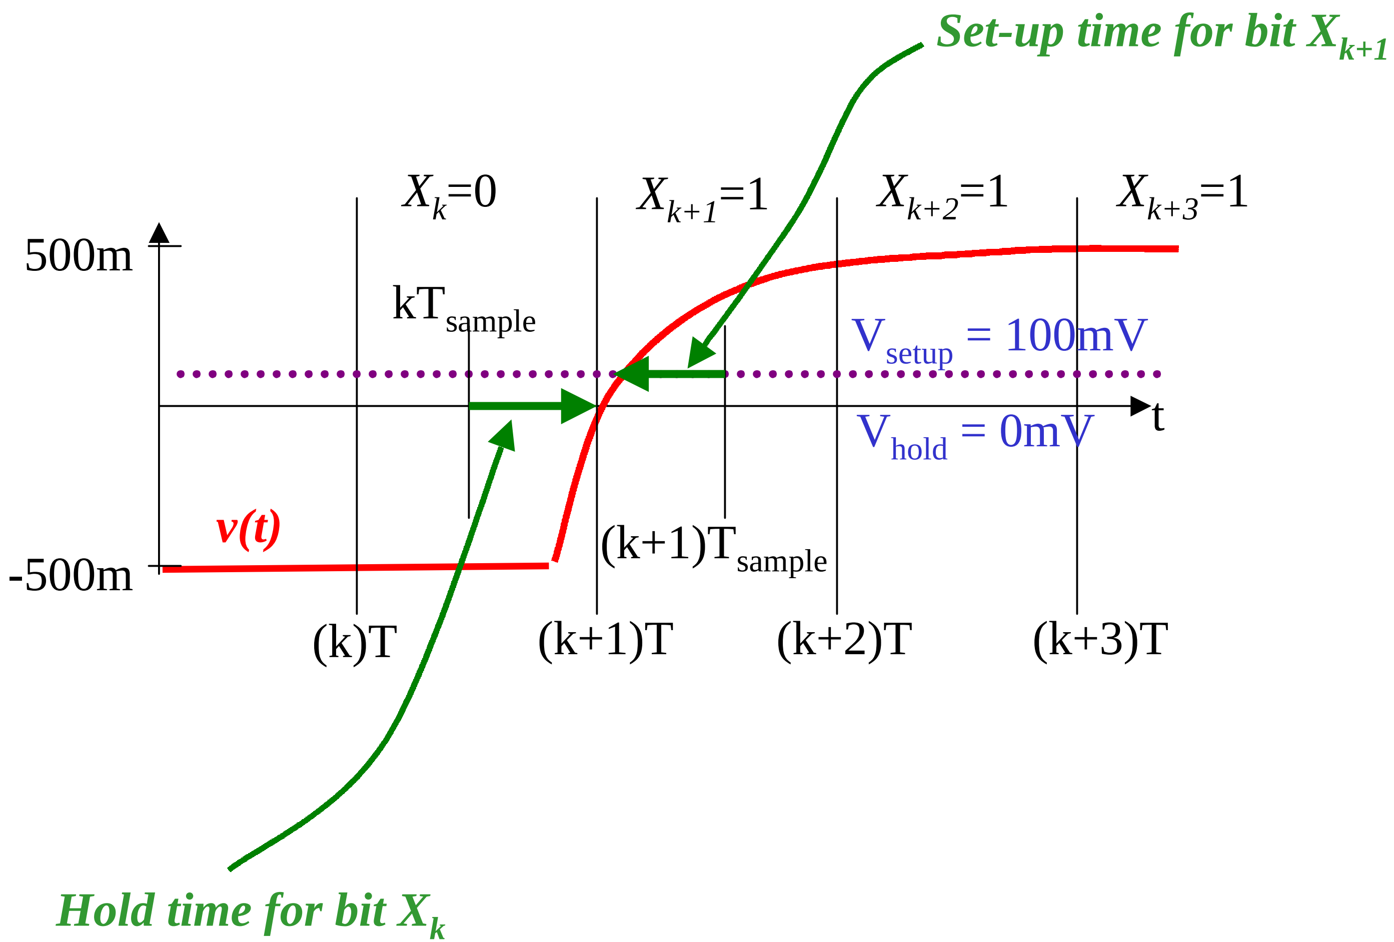

Receiver Set-up/Hold Time

- Numerical calculations for isolated pulse:

  $\tau \approx 1/ (2\ \pi\ 40\ \text{MHz})$, $\Delta\ t$ = *0ps* $\rightarrow$ 2ns by 500ps

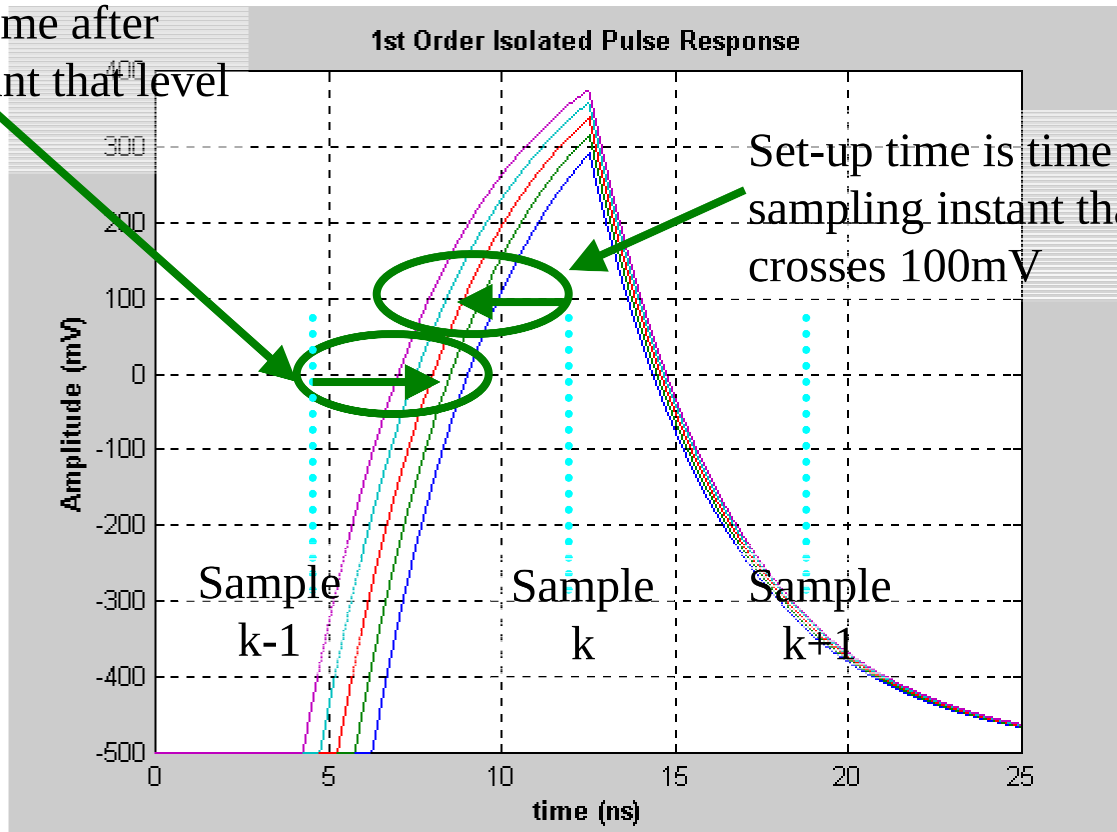

Fast and Slow Case Set-up/Hold Times

- Numerical calculations for isolated pulse:
  - slow case cable: $\tau \approx 1/(2\pi\ 40\ \text{MHz})$, Δt = 0ps → 2ns by 100ps (dashed)
  - fast case cable: $\tau \approx 1/(2\pi\ 120\ \text{MHz})$, Δt = 0ps → 2ns by 100ps (solid)

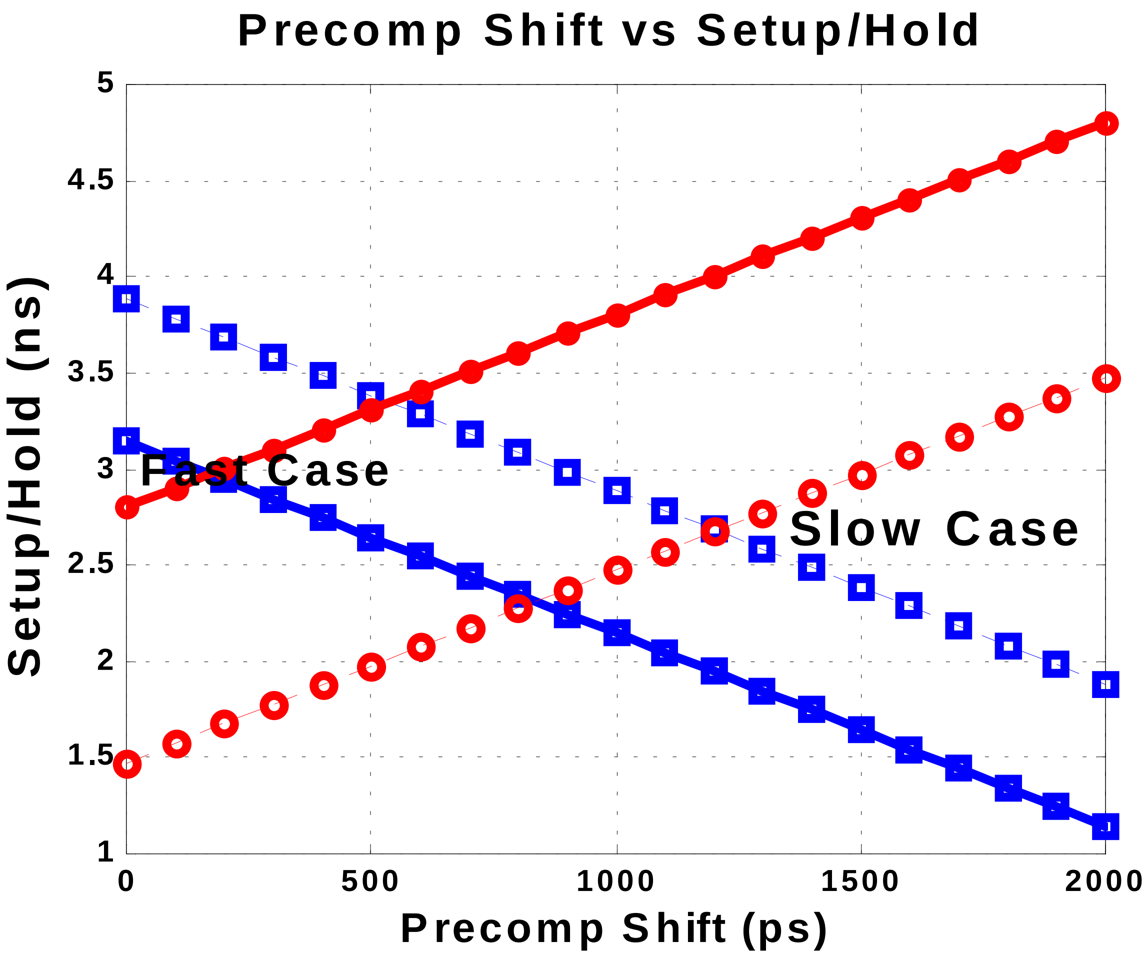

High Frequency Pattern Set-up Time

- For the high-frequency pattern, shifting each edge by Δt is the same as no pre-comp; all edges move over by the same amount (i.e., timing pre-comp just defeats de-skew)
- What is the set-up/hold time for the high frequency pattern.
- The first-order waveform equation in bit interval (k+1):

$$v_{k+1}(t) + 500mV = \begin{cases} X_{k+1} = 0 & V_k \exp(-t/2\tau) \\ X_{k+1} = 1 & 1000mV - (1000mV - V_k)\exp(-t/2\tau) \end{cases}$$

- Where V_k is the amplitude at the start of bit interval (k+1) (i.e., ISI)
- Can calculate V_k for $\{X_k, X_{k-1}, X_{k-2}, \ldots\}$ = {…, 1, 0, 1, 0, …}

High Frequency Pattern ISI Voltage

- The residual ISI voltage (V_k) for a high frequency pattern is the peak voltage of the high frequency waveform.

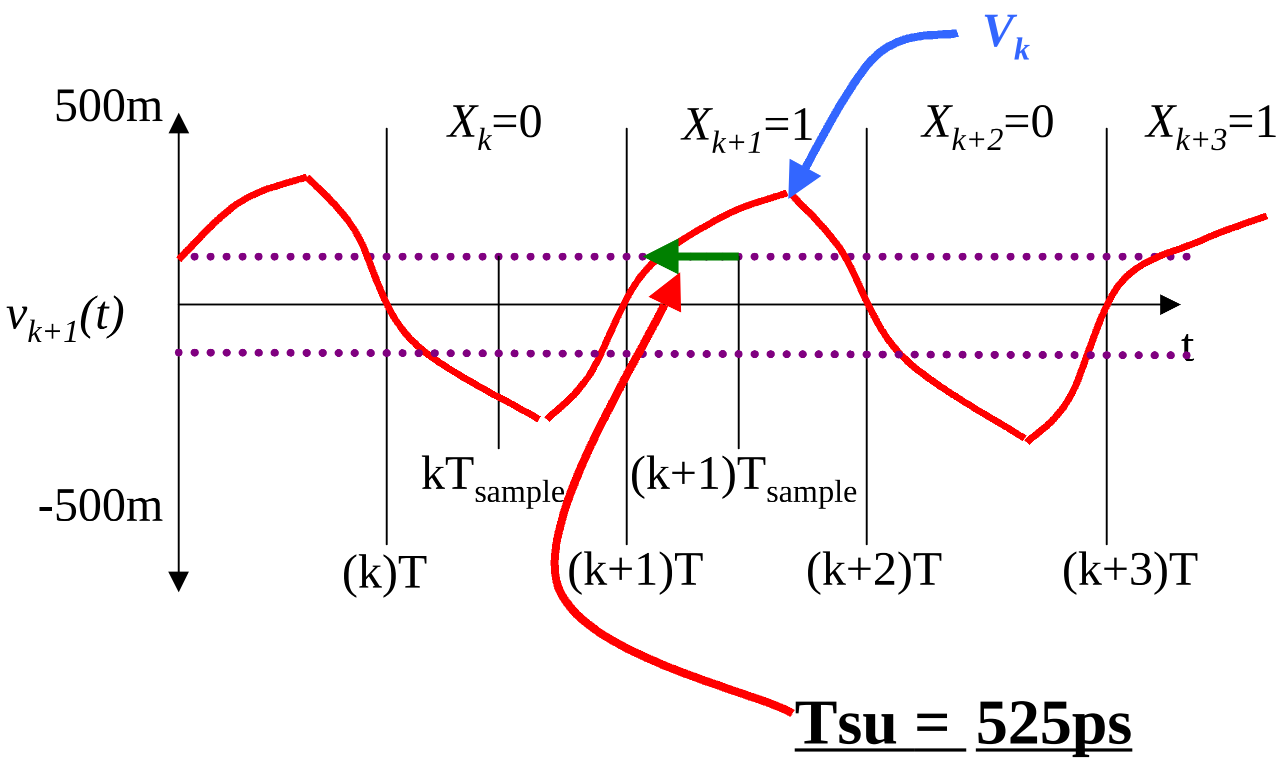

Assuming you have skew compensated perfectly!

High Frequency Pattern ISI Voltage

Quantum™

- The residual voltage (V_k) for a high frequency pattern depends on the cable cut-off frequency and the total receiver amplitude.
- At f_c = 40MHz V_k = 330mV for a $1V_{pk\text{-}pk}$ signal swing

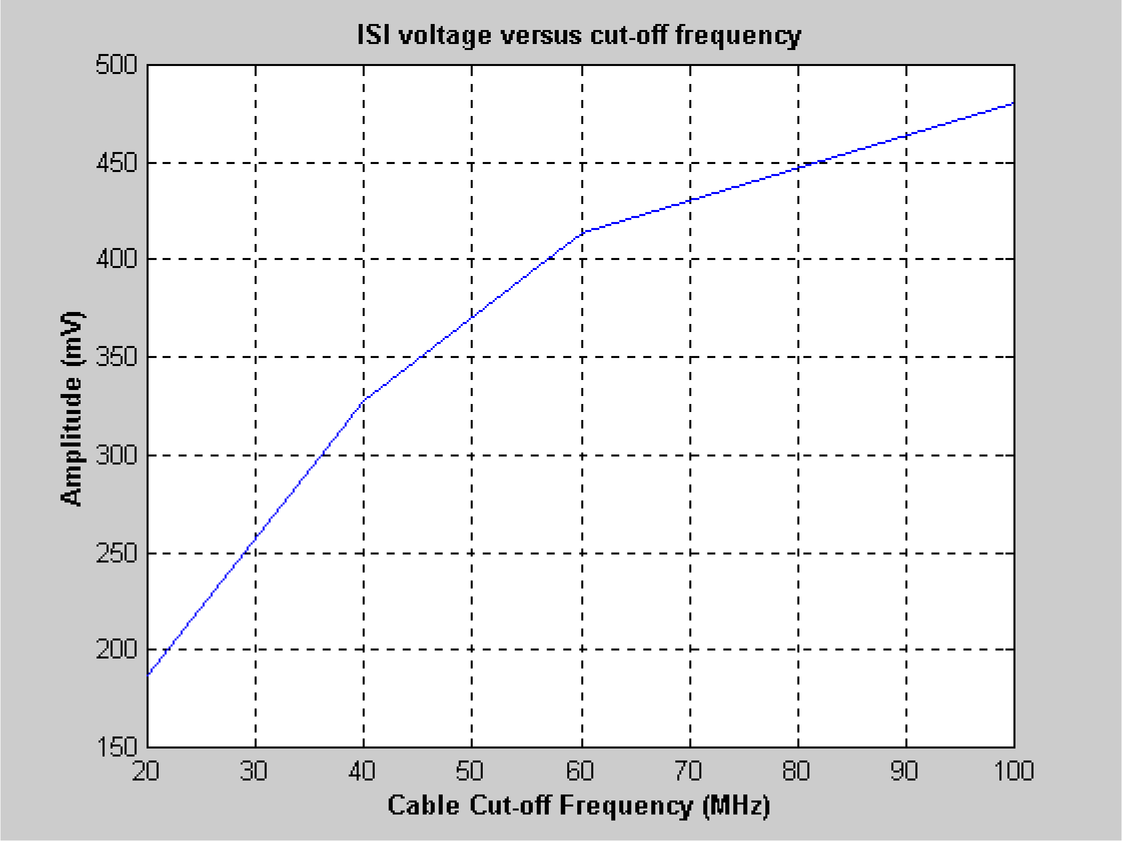

High Frequency Pattern Set-up Time

- With signal starting from ISI residual voltage the set-up time is:

$$t_{su}(V_{setup}) = T_{sample} + \tau \ln\left(\frac{500mV - V_{setup}}{1000mV - V_k}\right)$$

- The threshold voltage (V_{setup}) is specified at 100mV.

High Frequency Pattern Set-up Time

- Set-up time versus cable cut-off frequency.
  - For f_c = 60Mhz (=1/2π τ), set-up time = 1.75ns
  - For f_c = 40Mhz (=1/2π τ), set-up time = 525ps

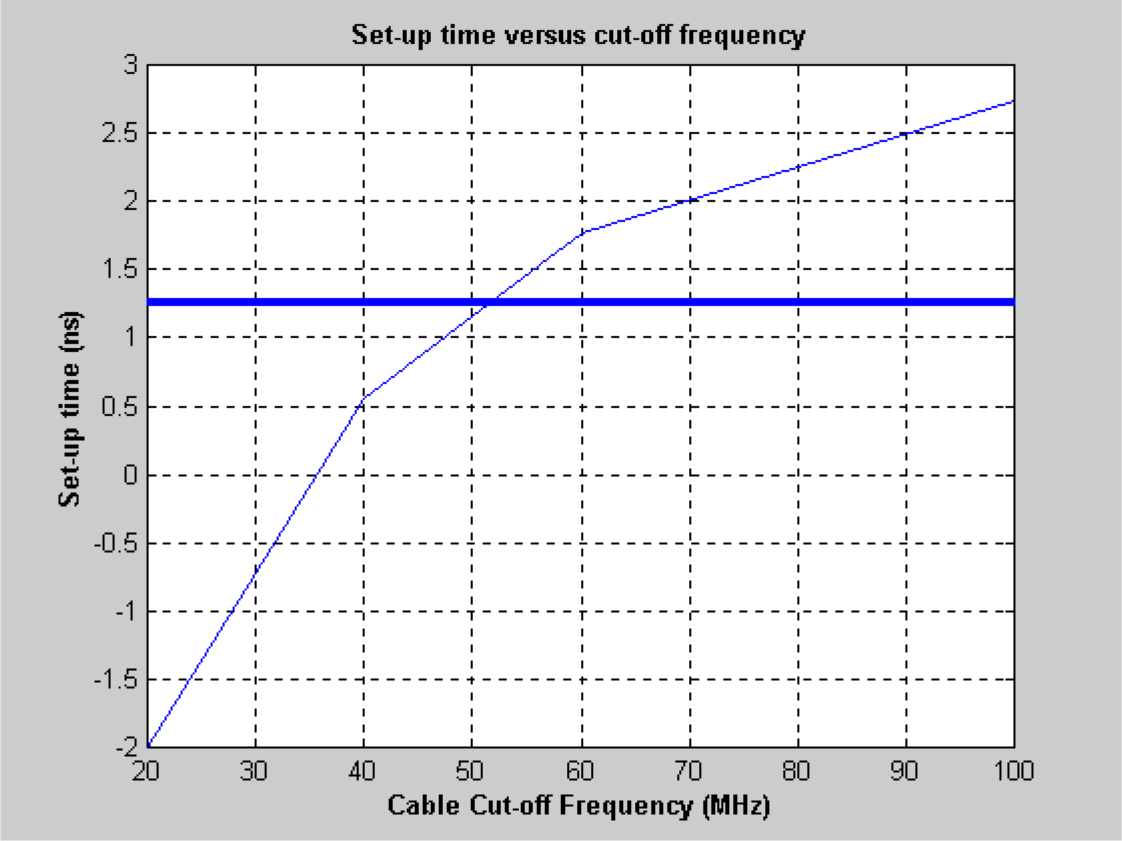

High Frequency Pattern Hold Time

- Hold time versus cable cut-off frequency

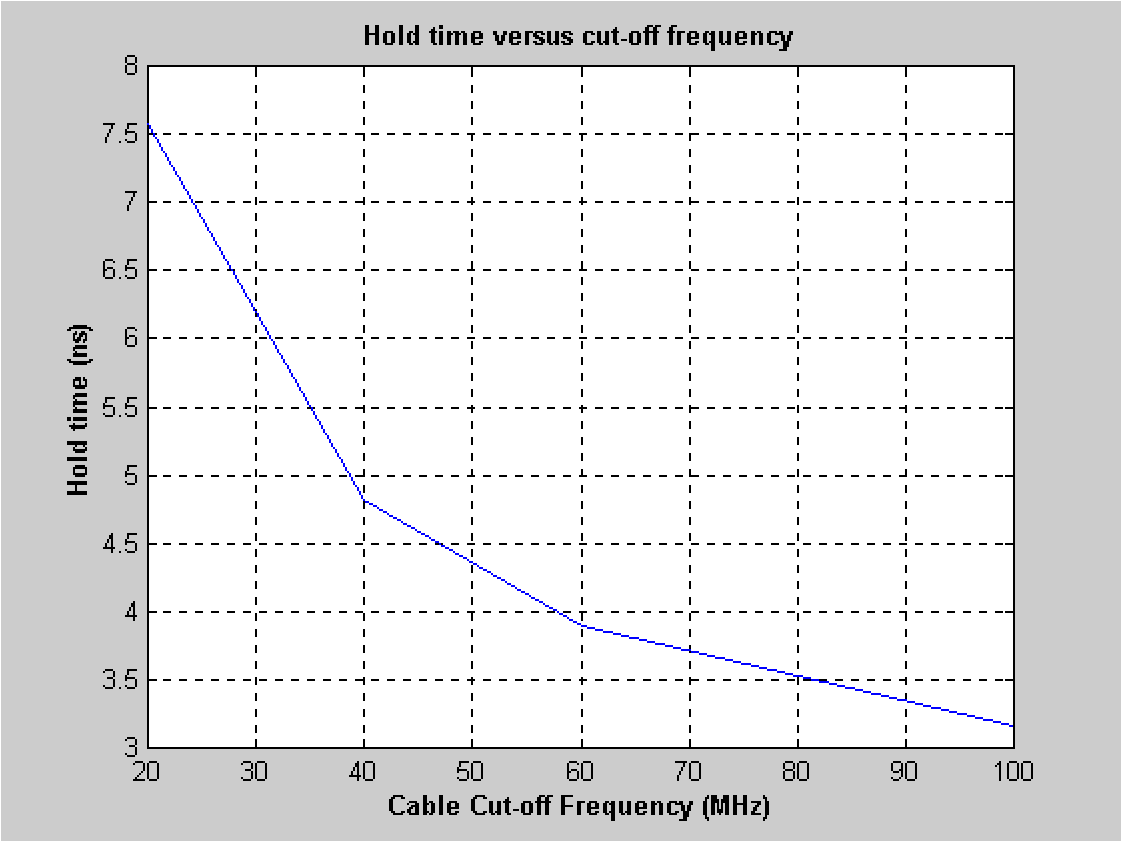

Quantum™

Conclusions on Timing Pre-Comp

- Quantum believes that timing pre-compensation alone is insufficient to compensate for ISI at 320 MB/sec SCSI transfer rates

Transmitter Pre-Compensation for 320 MB/sec SCSI

Part II

Amplitude Pre-Compensation Study

Quantum™

Amplitude Pre-Compensation Basics

- Increase the amplitude of transmitted signal on signal transitions
- If there is no transition, return amplitude to nominal level (normalized to 1)
- ΔA - amplitude increase is design parameter

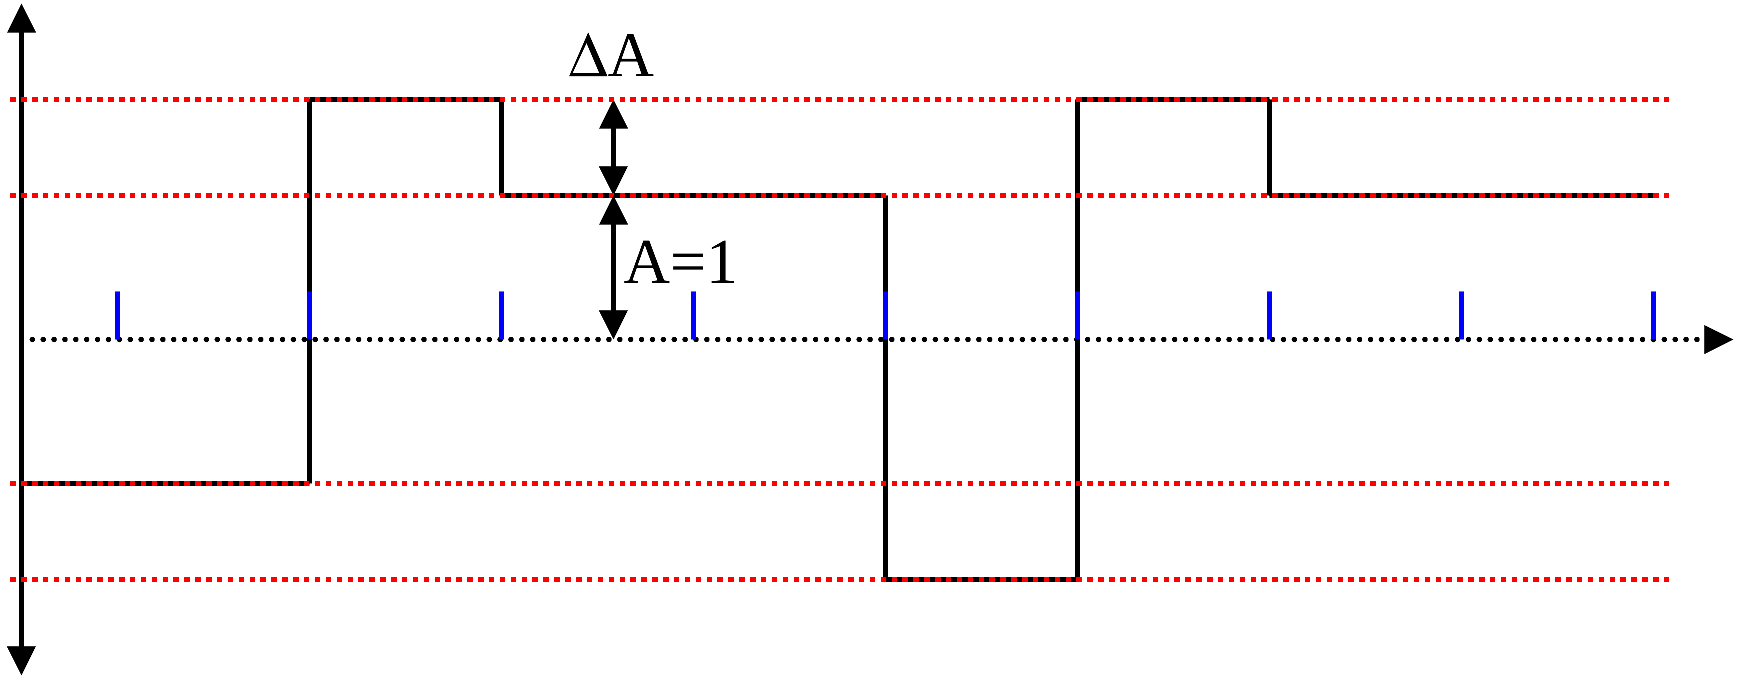

Experimental Setup #1 - Short Cable

- Hitachi 0.625 meter, 32AWG twisted-flat ribbon cable, 12.5cm load spacing, plus 6-slot back-plane.
- Waveforms captured @ 4Gs/s
  - no pre-comp: ΔA = 0.0
  - amplitude pre-comp: ΔA = 0.2, 0.5, 0.8

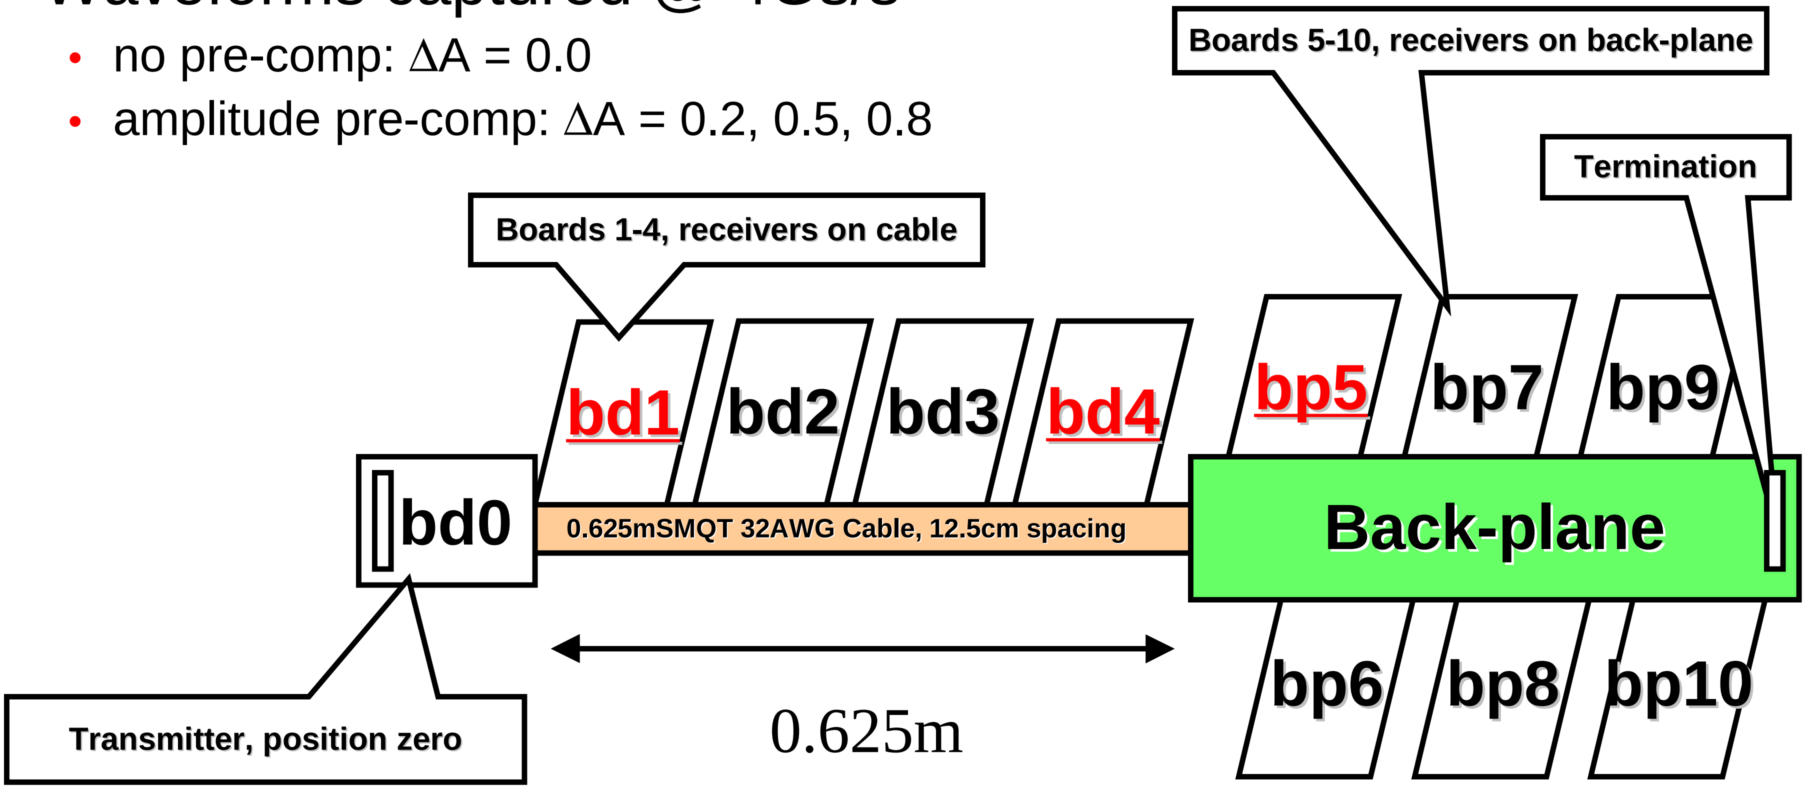

Electrical Setup #1 - Short Cable

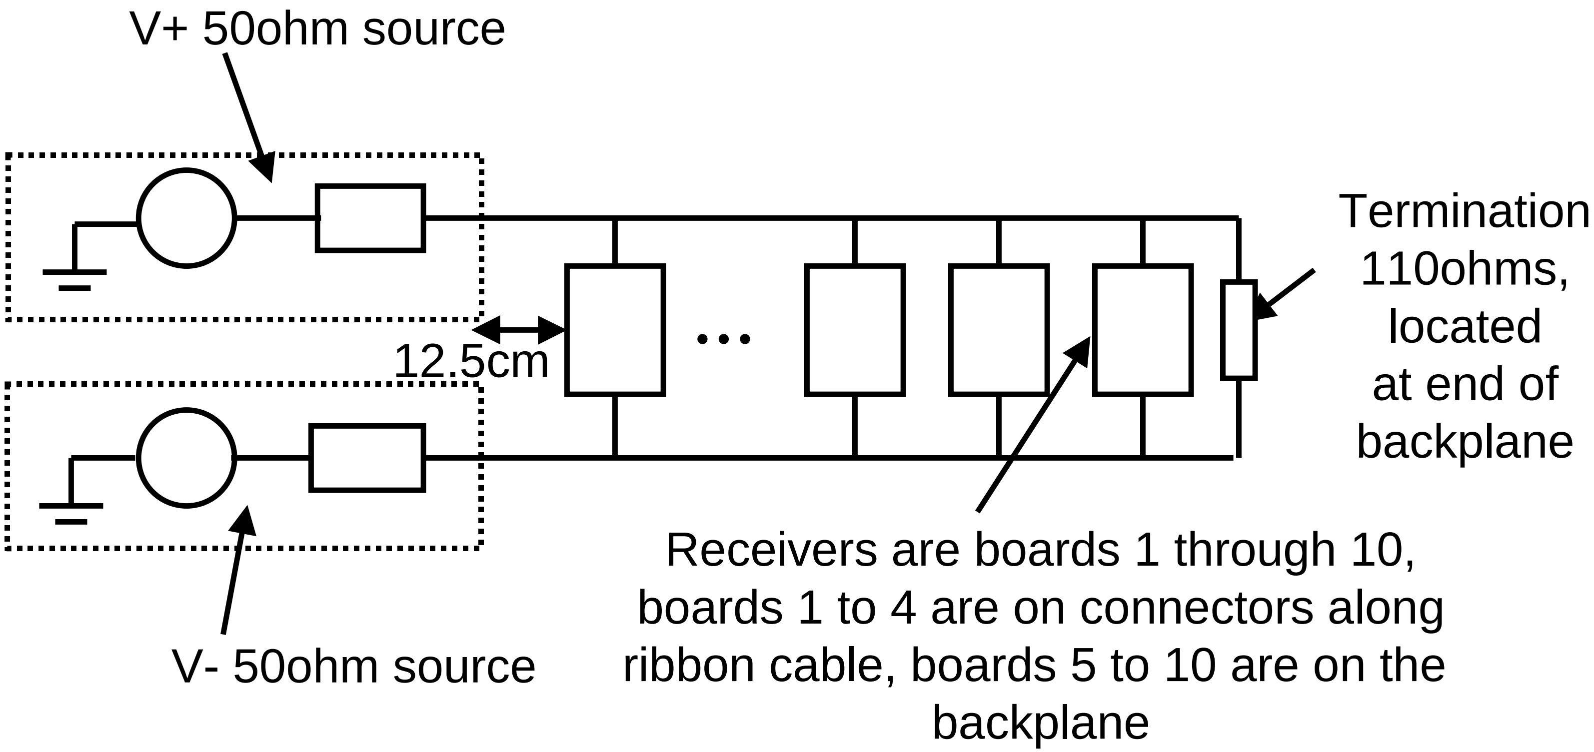

Waveform Data: bd1, ΔA = 0.0

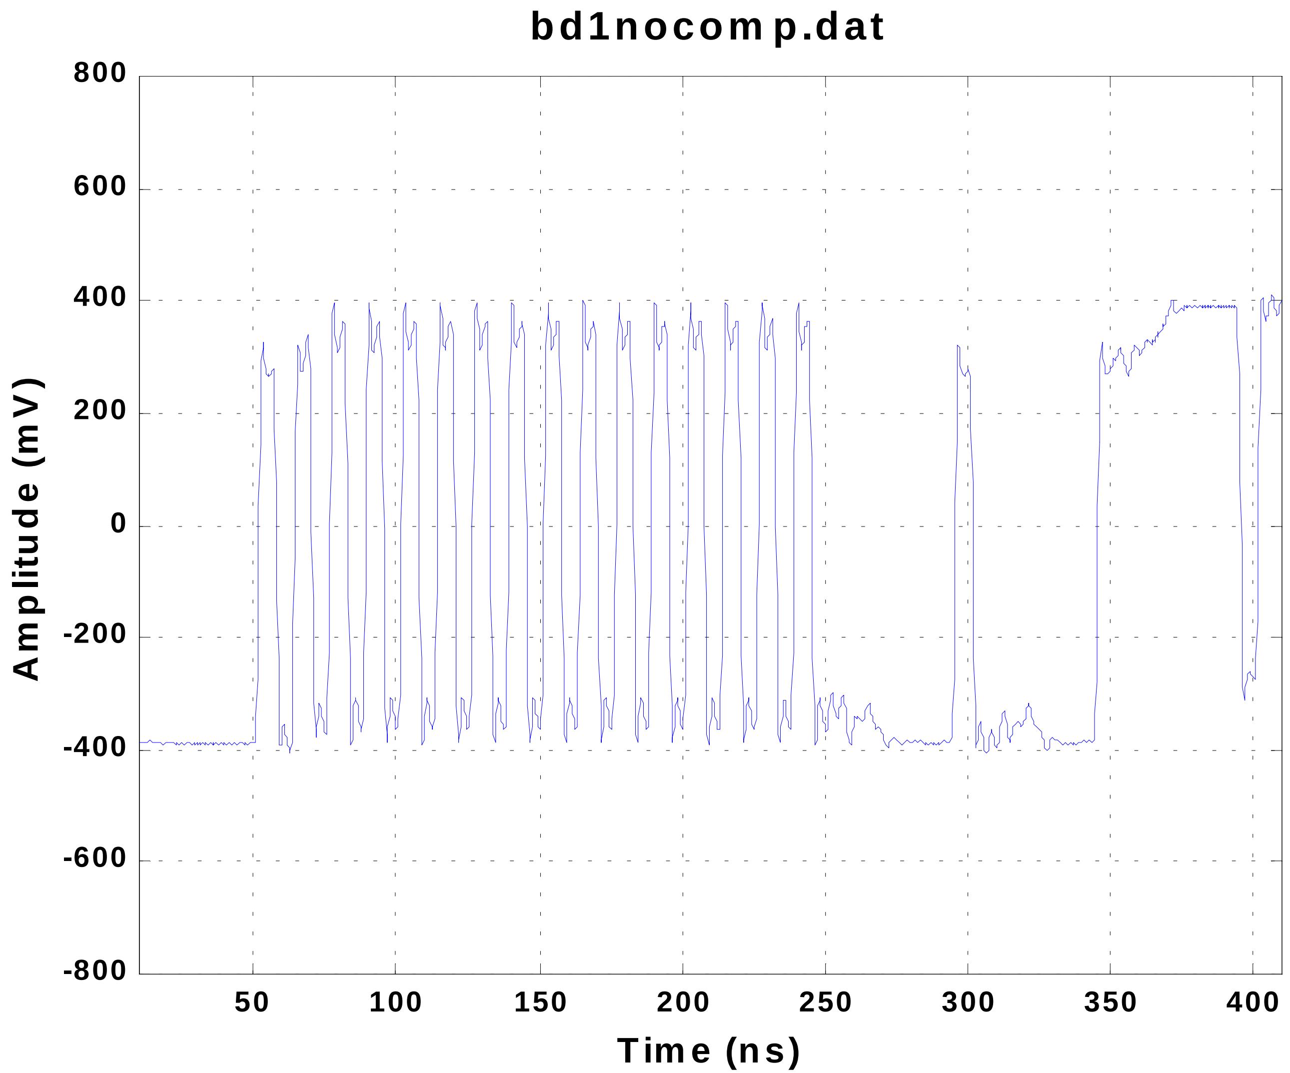

Eye Diagram: bd1, ΔA = 0.0

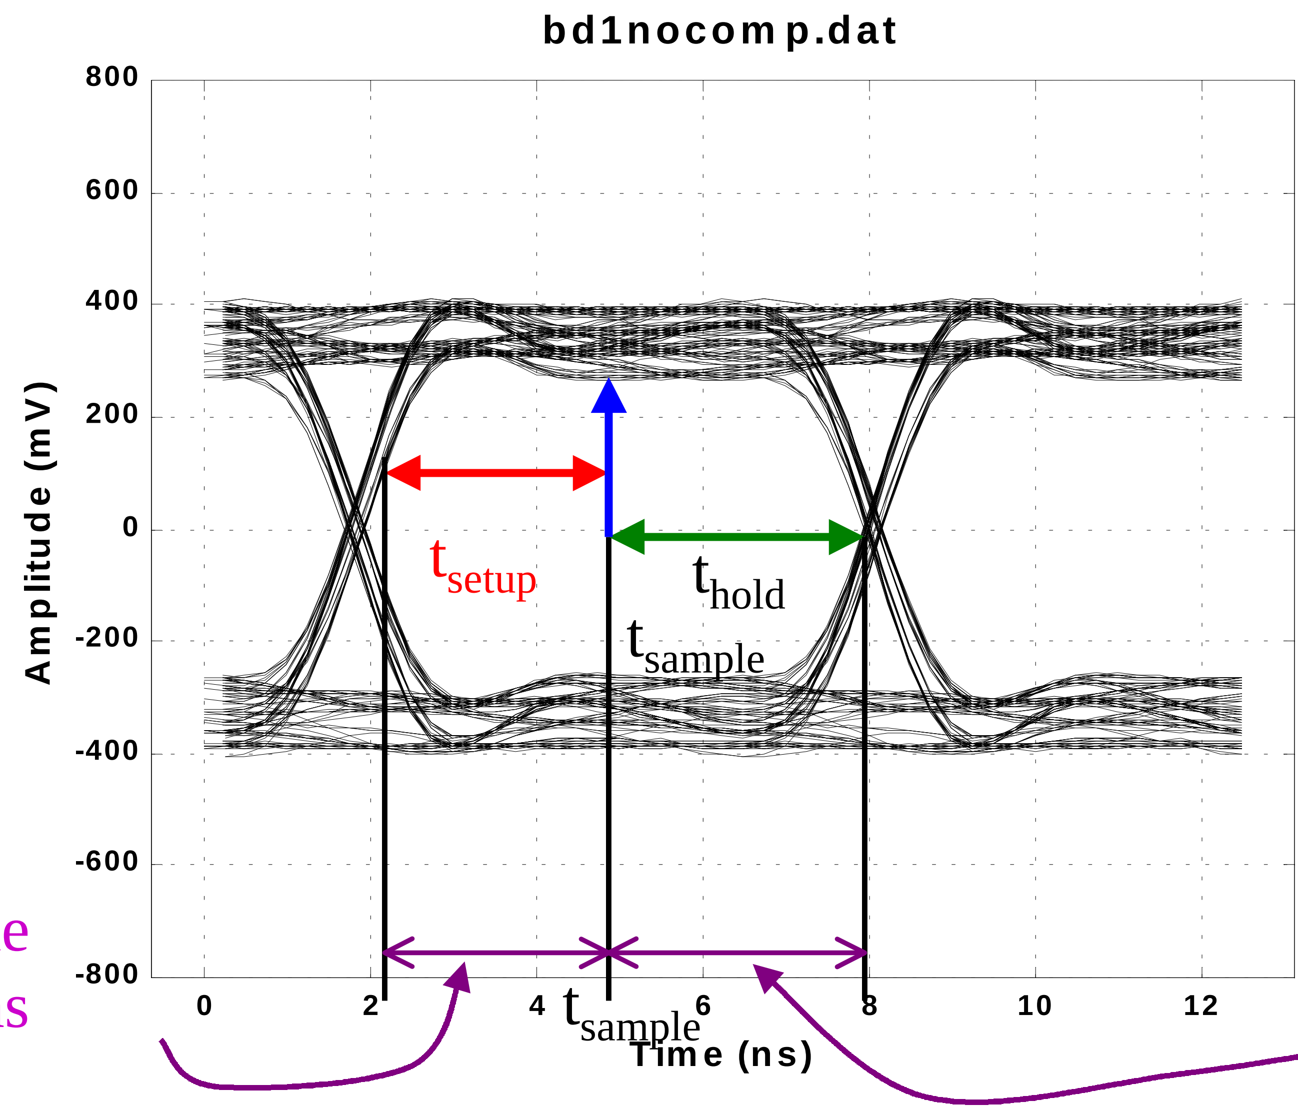

setup time
~2.8ns

hold time
~3.2ns

Waveform Data: bd1, ΔA = 0.2

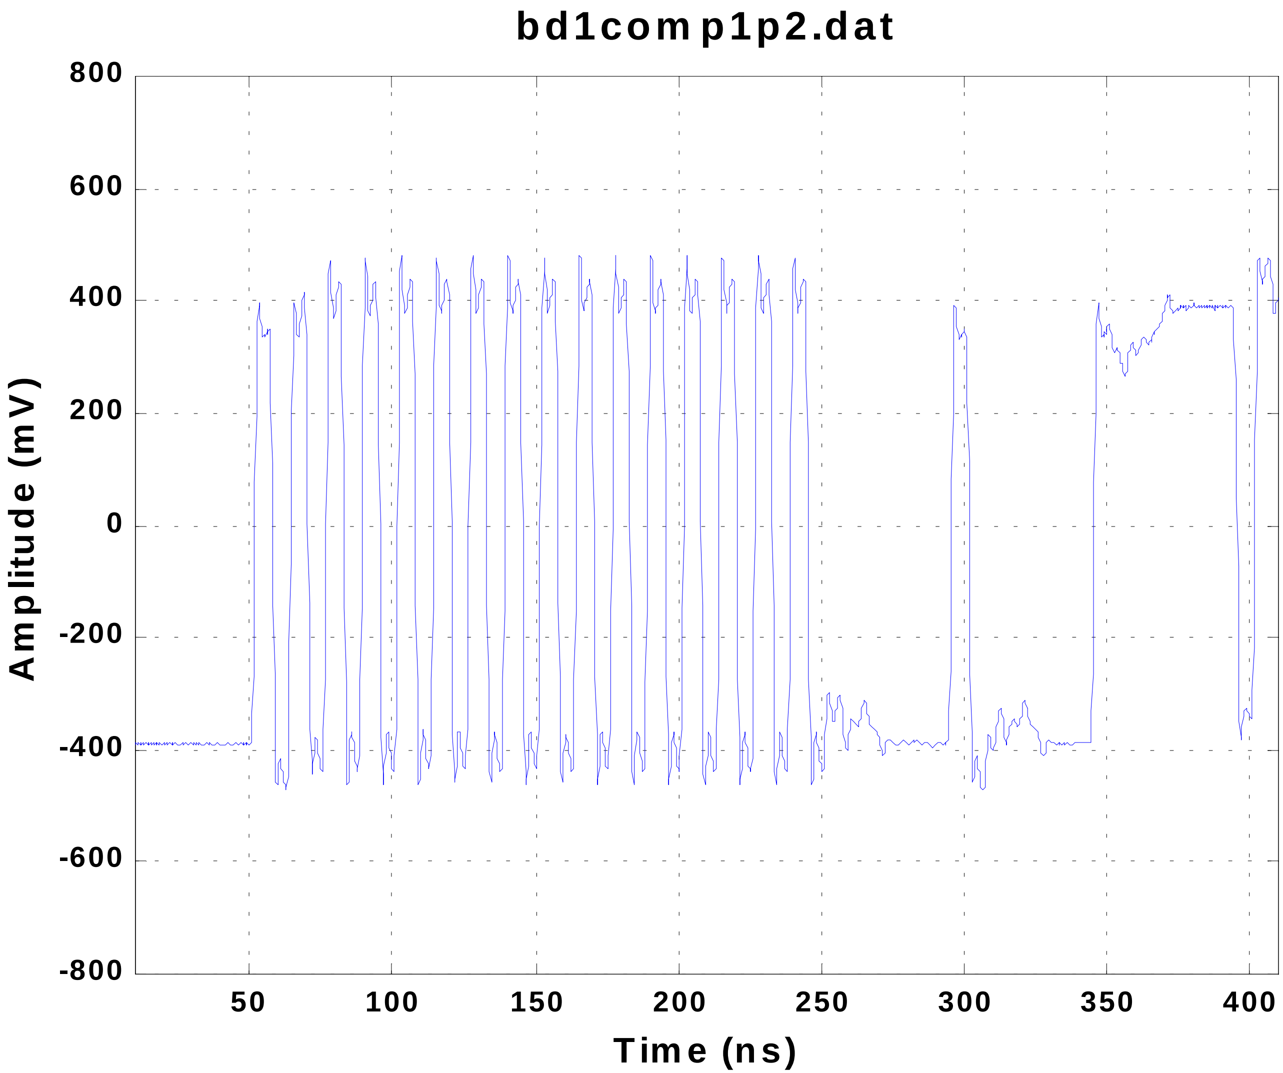

Eye Diagram: bd1, ΔA = 0.2

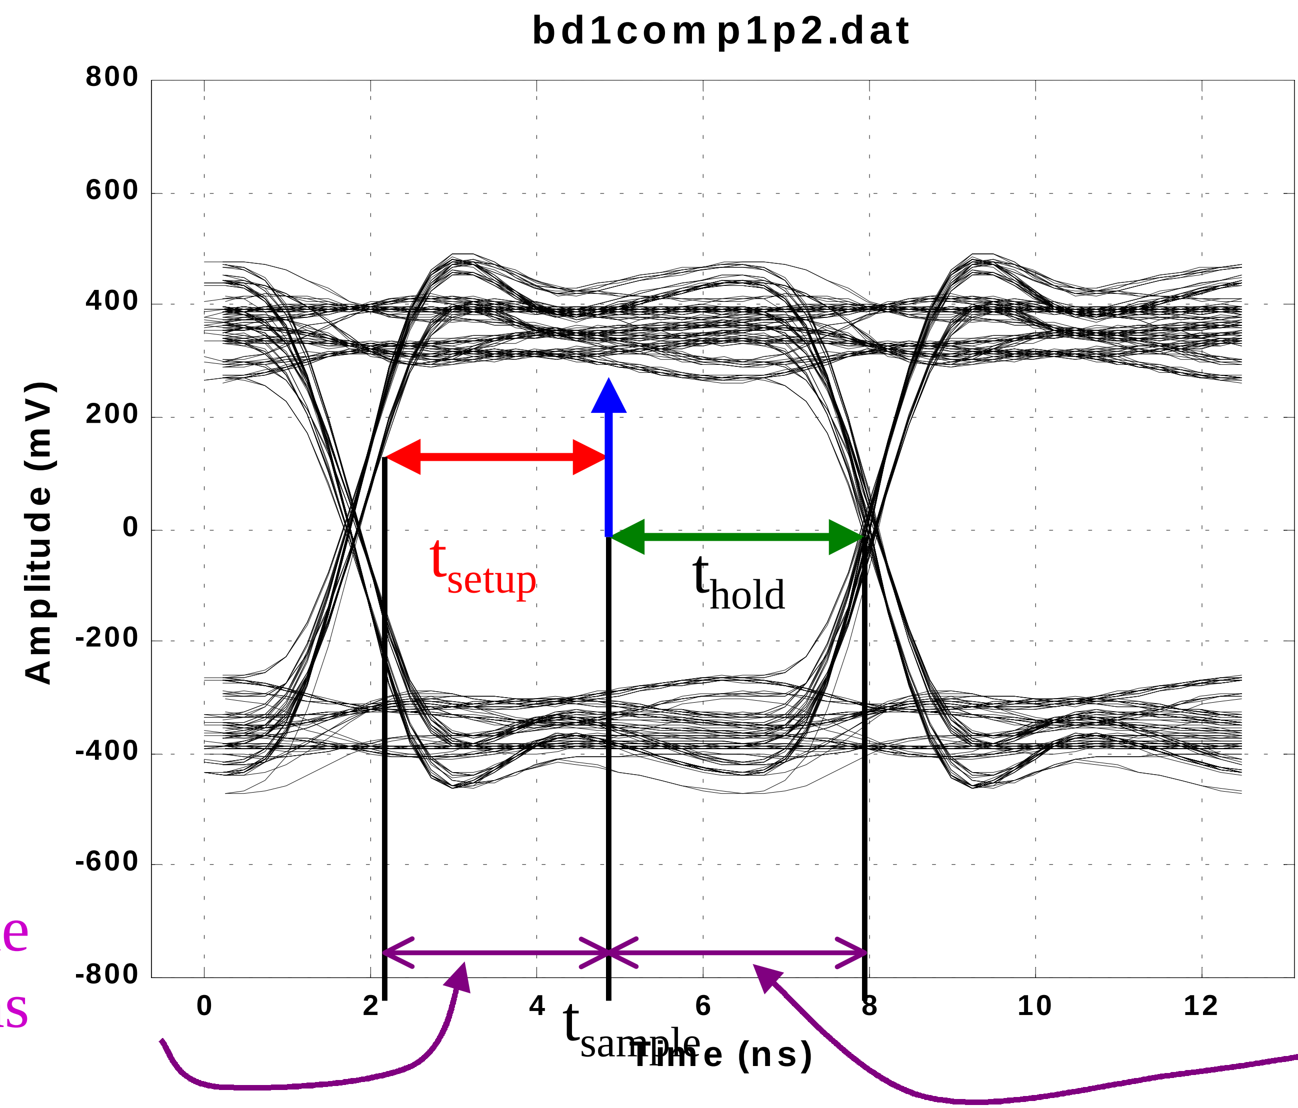

setup time
~2.8ns

hold time
~3.2ns

Waveform Data: bd1, ΔA = 0.5

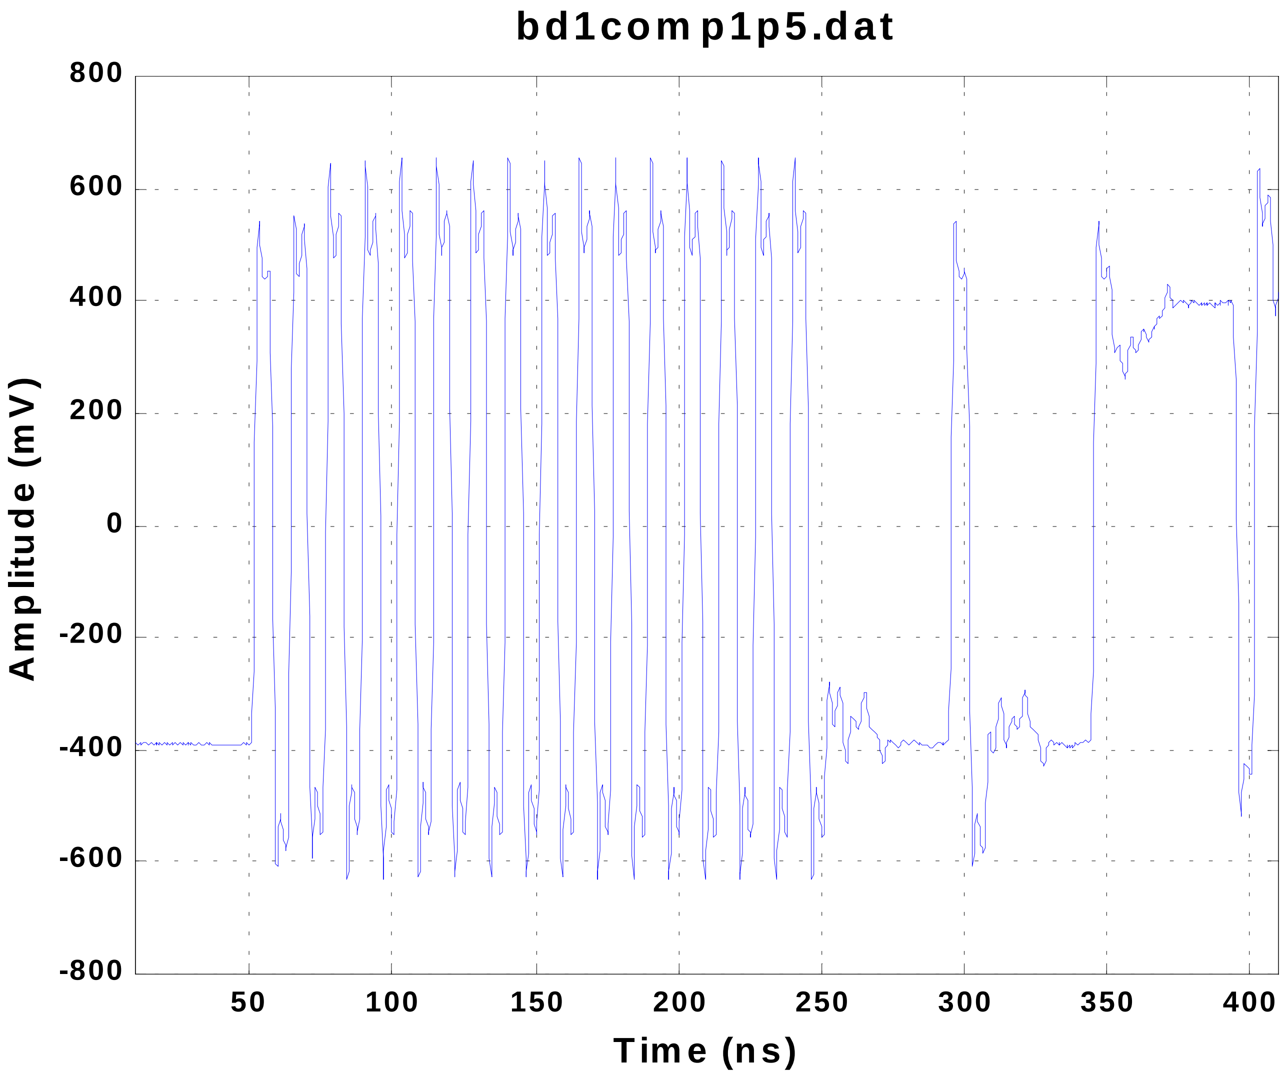

Eye Diagram: bd1, ΔA = 0.5

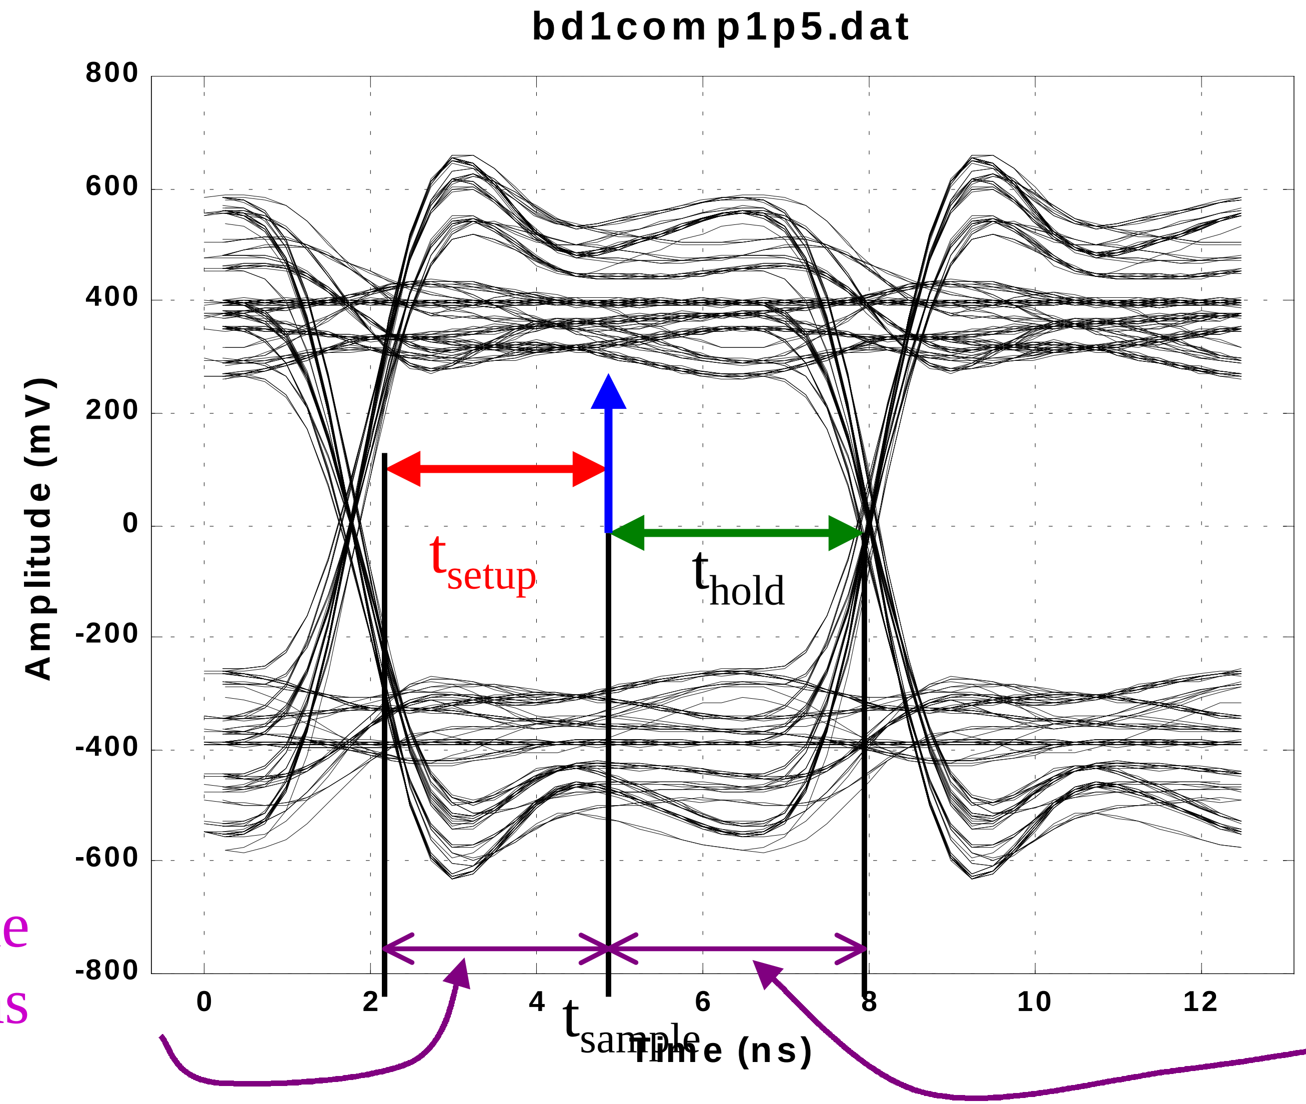

setup time
~2.8ns

hold time
~3.2ns

Quantum™

Waveform Data: bd1, ΔA = 0.8

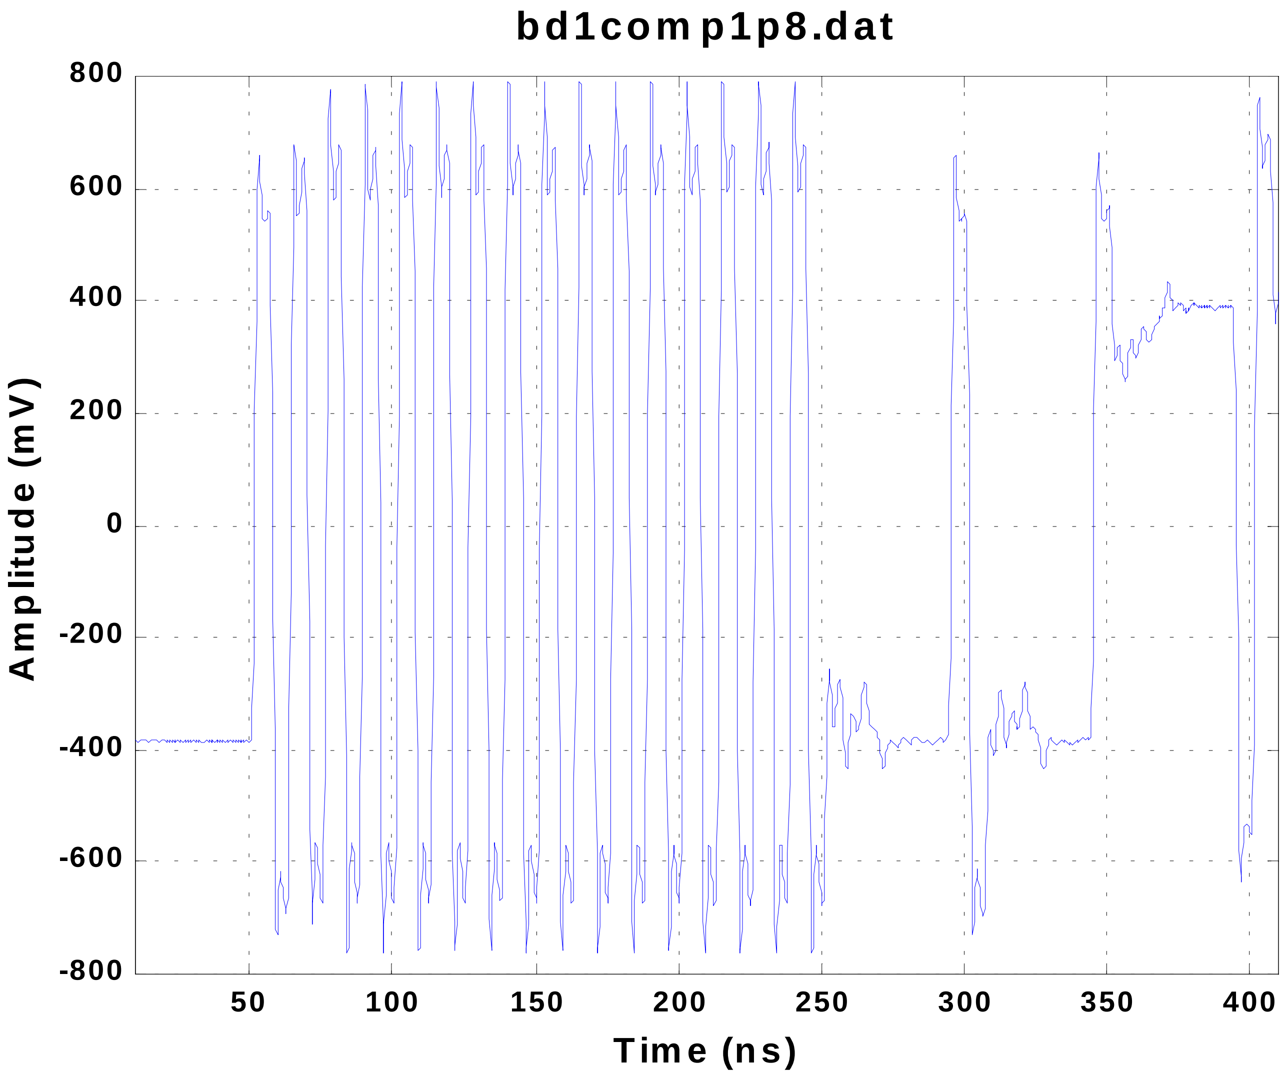

Quantum™

Eye Diagram: bd1, ΔA = 0.8

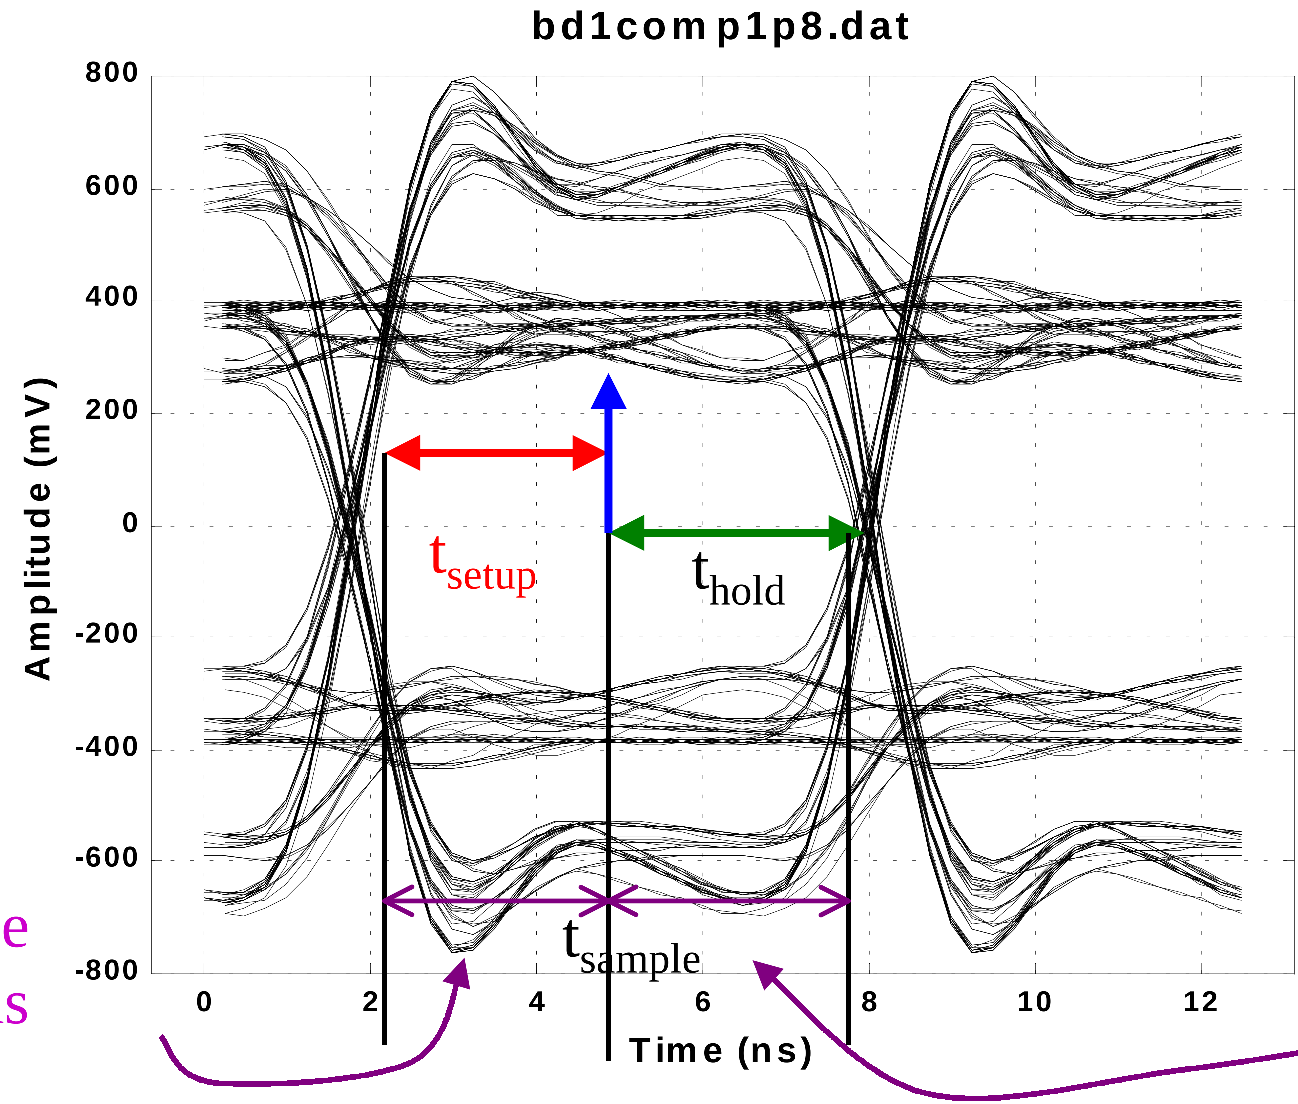

setup time ~2.8ns

hold time ~3.0ns

Waveform Data: bd4, ΔA = 0.0

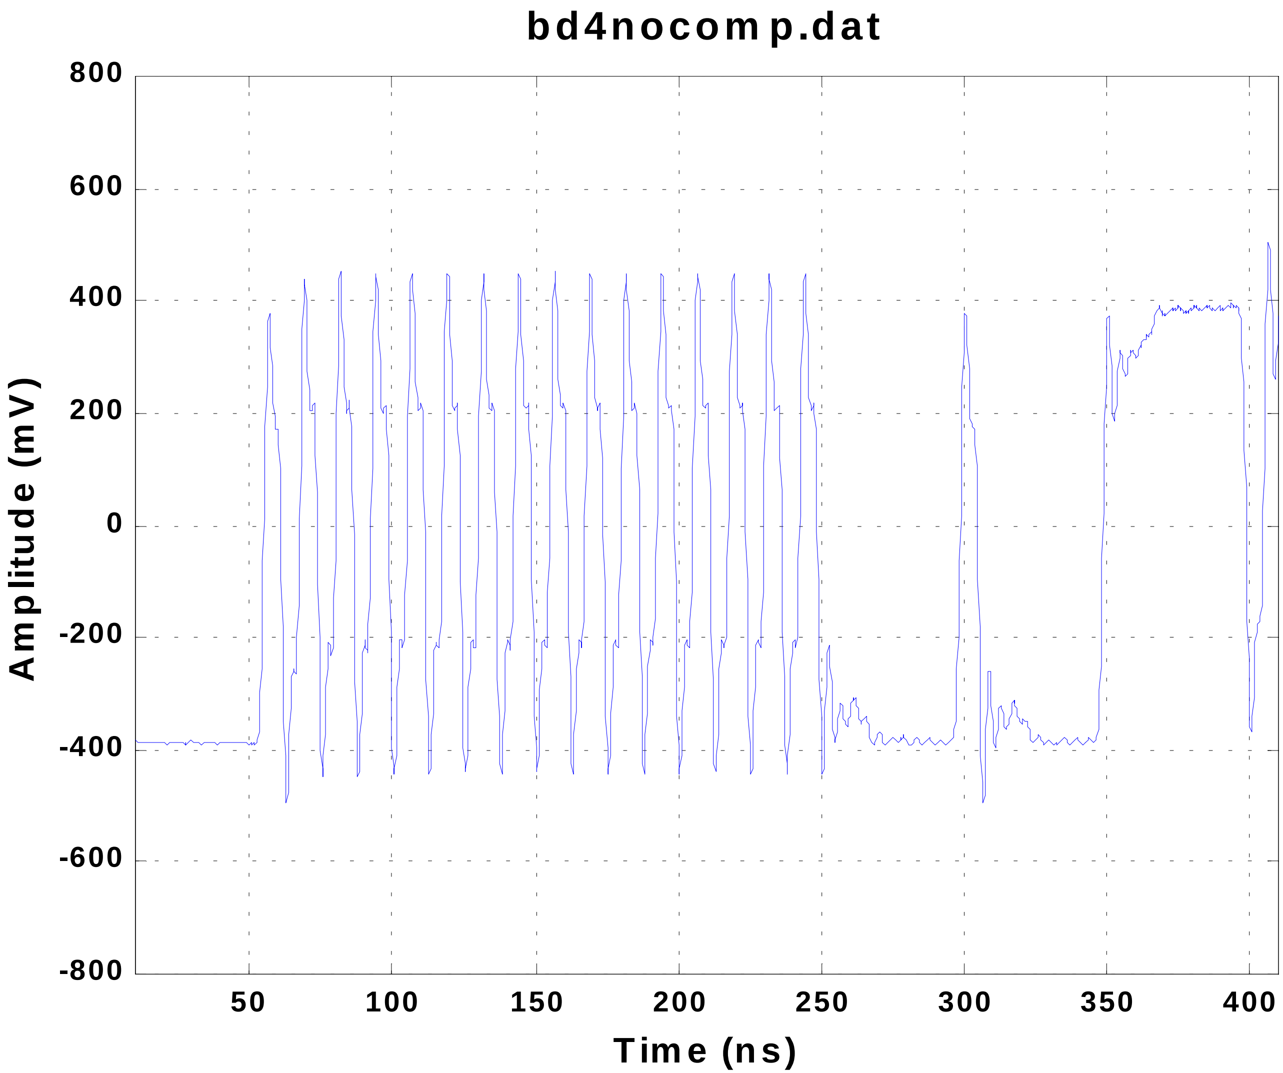

Eye Diagram: bd4, ΔA = 0.0

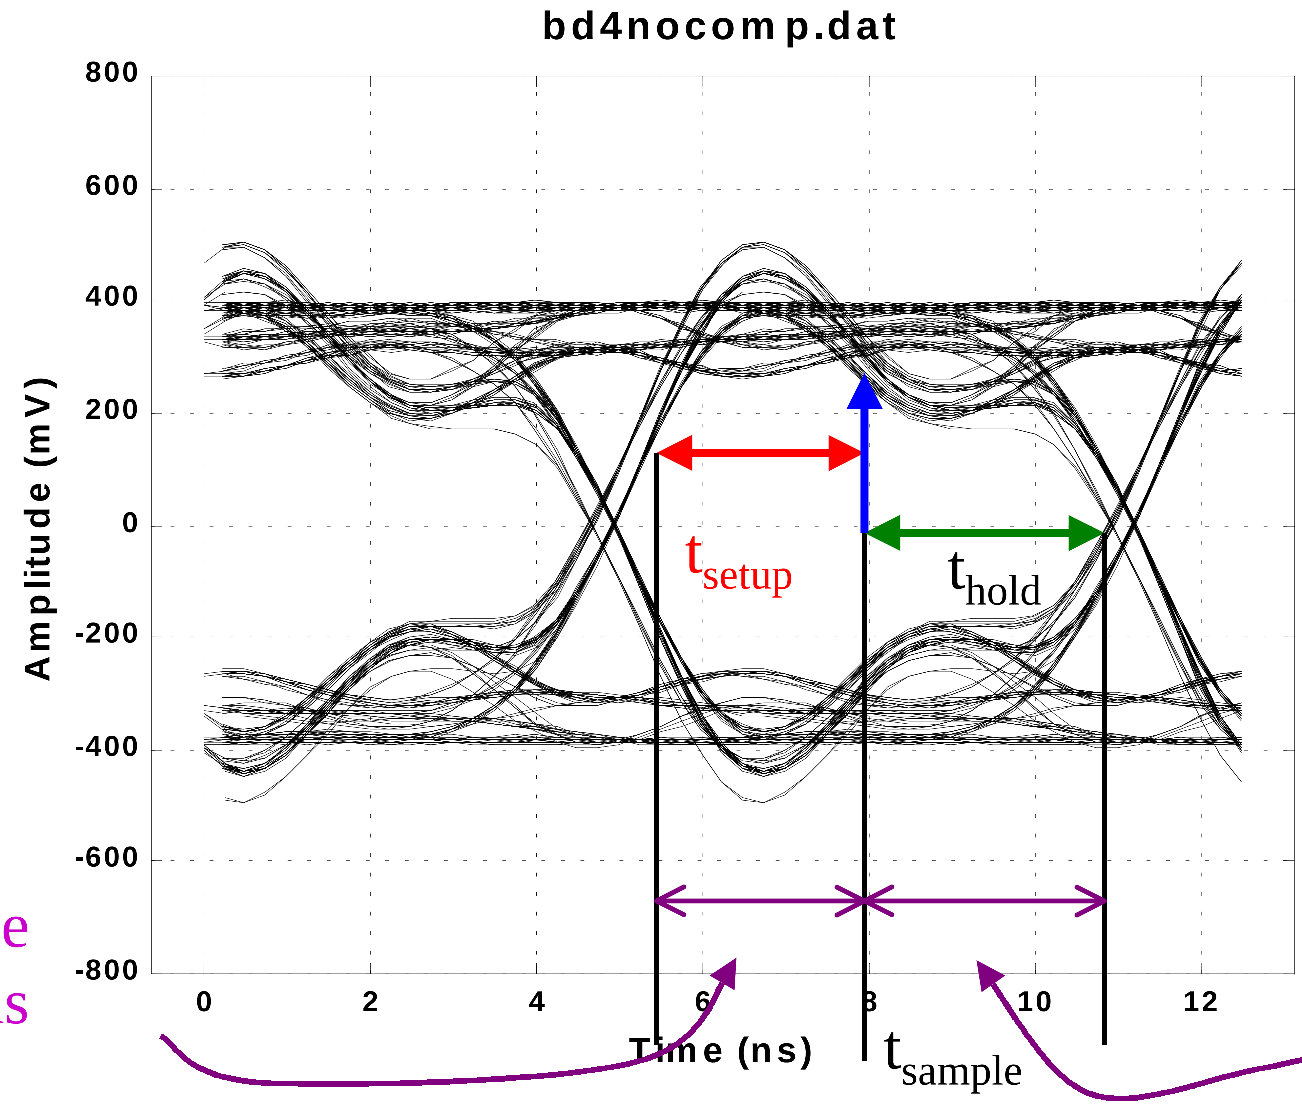

setup time
~2.6ns

hold time
~3.0ns

Waveform Data: bd4, ΔA = 0.2

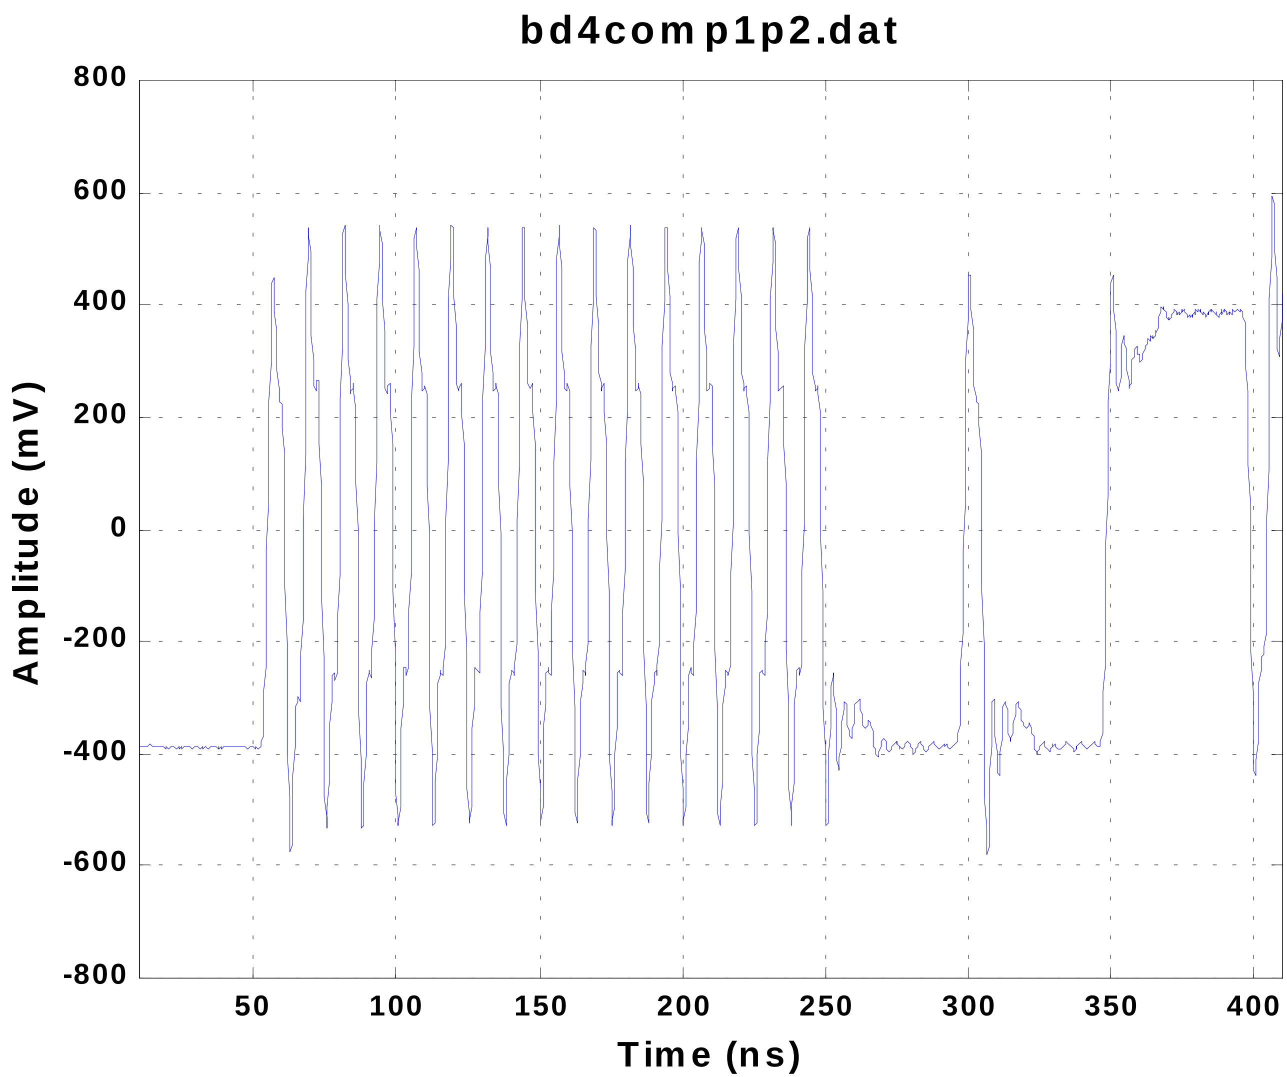

Eye Diagram: bd4, ΔA = 0.2

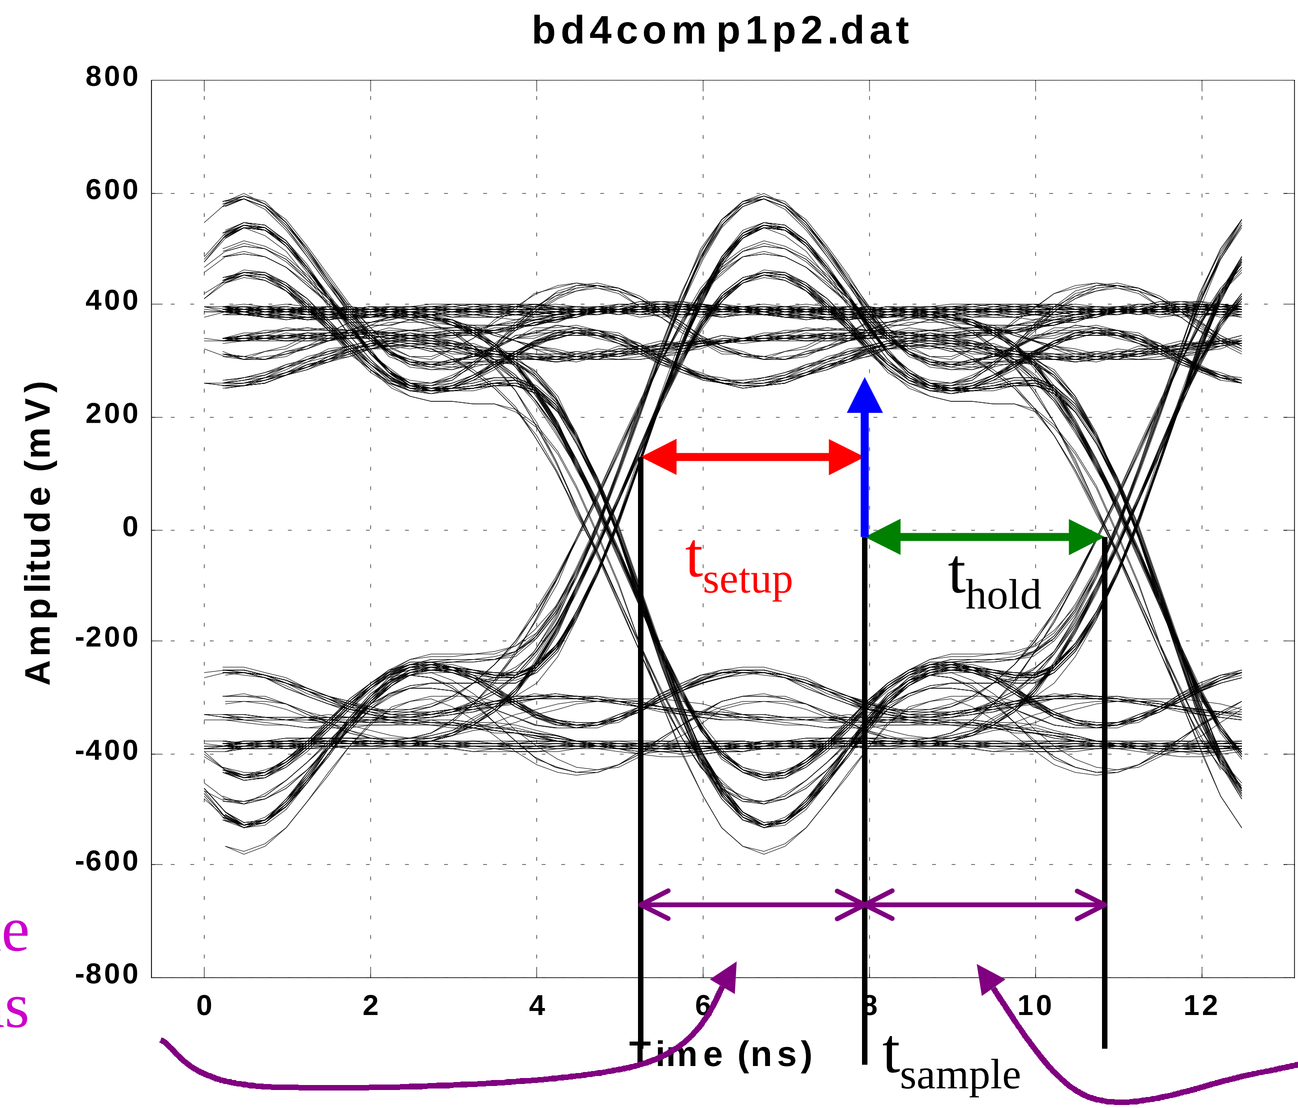

setup time
~2.8ns

hold time
~3.0ns

Waveform Data: bd4, ΔA = 0.5

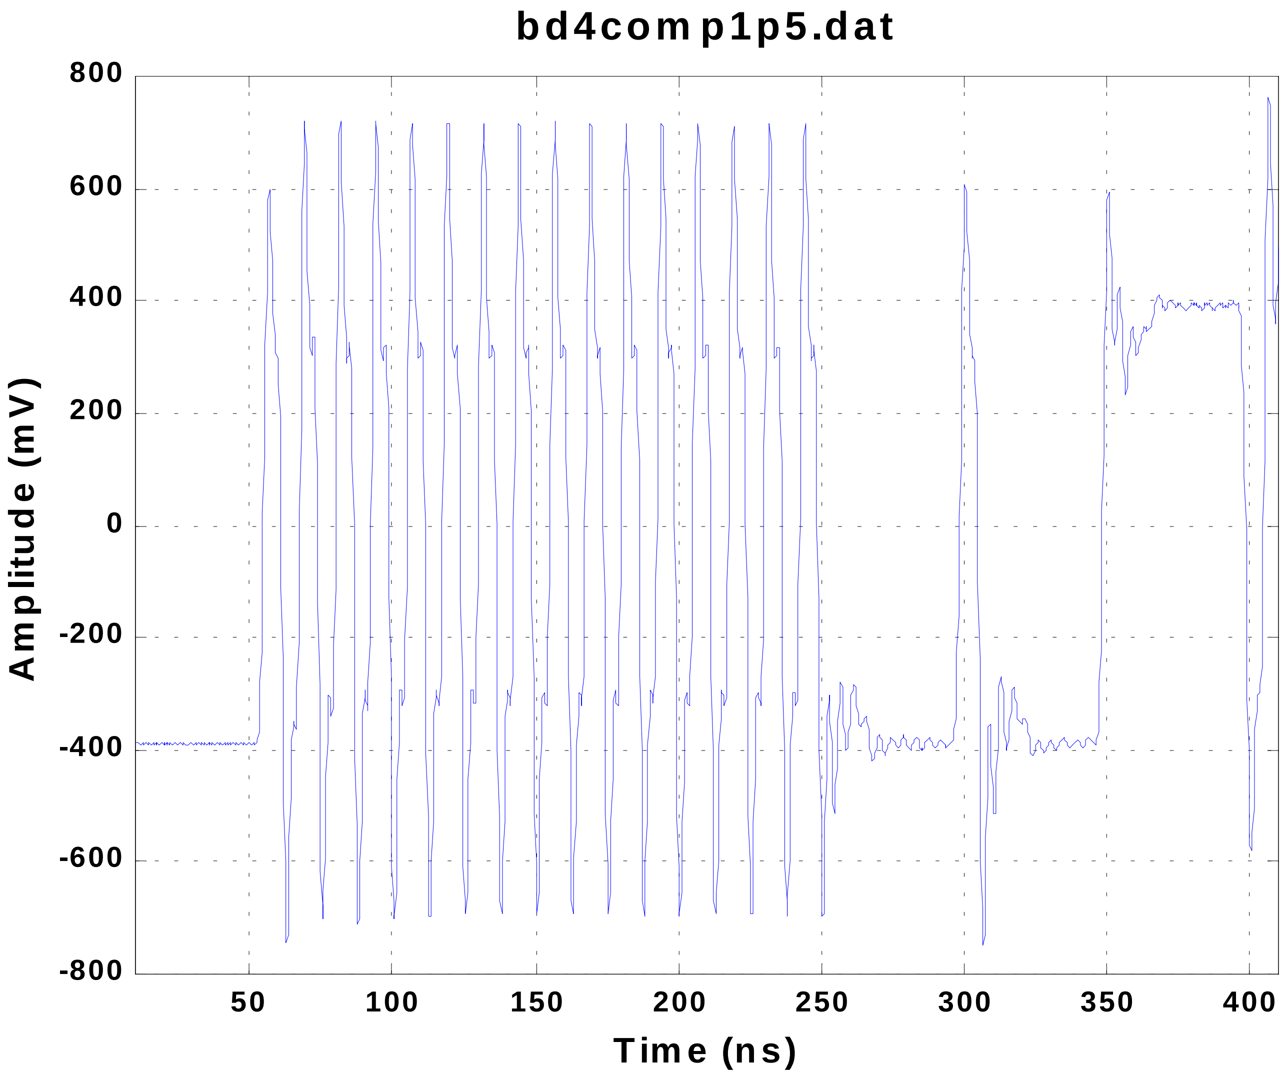

Eye Diagram: bd4, ΔA = 0.5

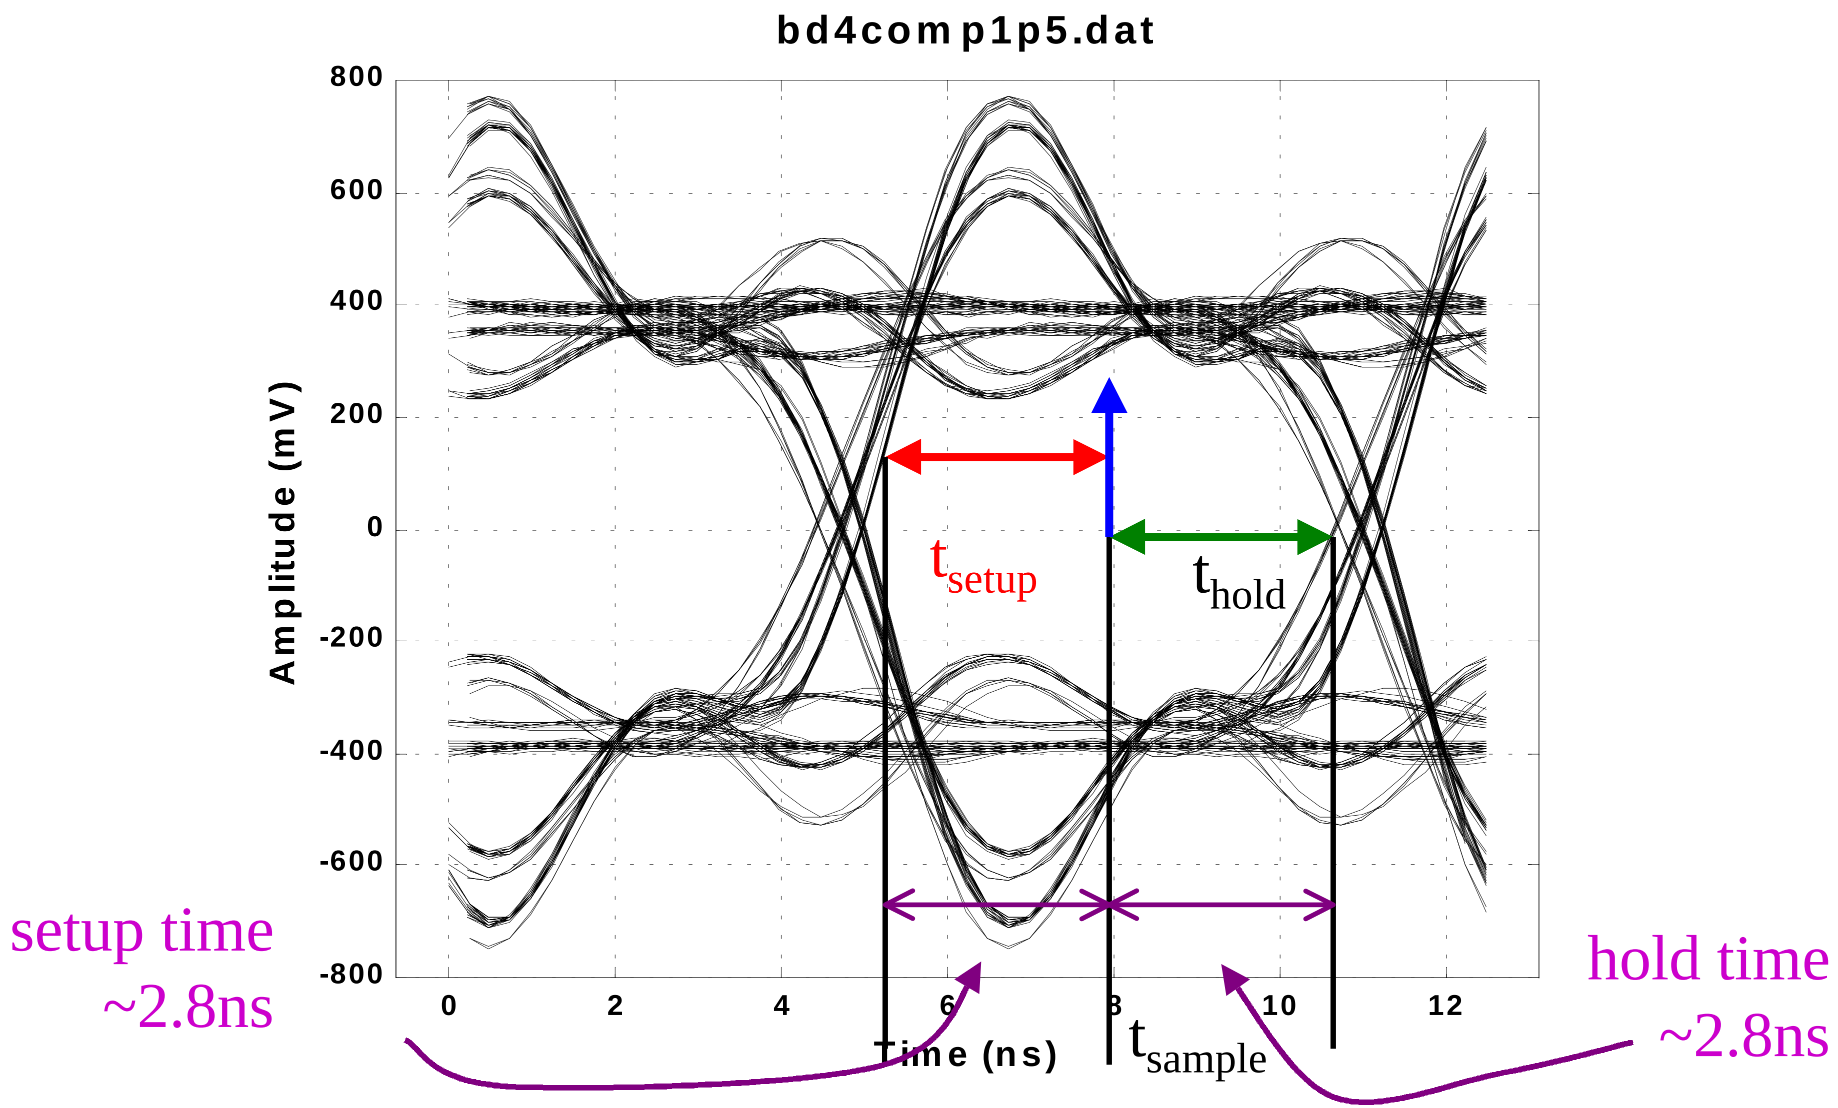

Waveform Data: bd4, ΔA = 0.8

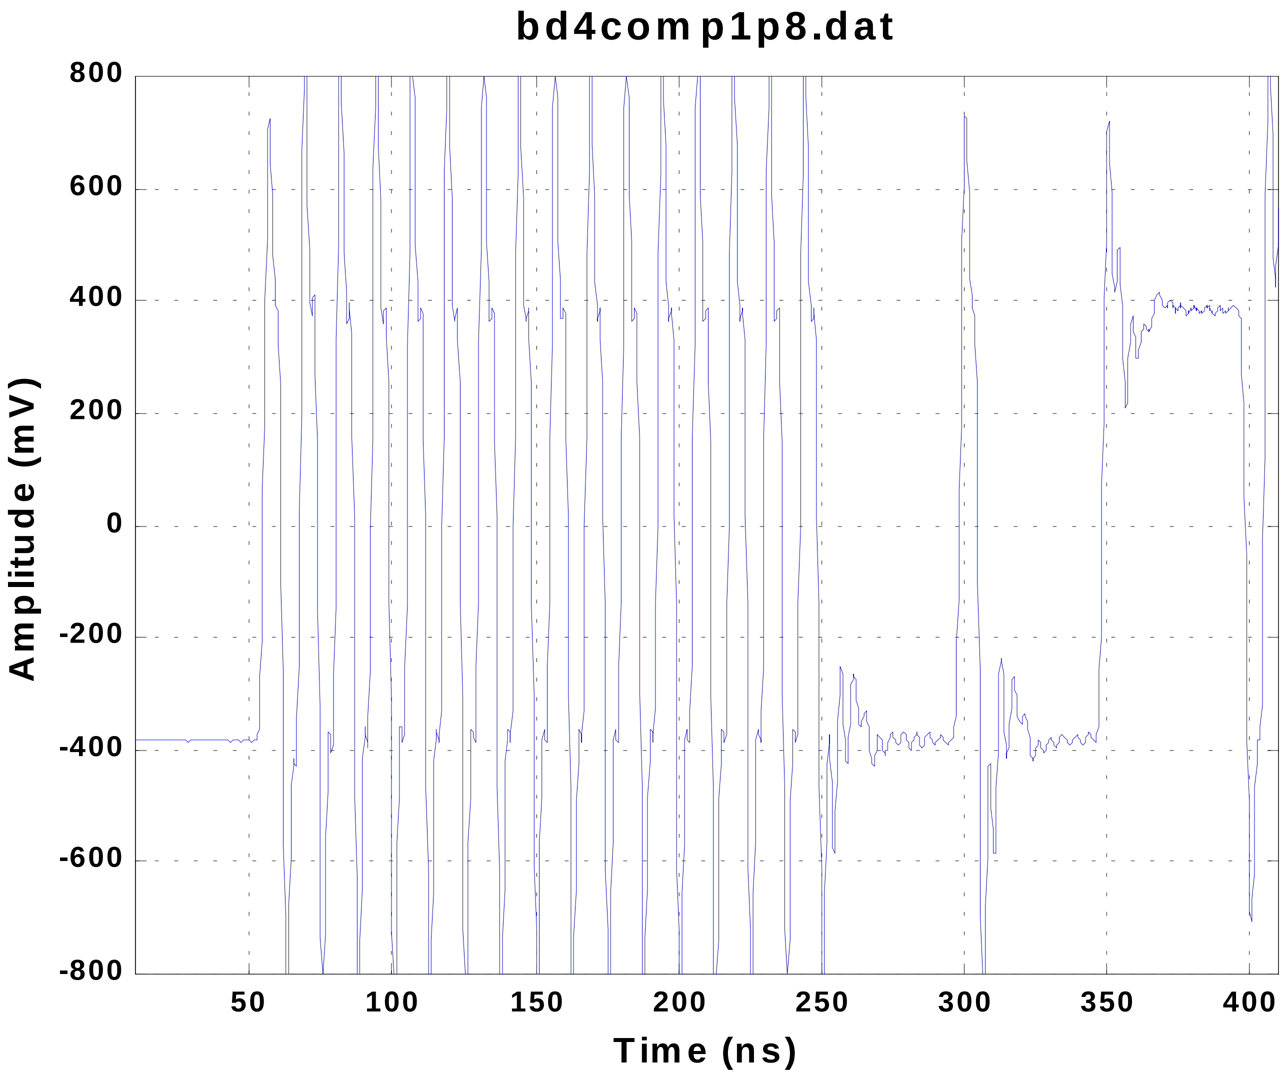

Eye Diagram: bd4, ΔA = 0.8

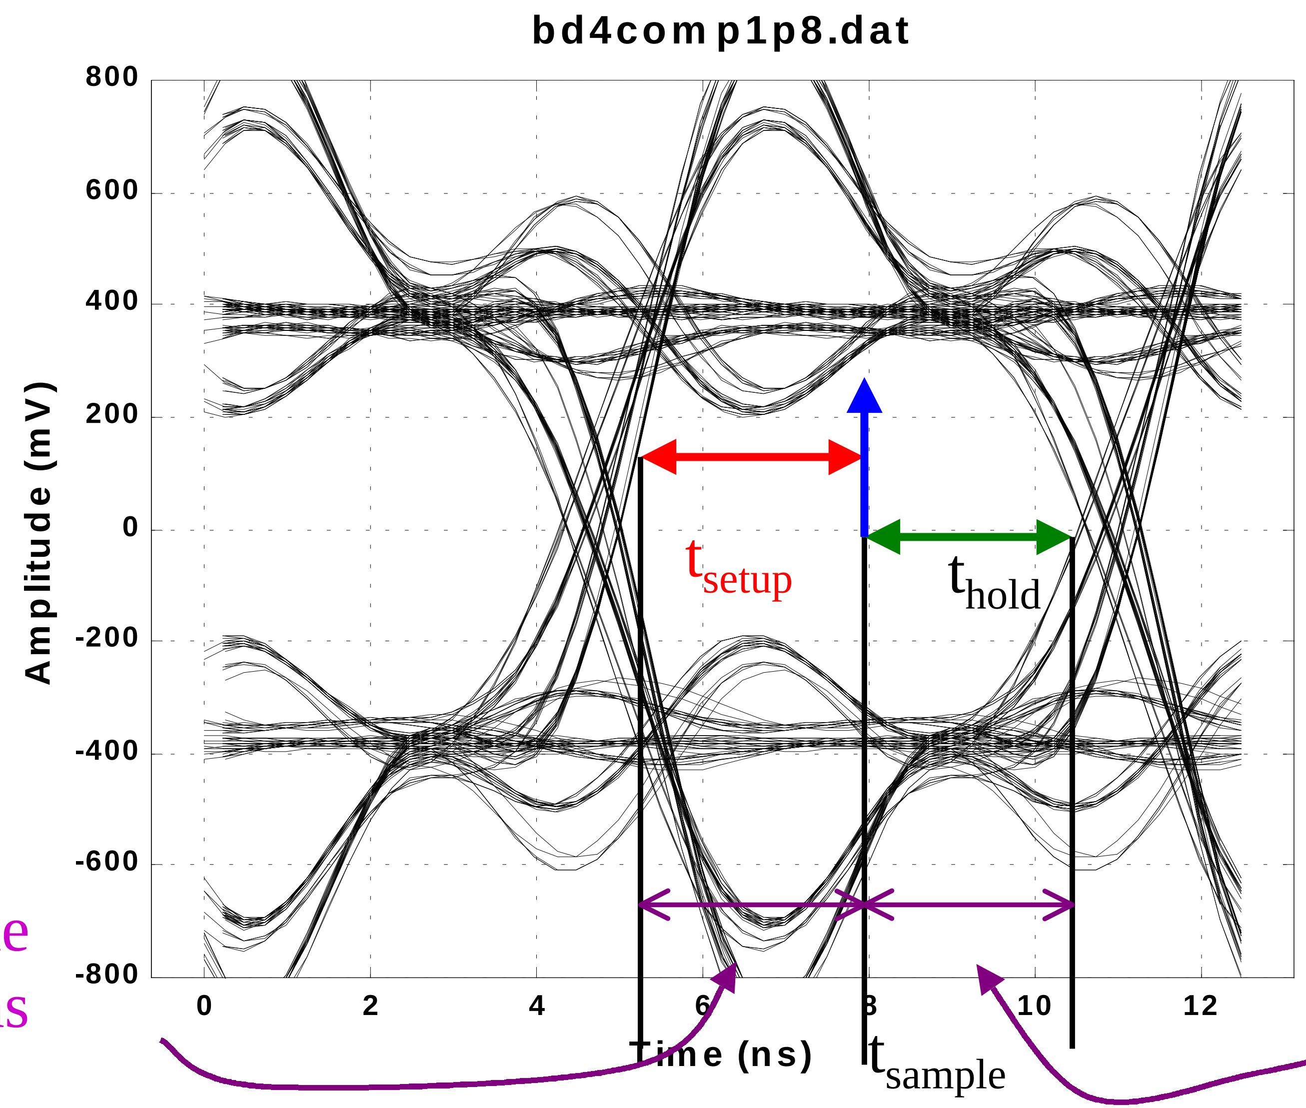

setup time
~2.8ns

hold time
~2.6ns

Waveform Data: bp5, ΔA = 0.0

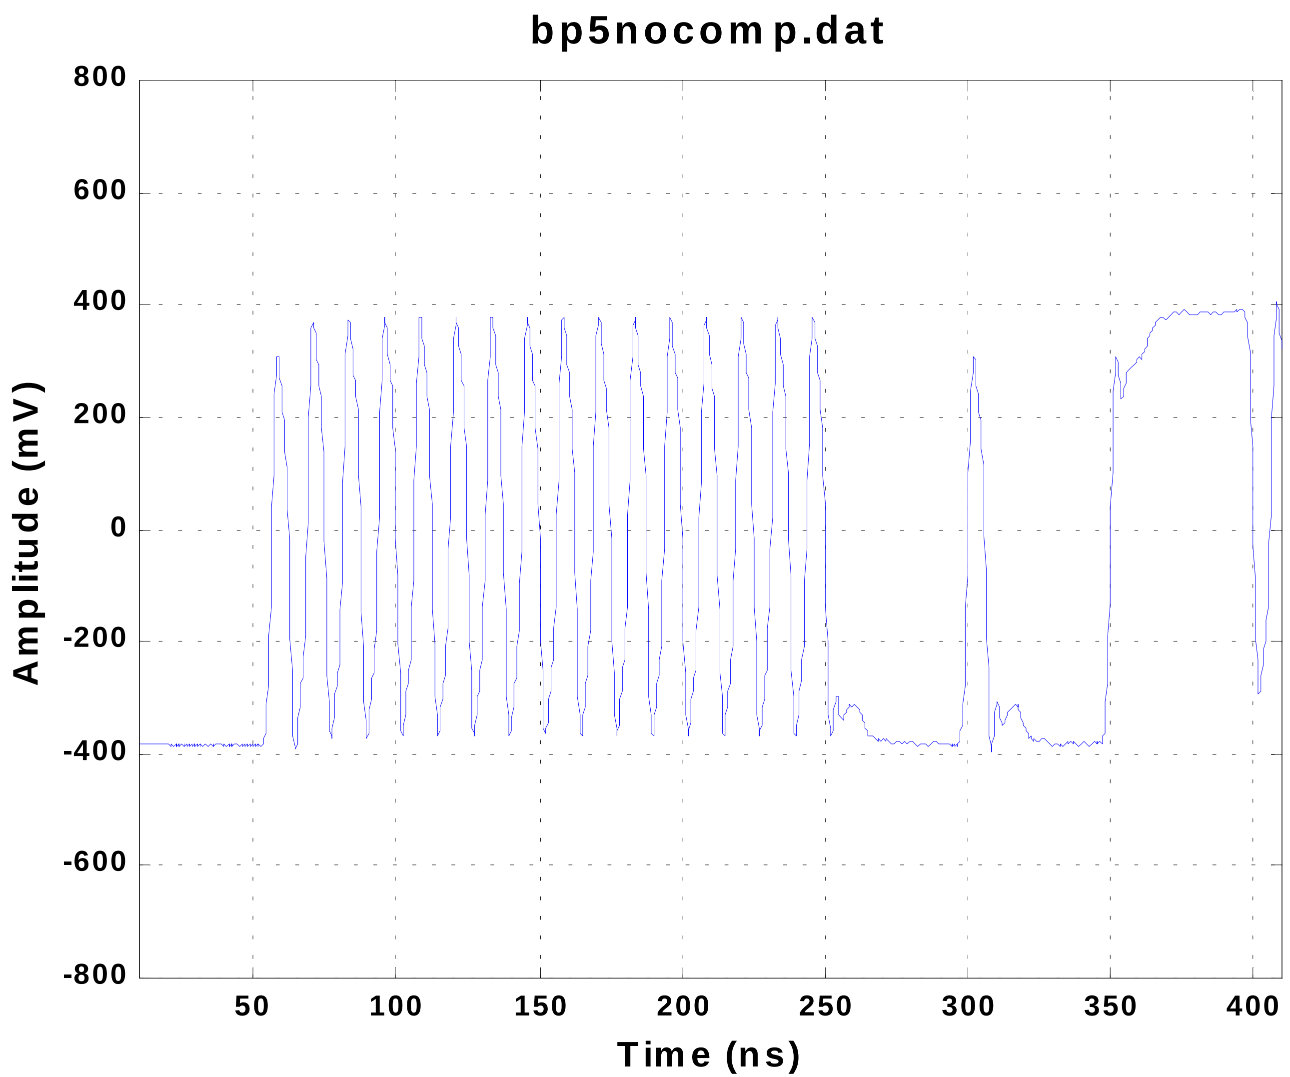

Eye Diagram: bp5, ΔA = 0.0

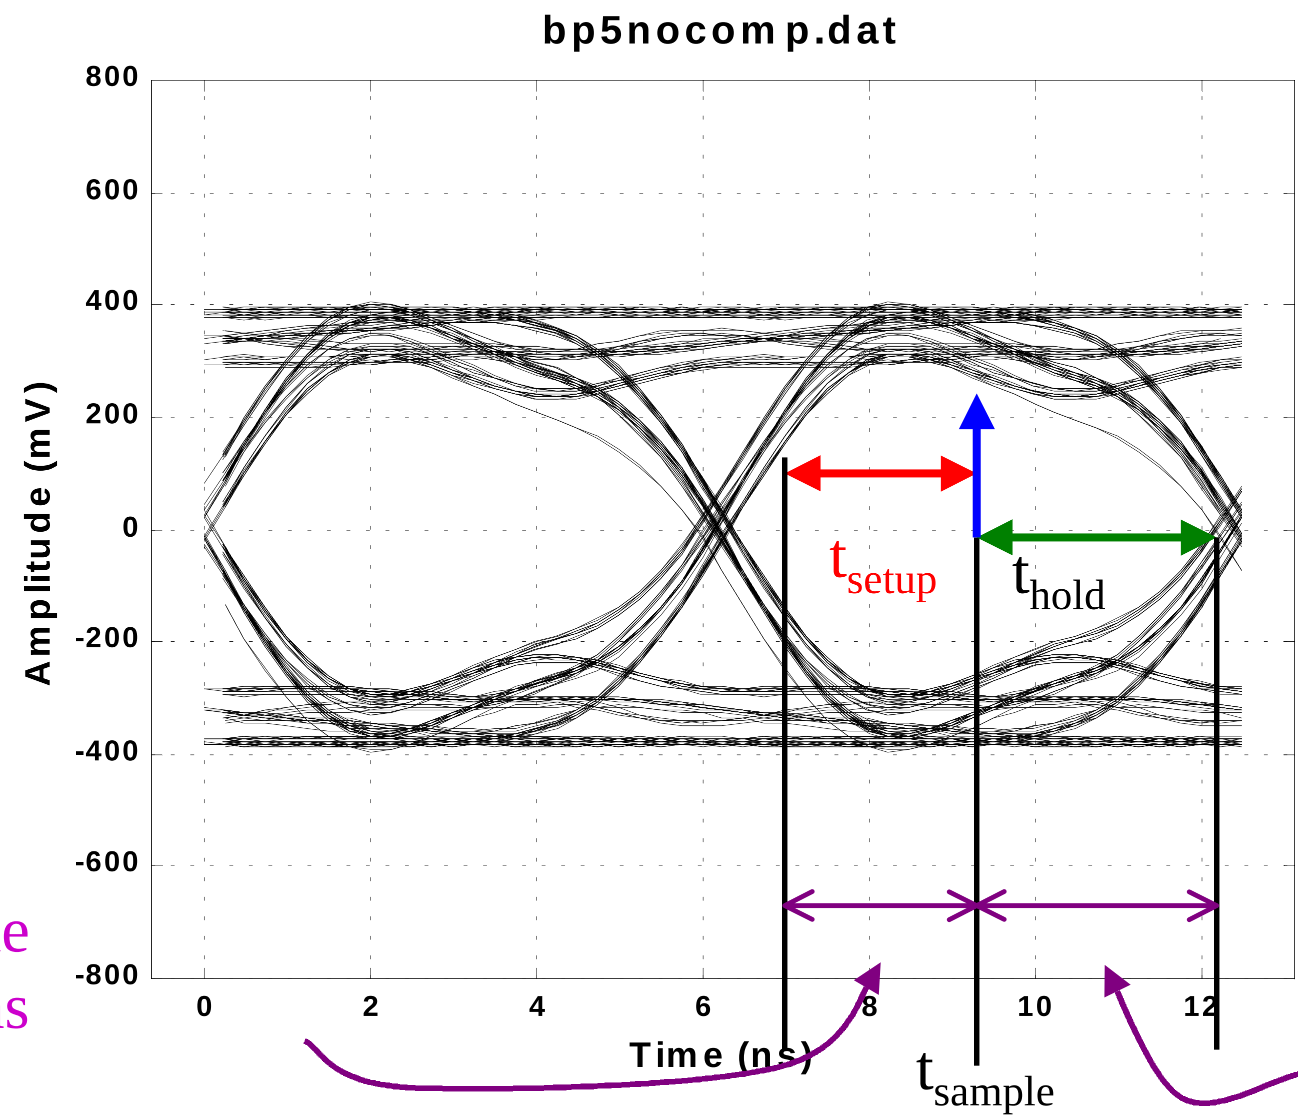

setup time
~2.4ns

hold time
~3.0ns

Waveform Data: bp5, ΔA = 0.2

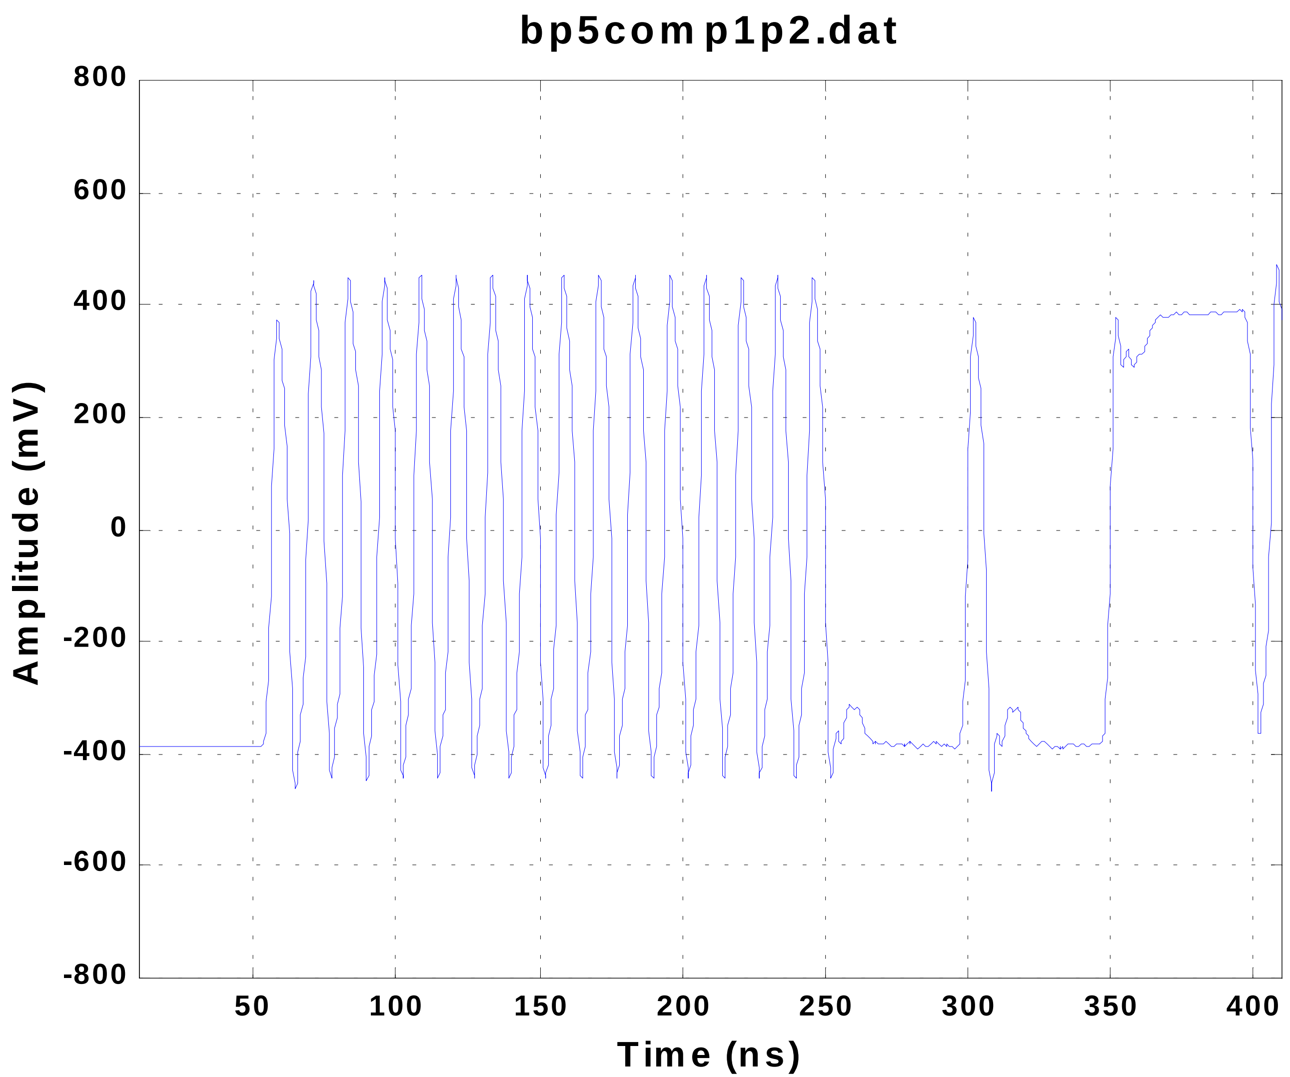

Eye Diagram: bp5, ΔA = 0.2

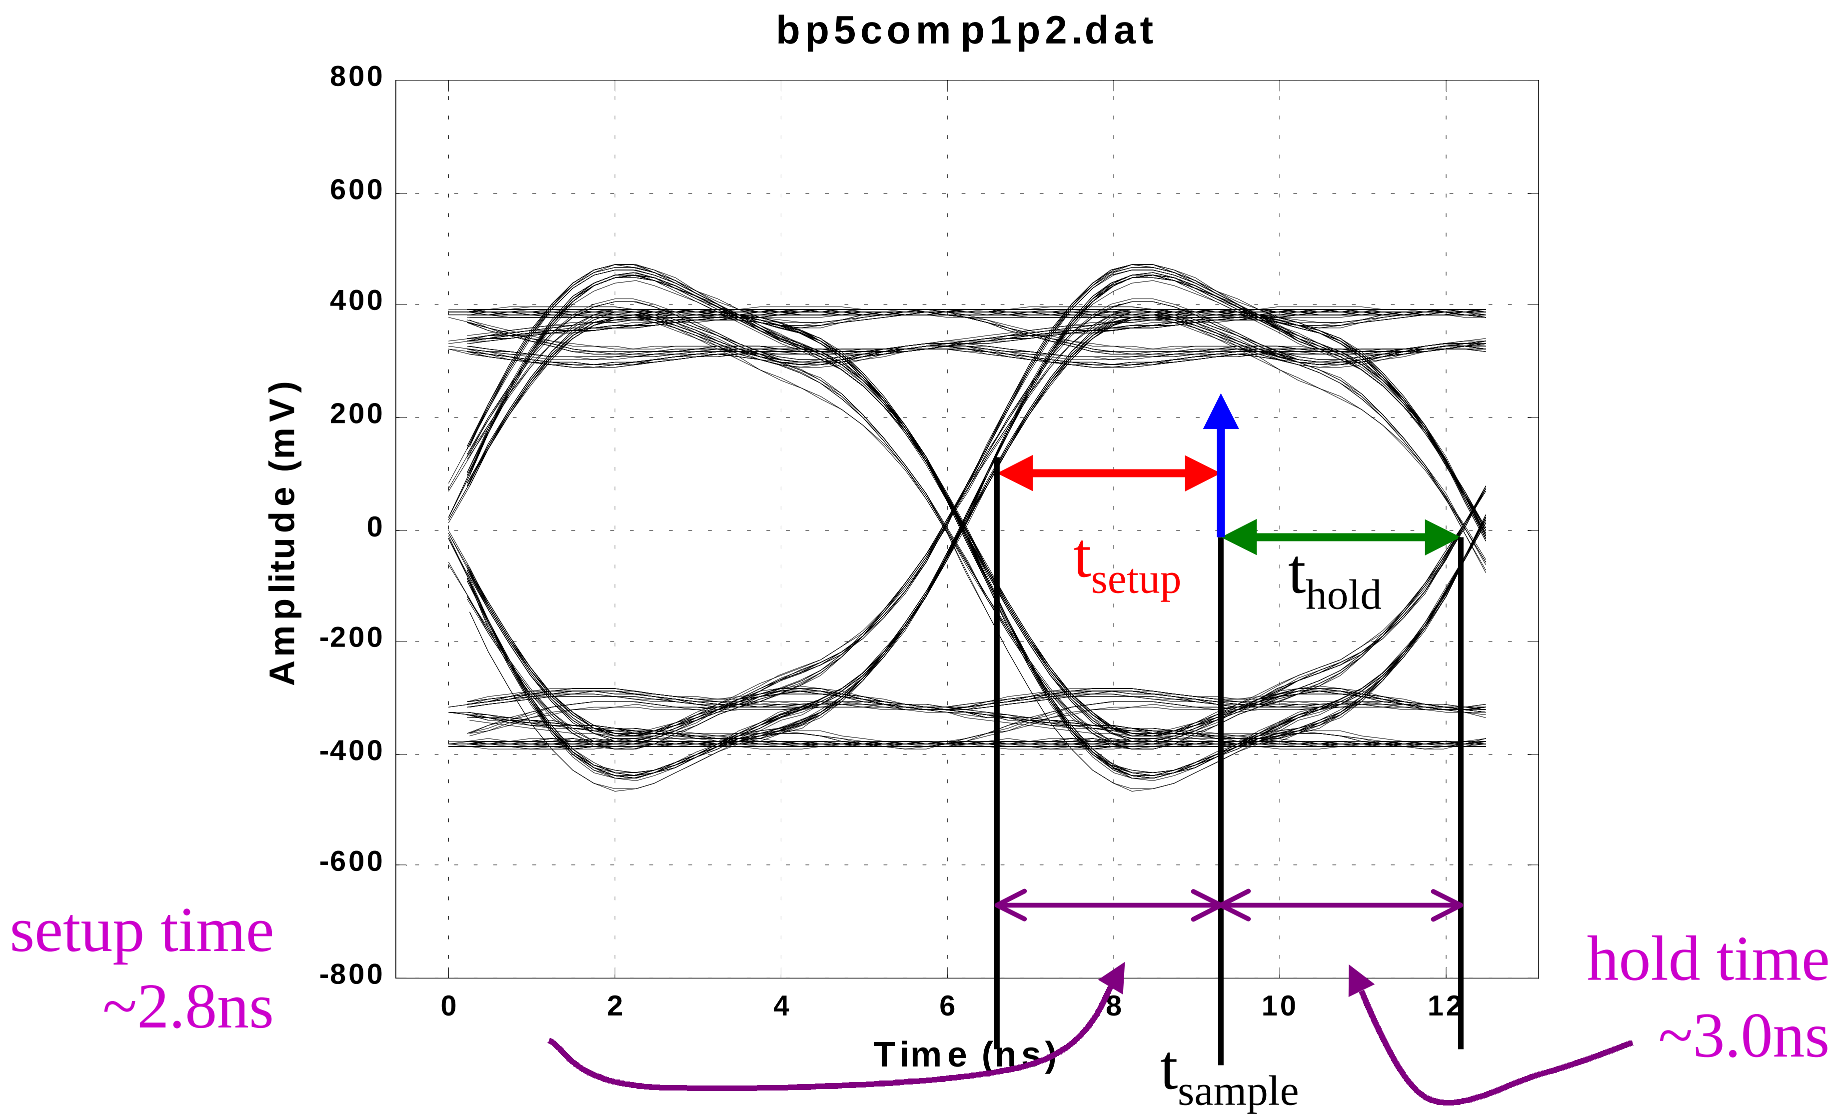

Waveform Data: bp5, ΔA = 0.5

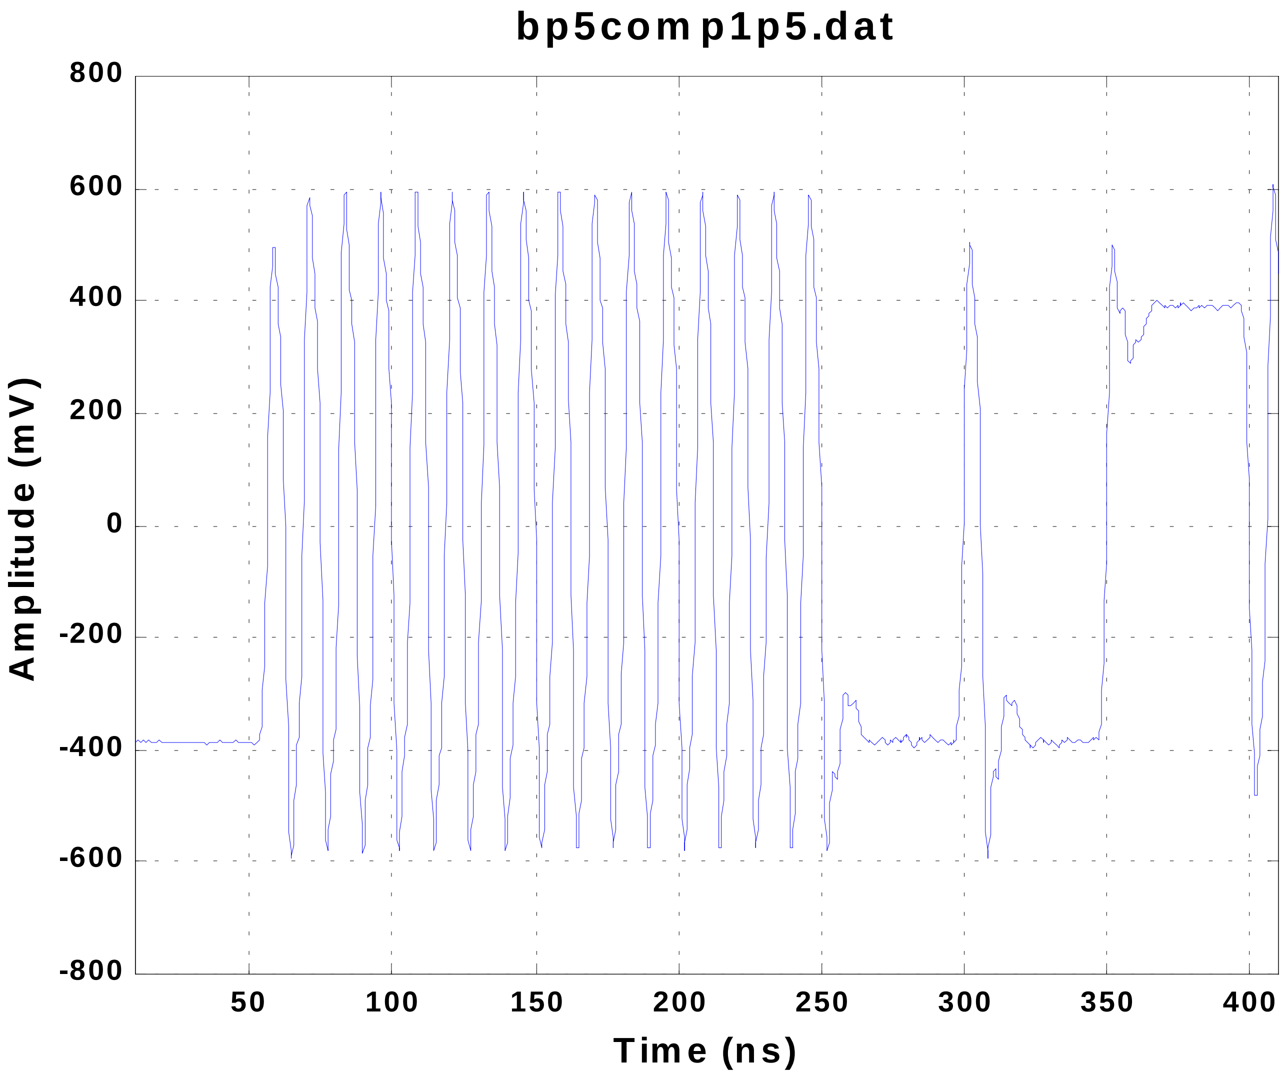

Eye Diagram: bp5, ΔA = 0.5

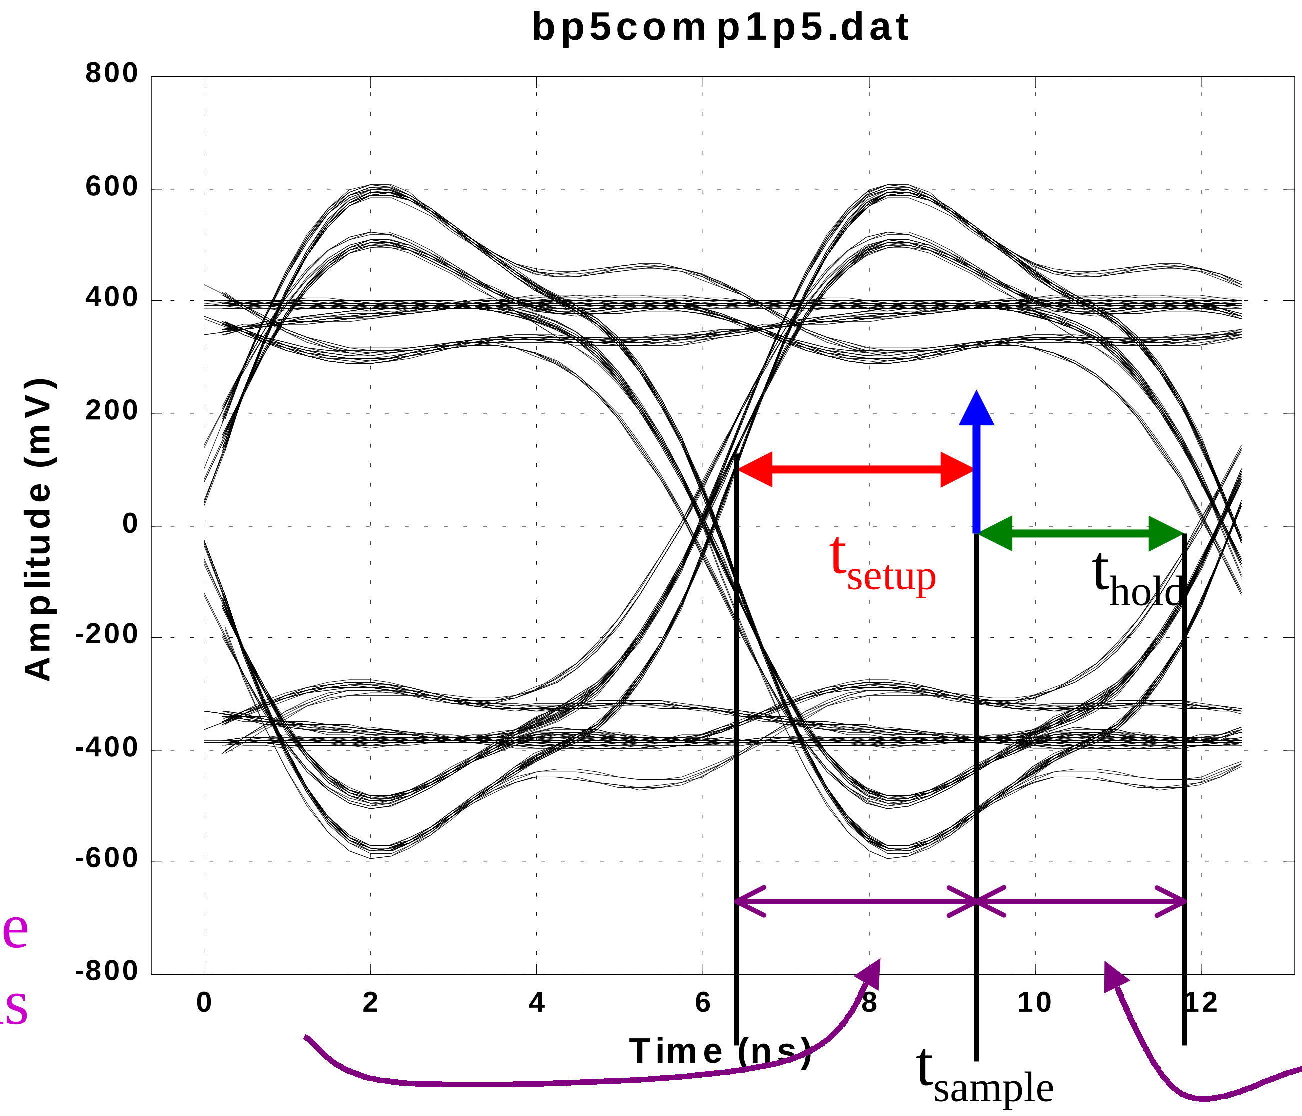

setup time
~3.0ns

hold time
~2.6ns

Quantum™

Waveform Data: bp5, ΔA = 0.8

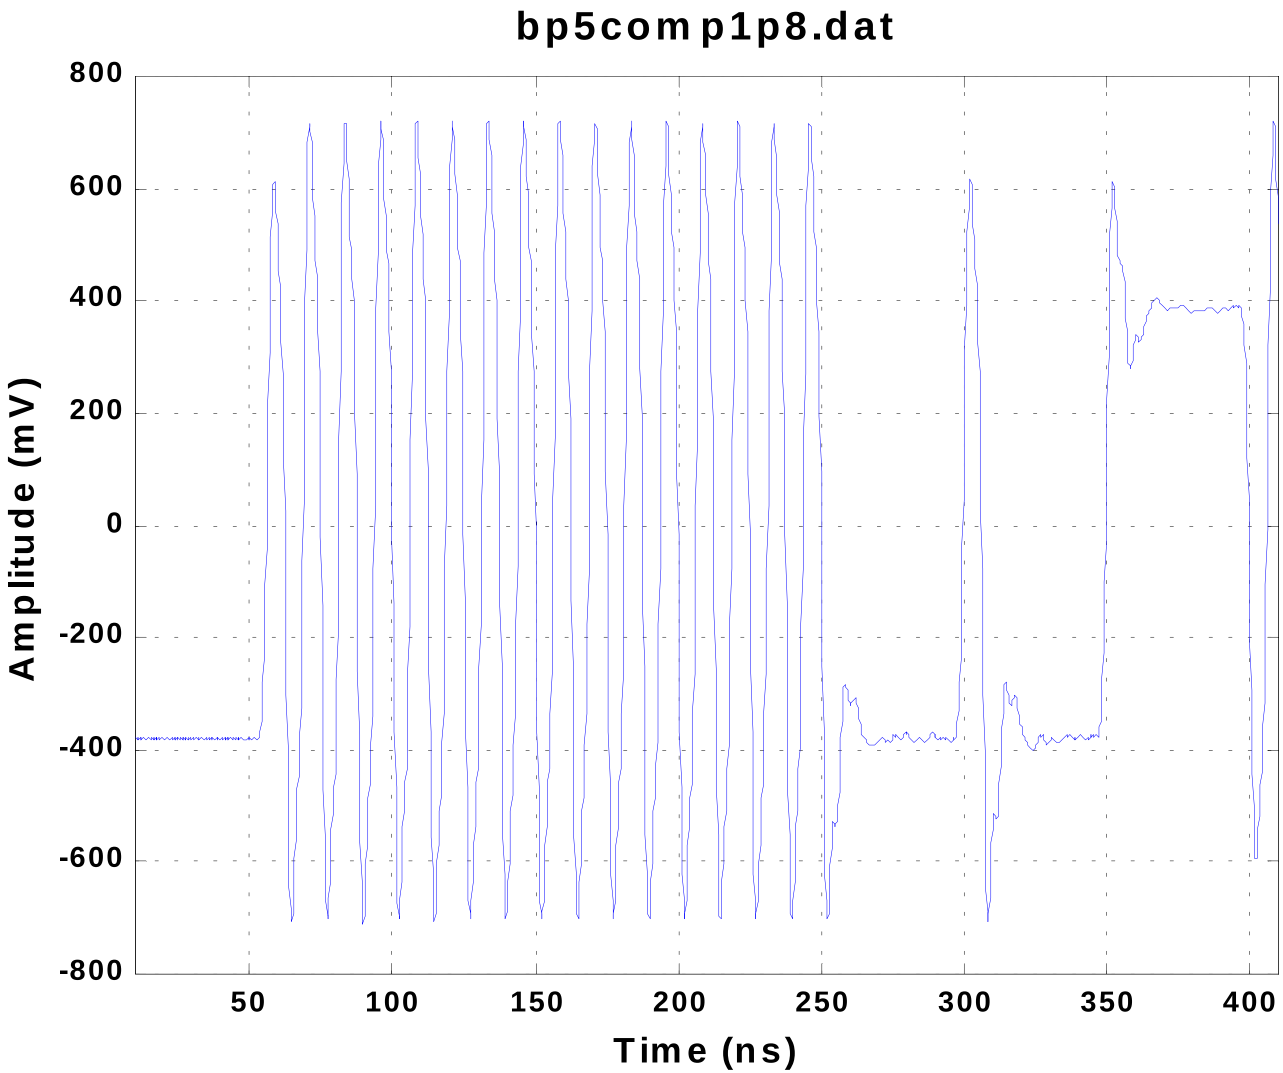

Eye Diagram: bp5, ΔA = 0.8

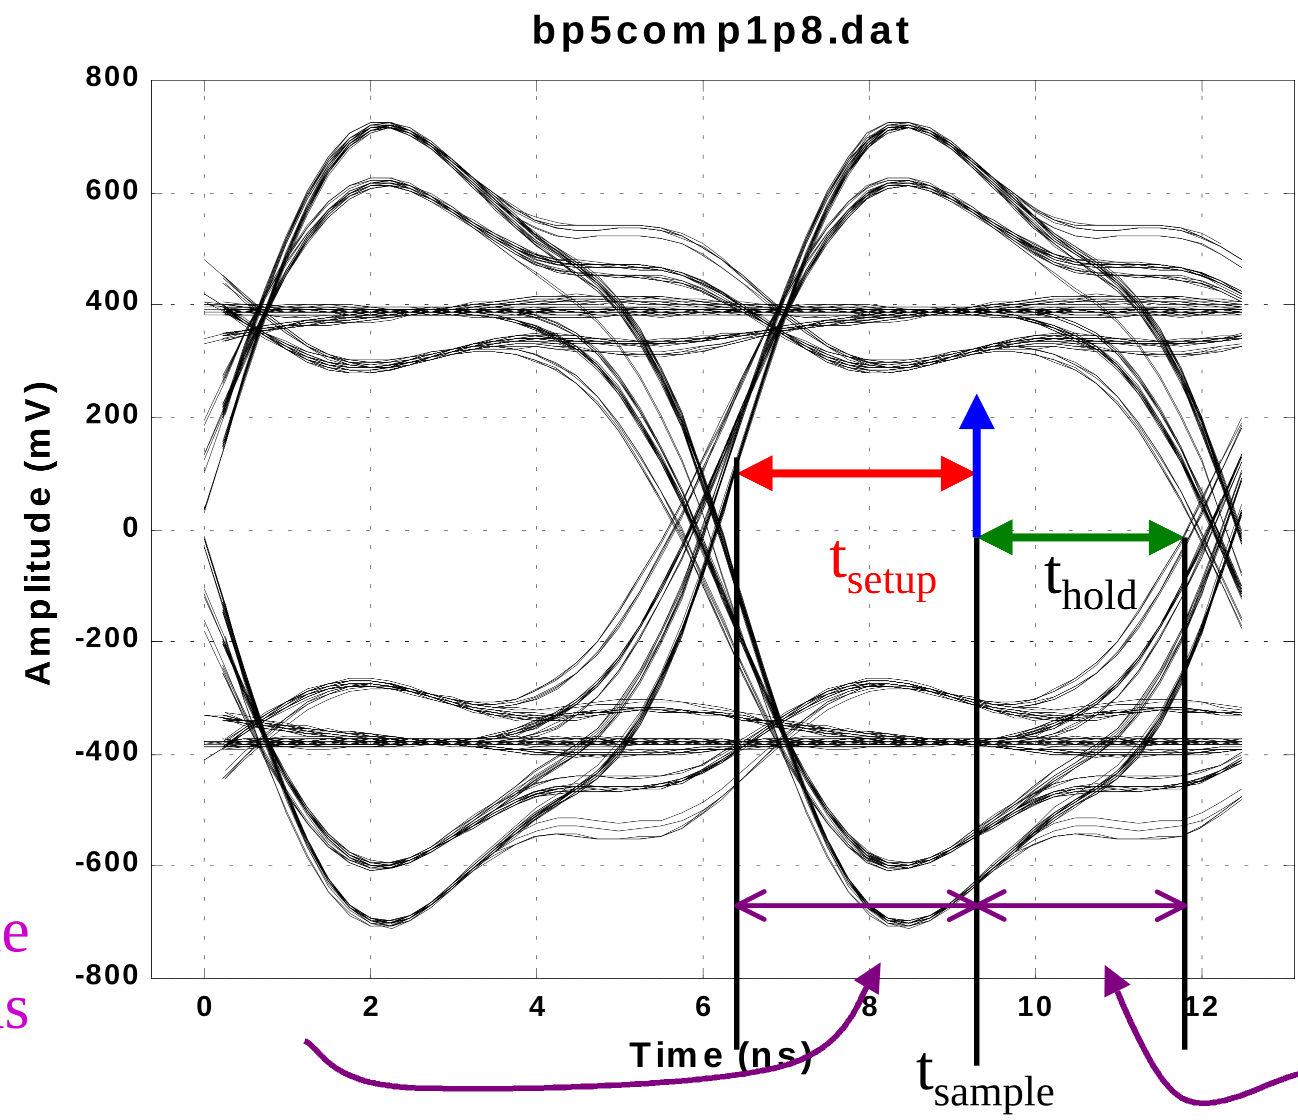

setup time
~3.0ns

hold time
~2.6ns

Experimental Setup #2 - Long Cable

- Hitachi 10 meter, 32AWG twisted-flat ribbon cable, 25cm load spacing, plus 6-slot back-plane.
- Waveforms captured @ 4Gs/s:
  - no pre-comp: ΔA = 0.0
  - amplitude pre-comp: ΔA = 0.2, 0.5, 0.8

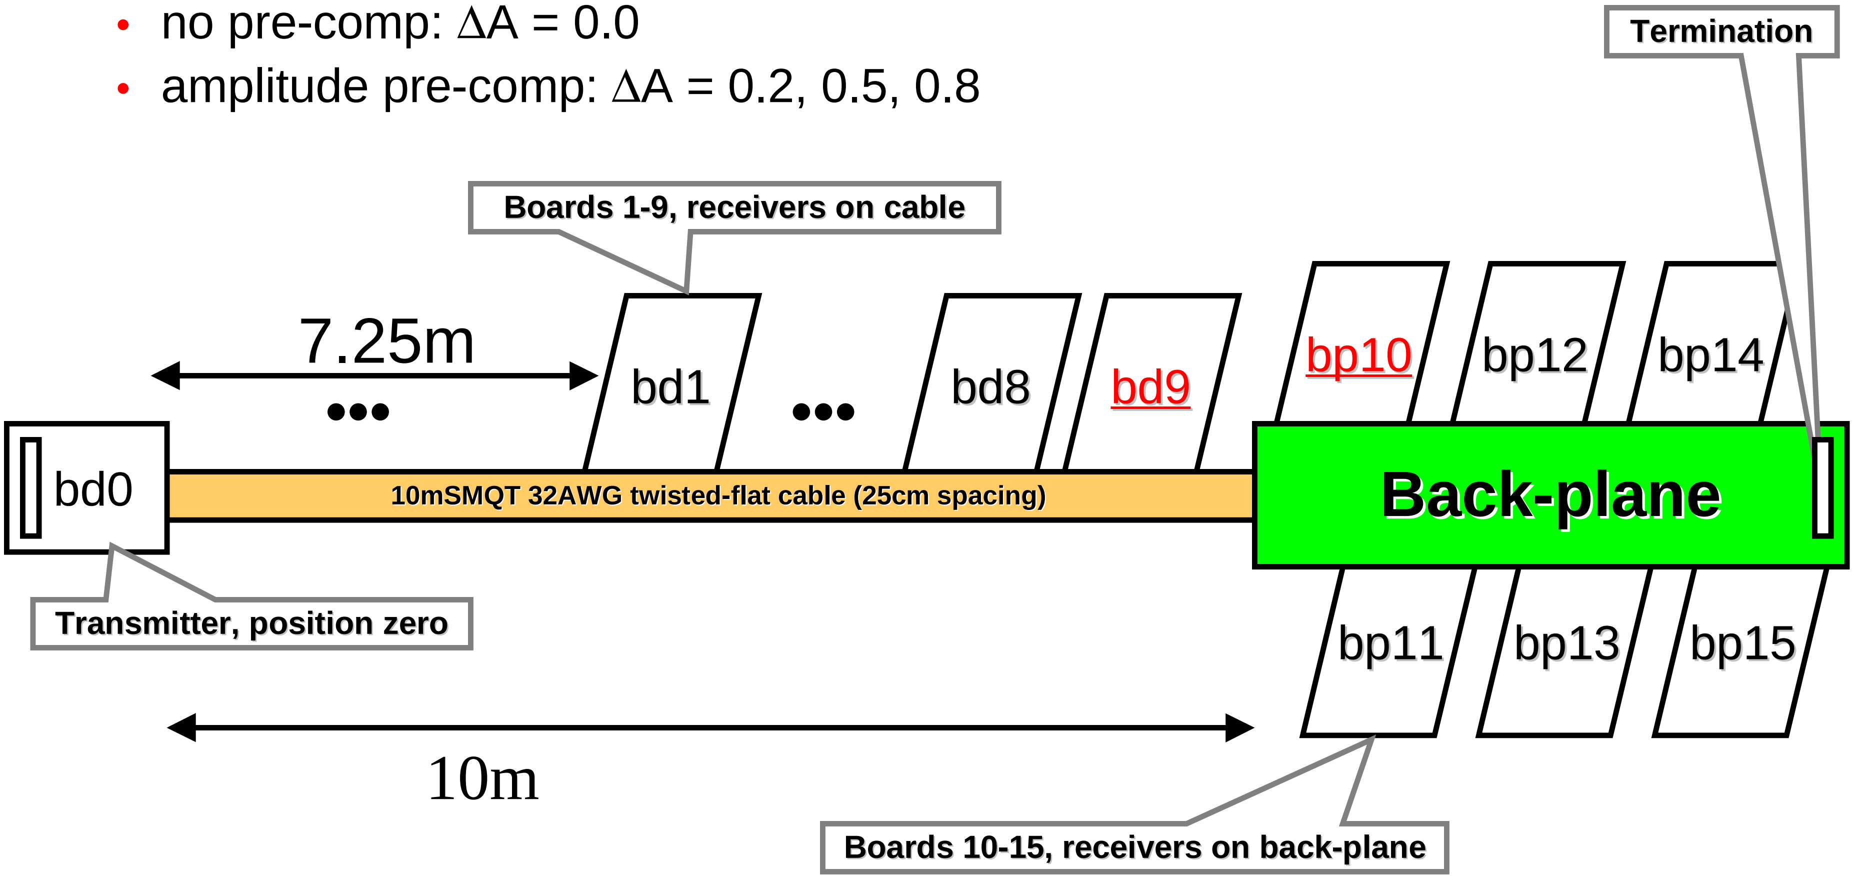

Electrical Setup #2 - Long Cable

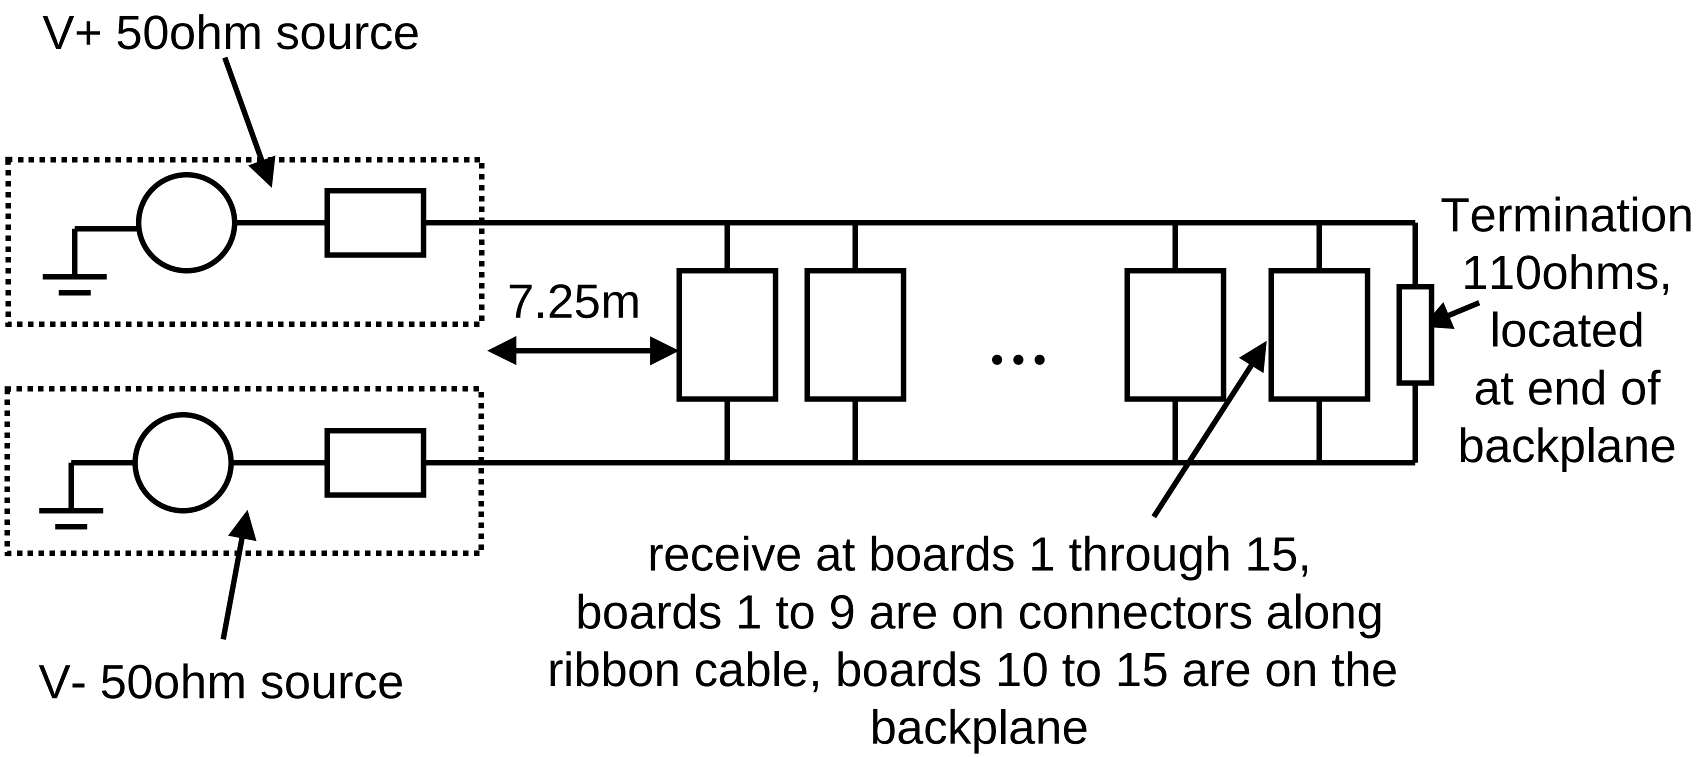

Waveform data: bd9, ΔA = 0.0

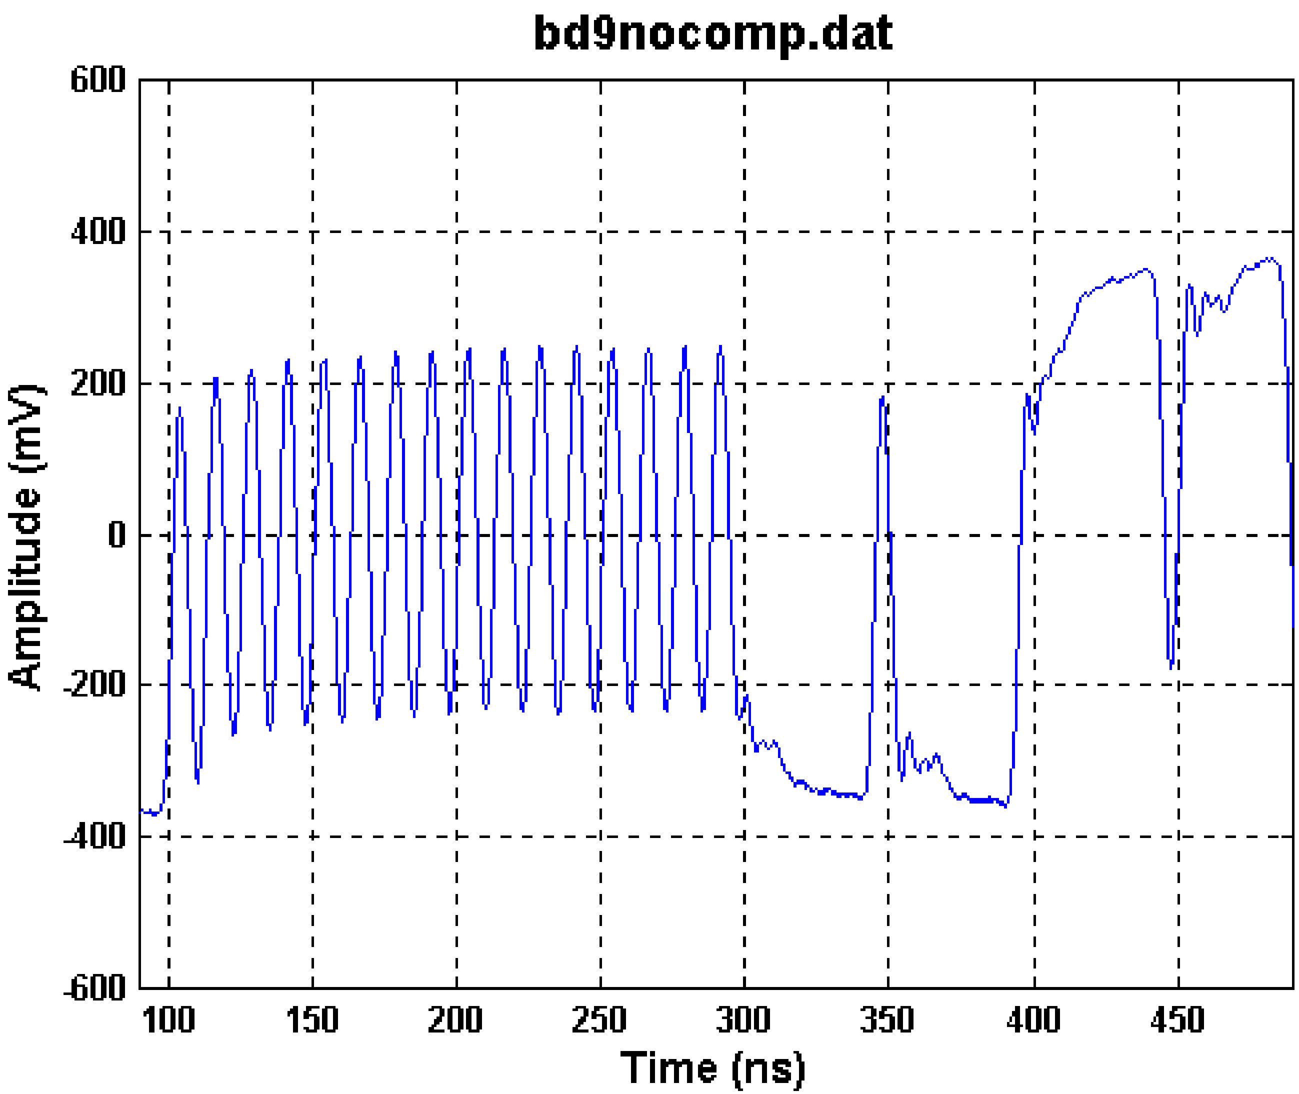

Eye Diagram: bd9, ΔA = 0.0

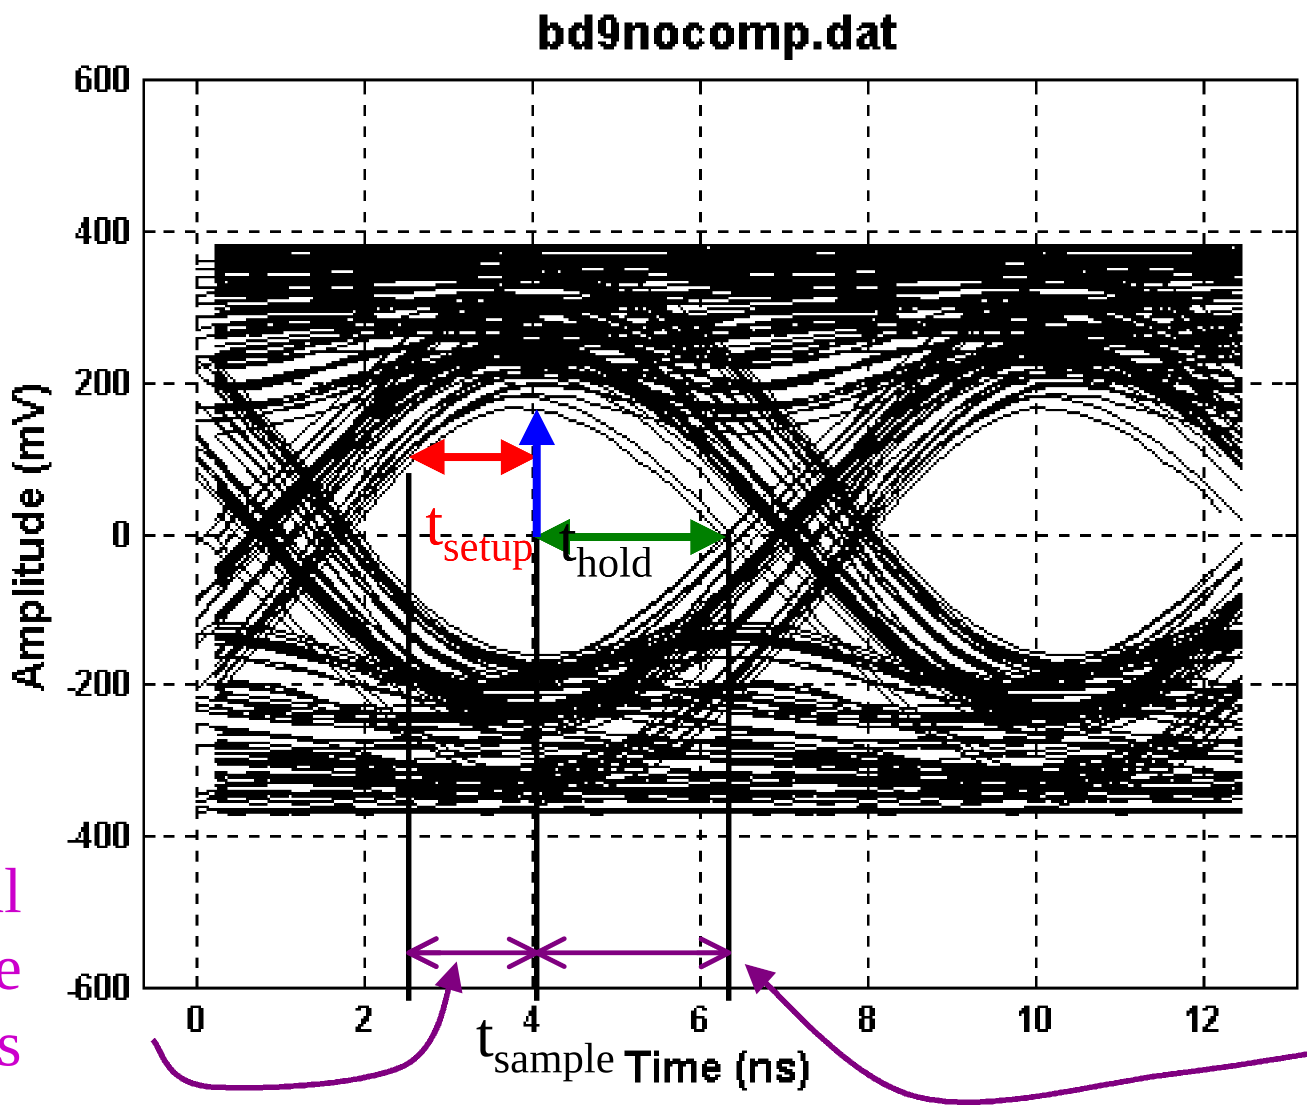

Nominal setup time ~1.6ns

Nominal hold time ~2.4ns

Waveform data: bd9, ΔA = 0.2

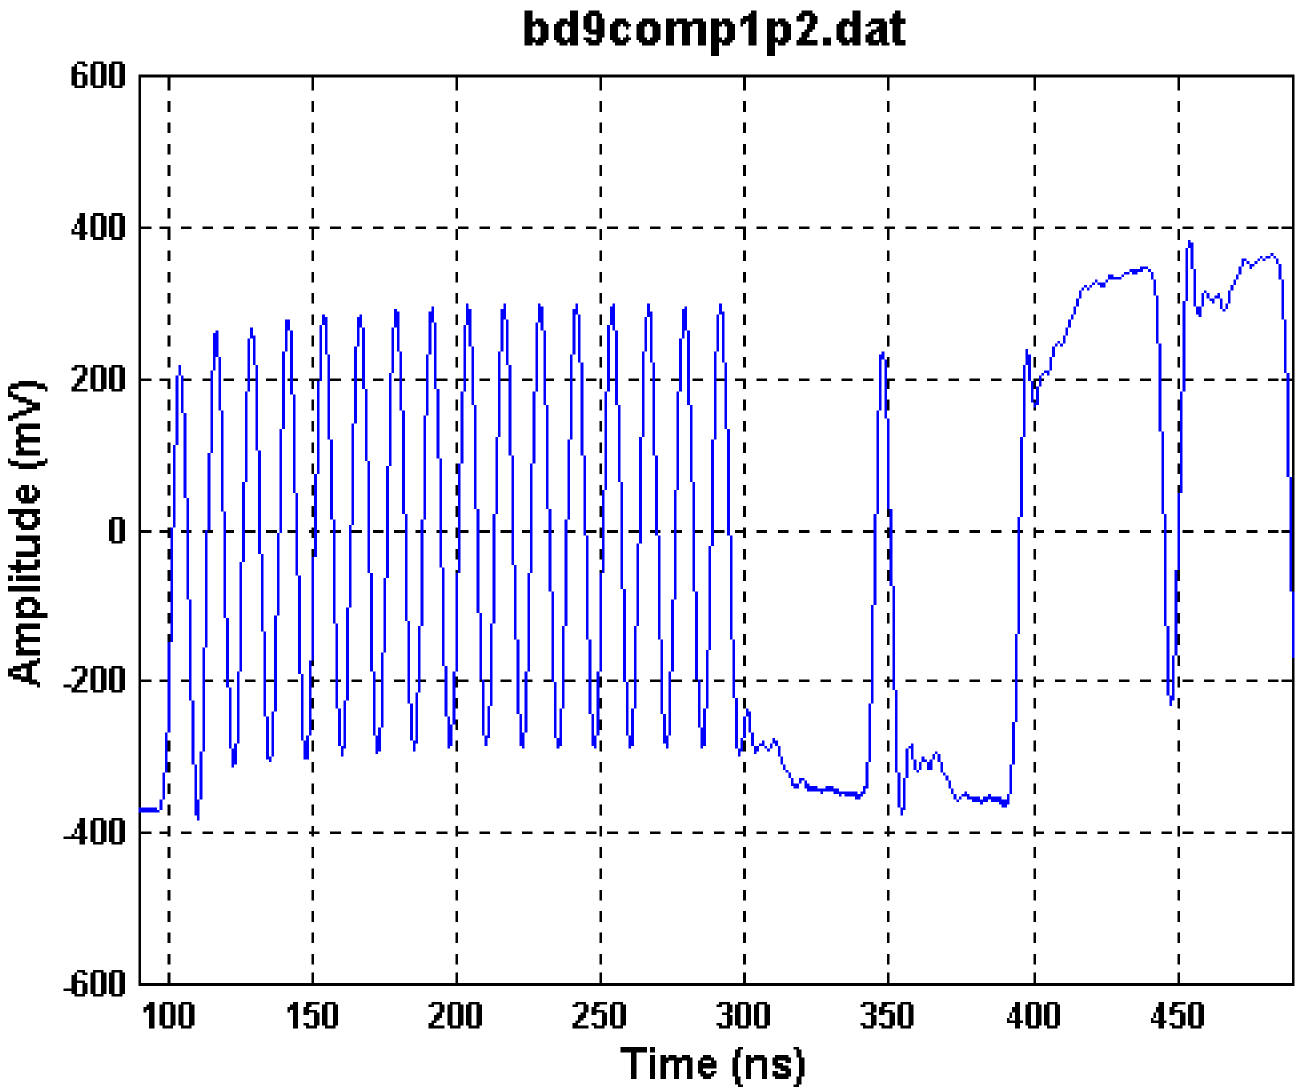

Eye Diagram: bd9, ΔA = 0.2

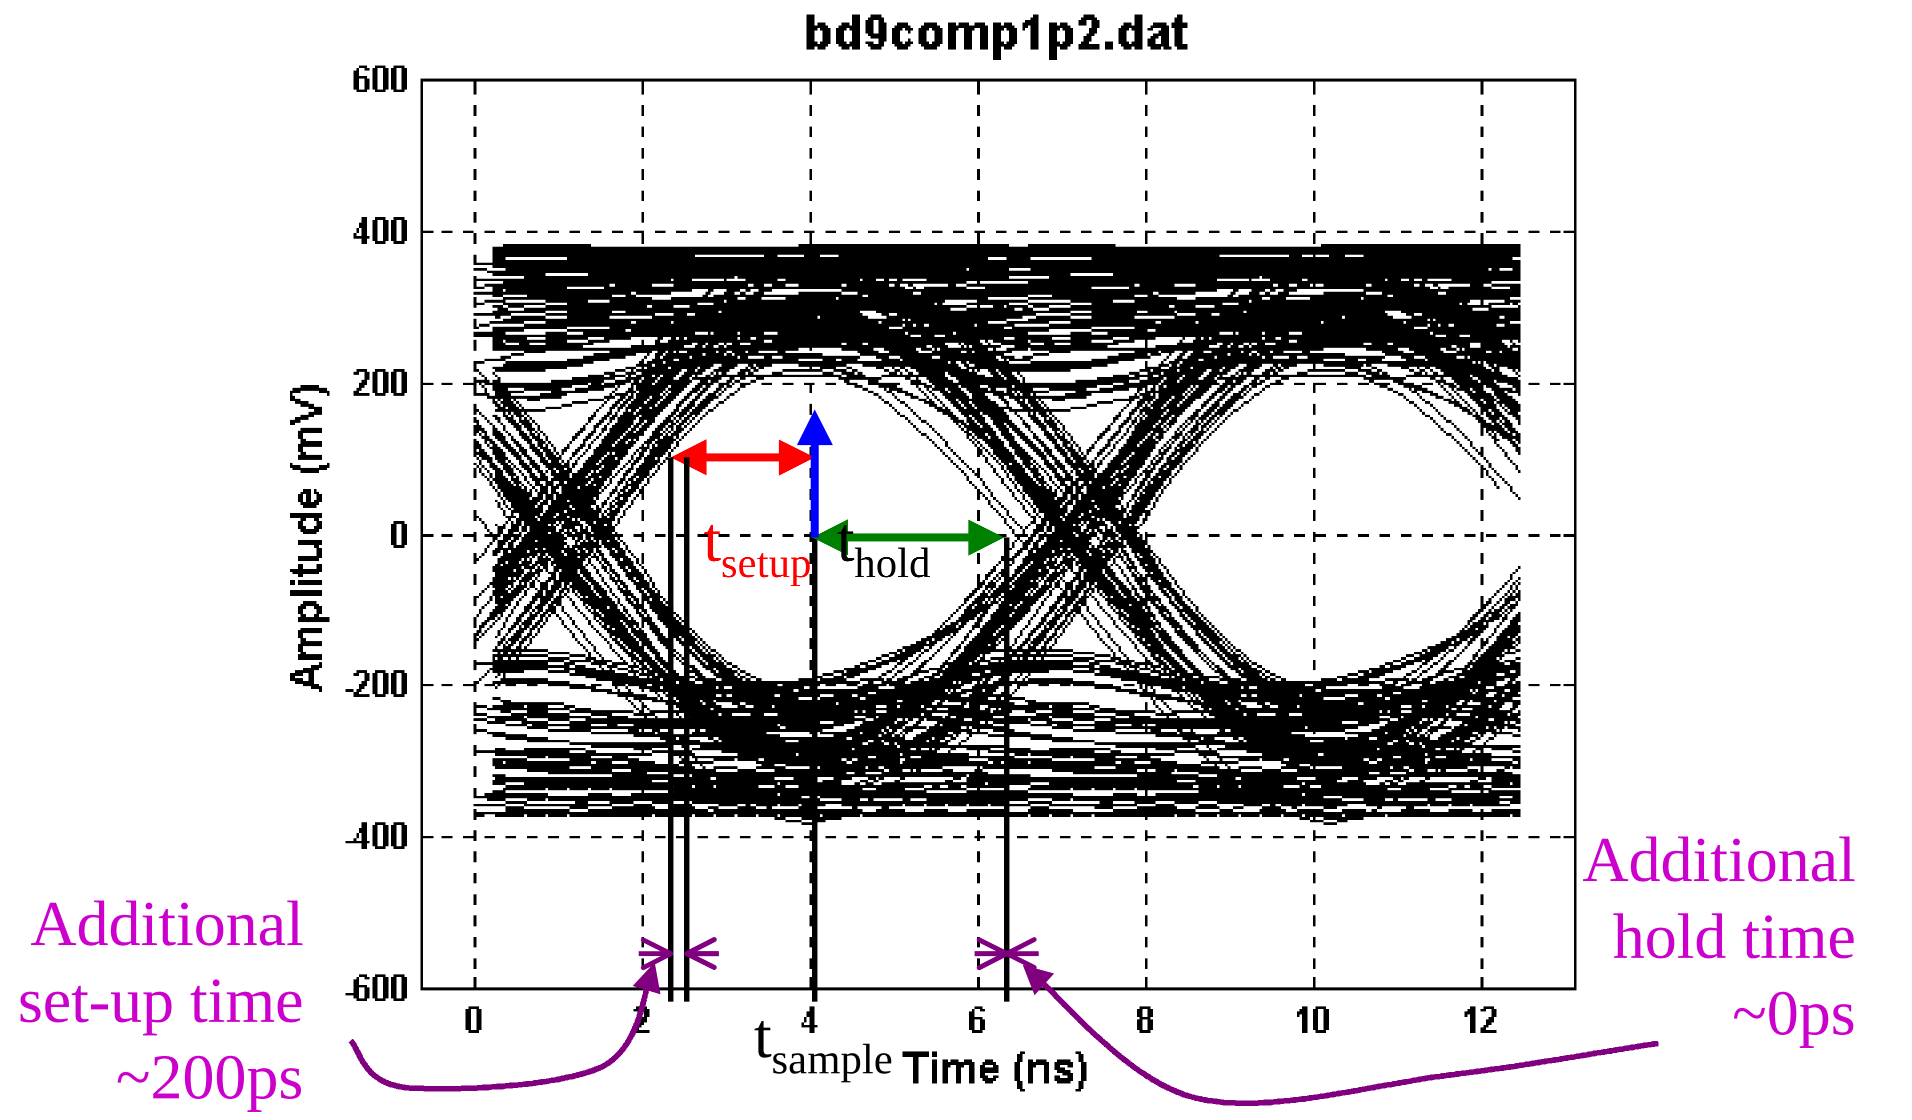

Waveform data: bd9, ΔA = 0.5

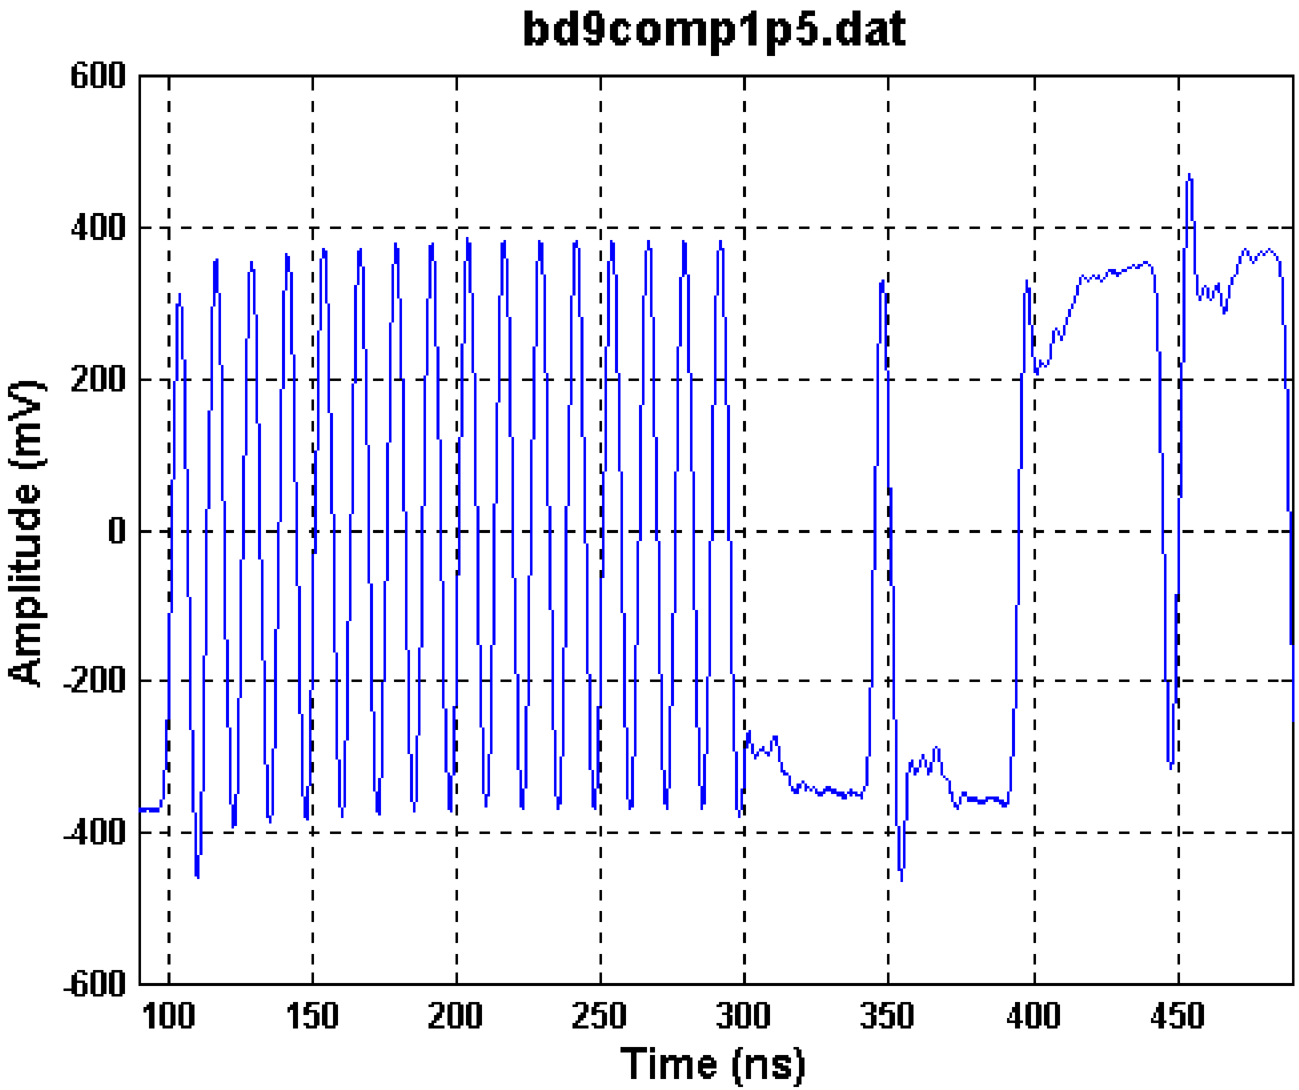

Quantum™

Eye Diagram: bd9, ΔA = 0.5

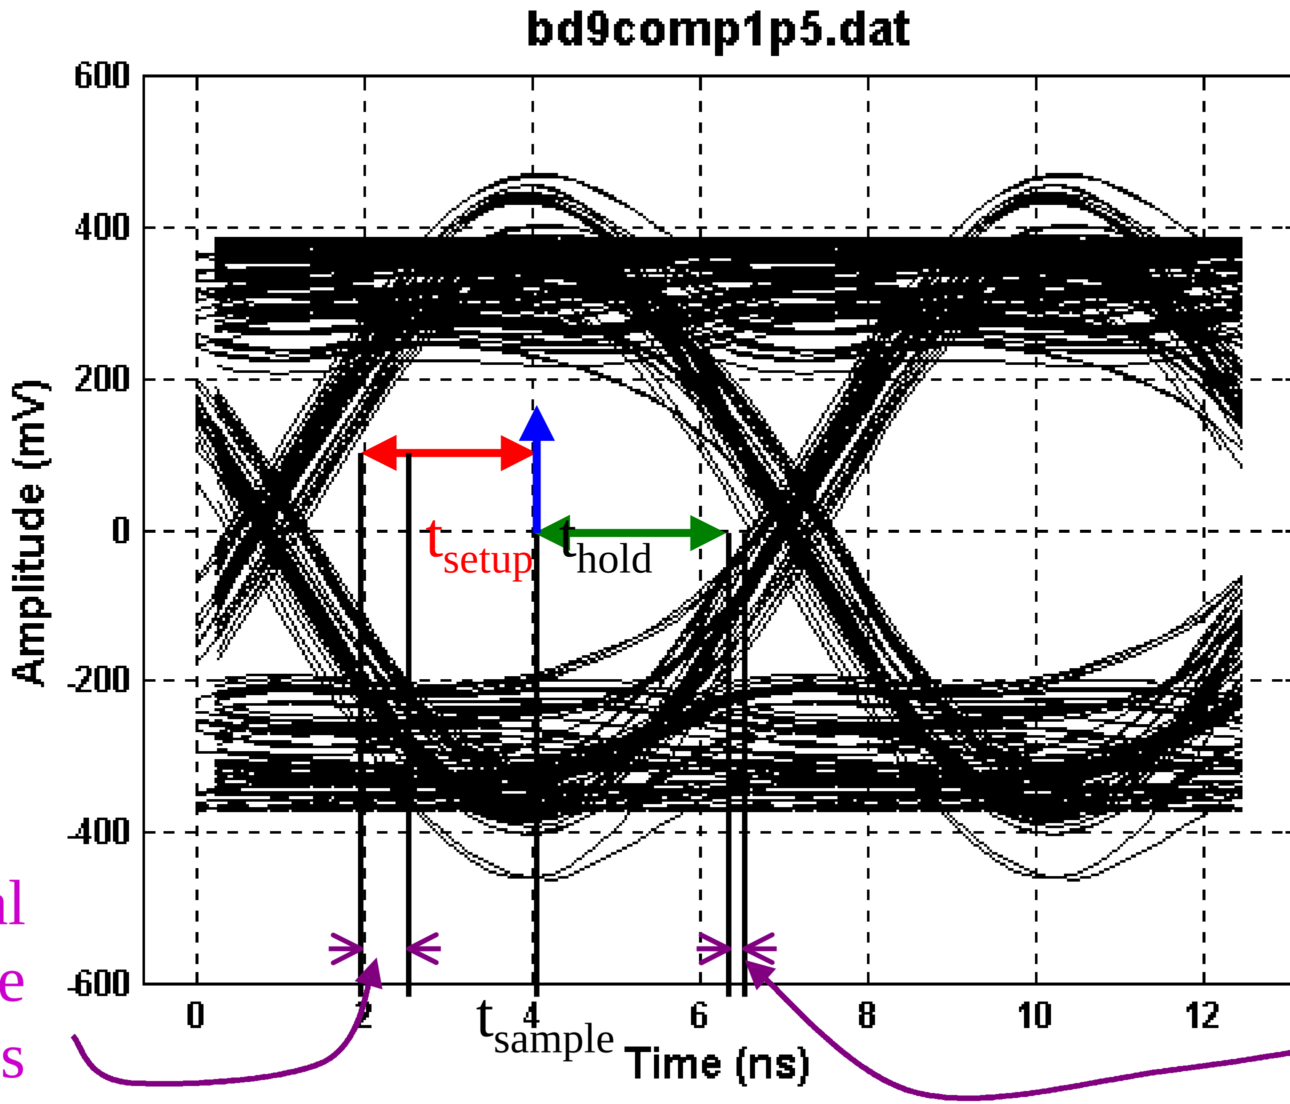

Additional set-up time ~600ps

Additional hold time ~200ps

Waveform data: bd9, ΔA = 0.8

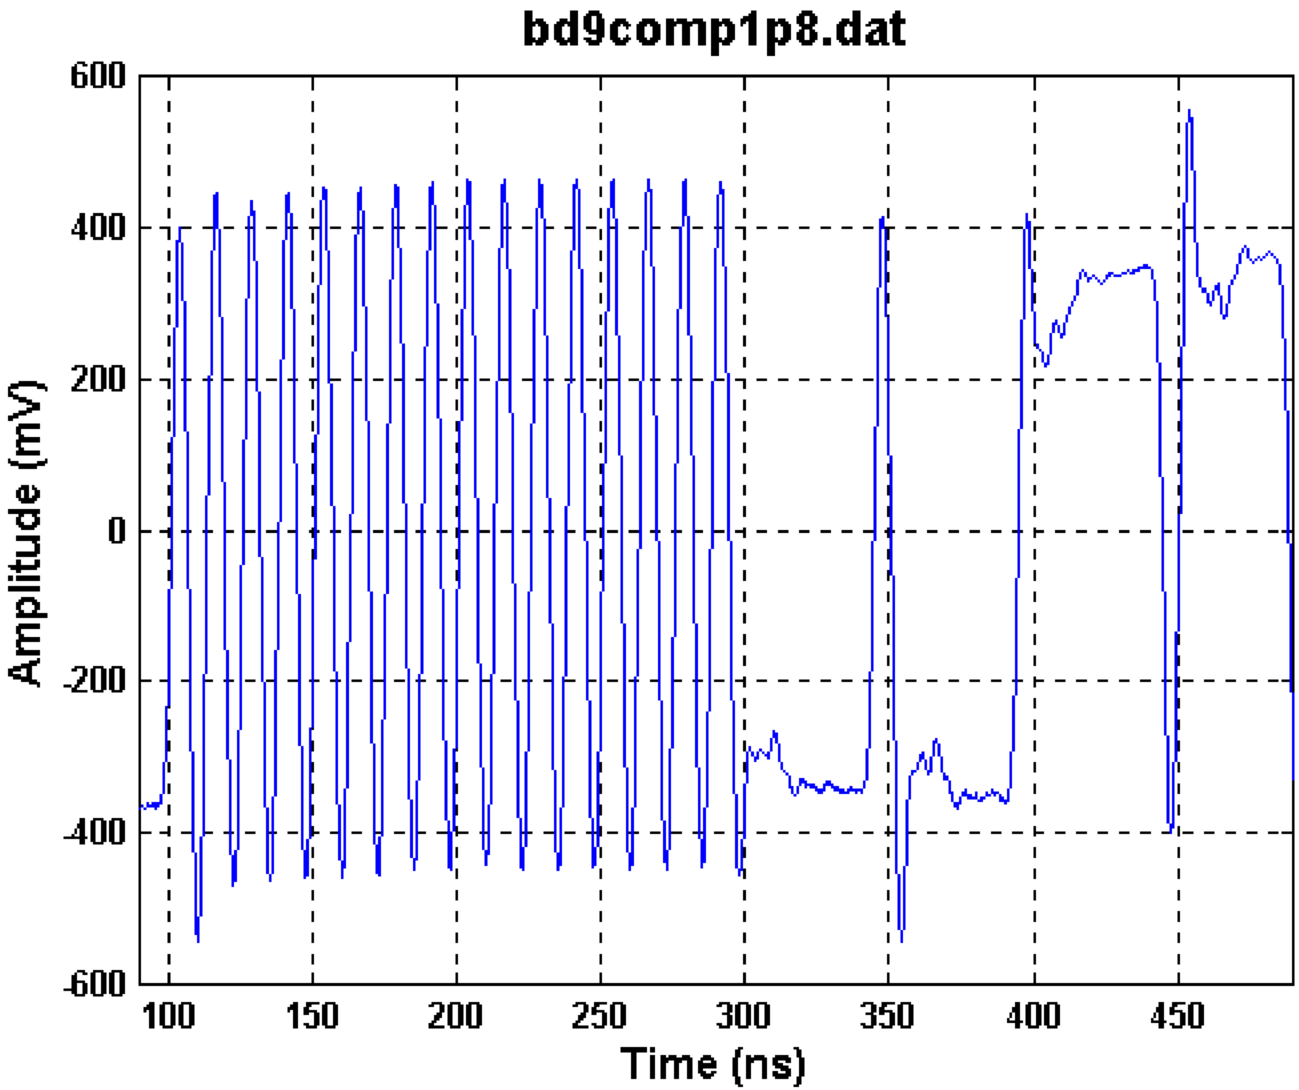

Eye Diagram: bd9, ΔA = 0.8

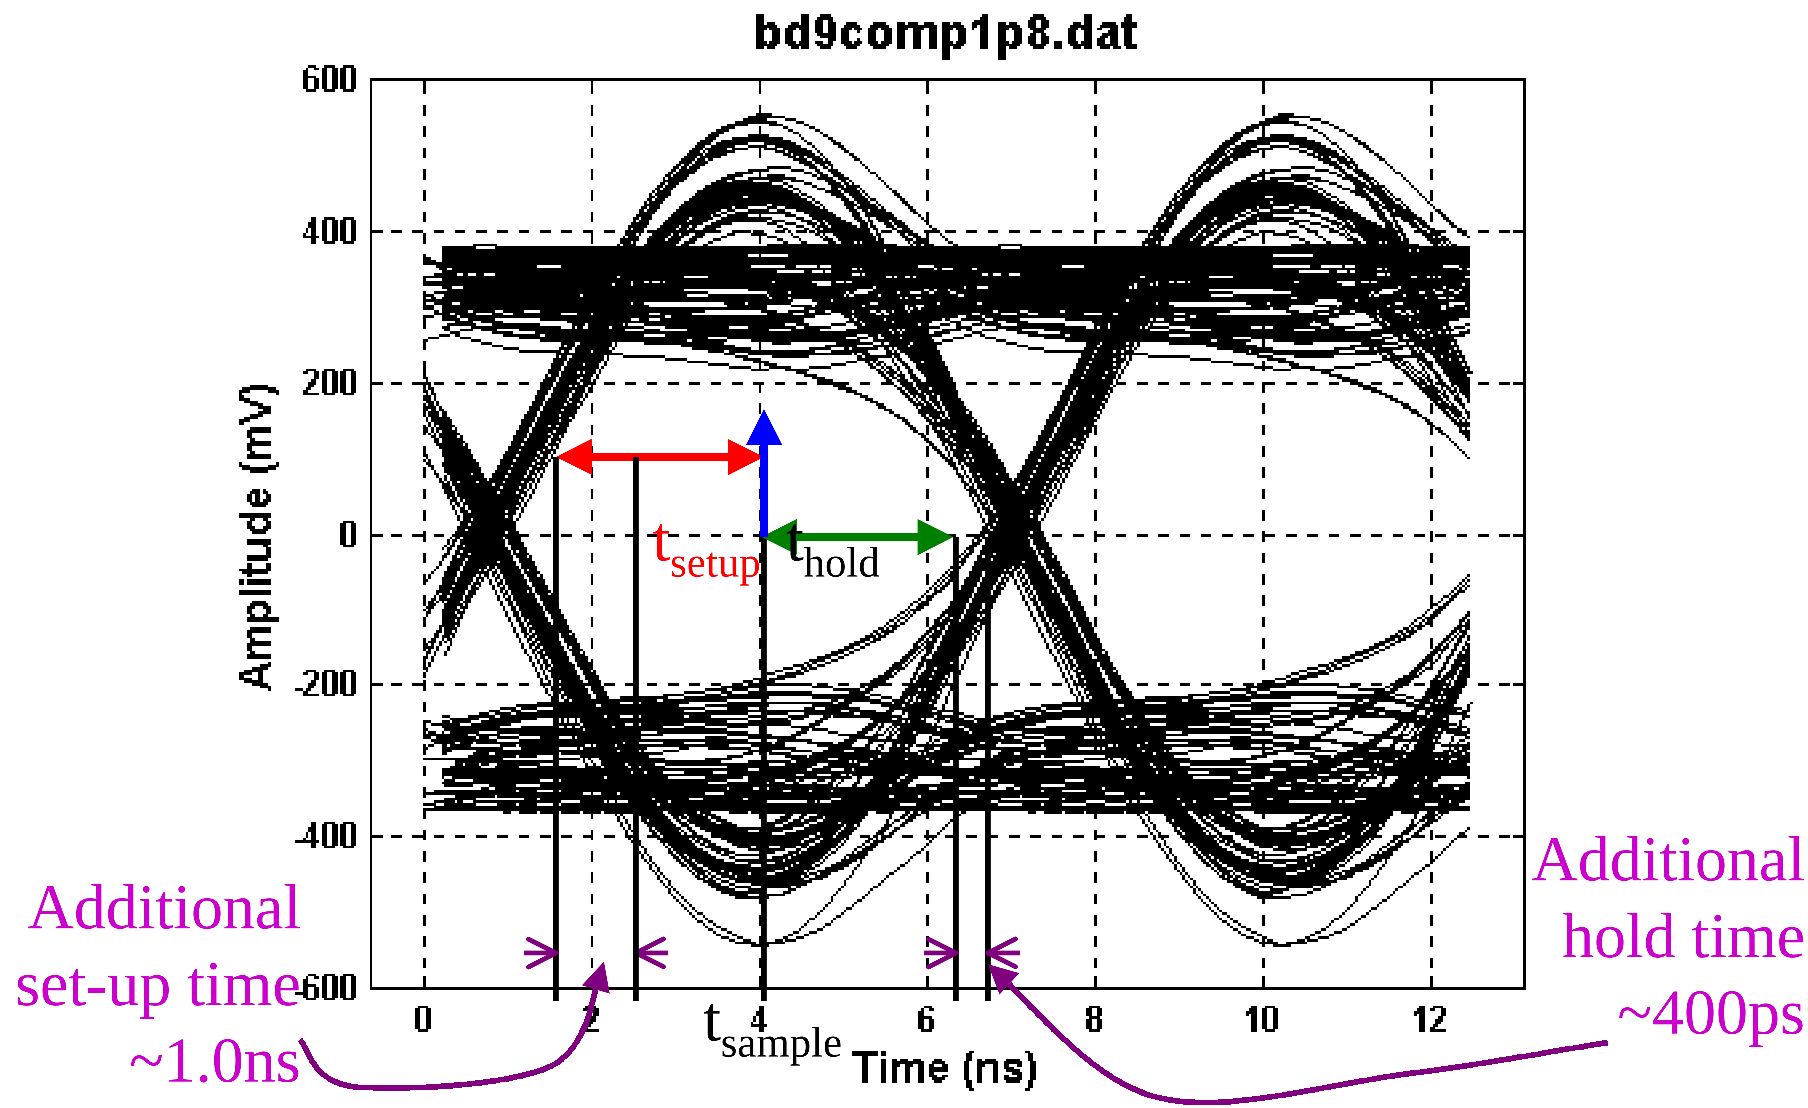

Waveform data: bp10, ΔA = 0.0

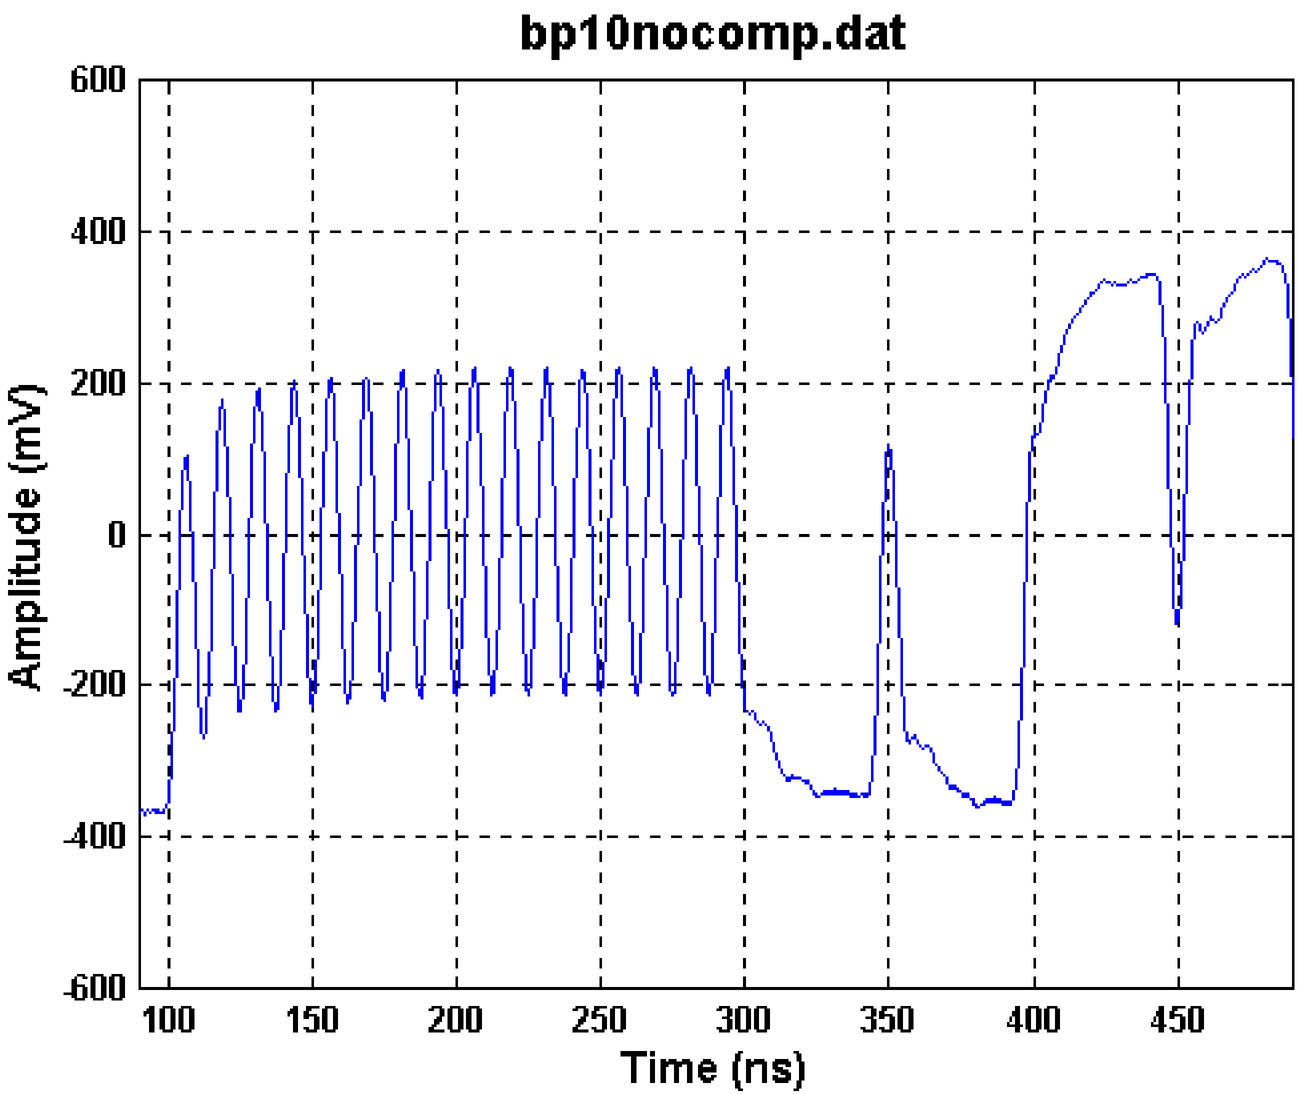

Eye Diagram: bp10, ΔA = 0.0

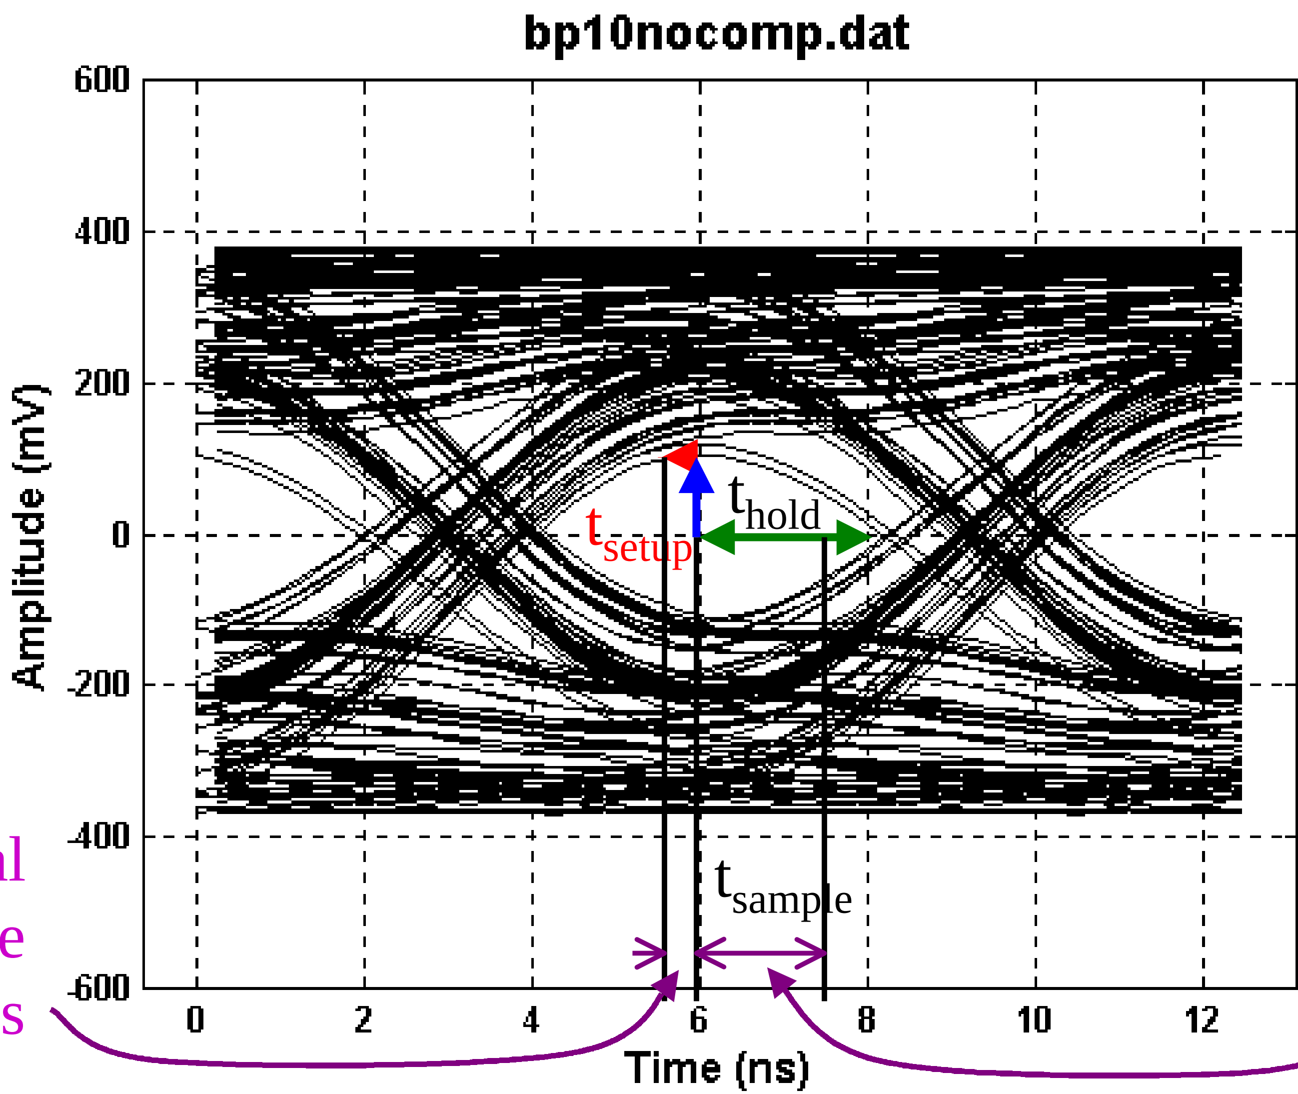

Nominal set-up time ~400ps

Nominal hold time ~1.6s

Waveform data: bp10, ΔA = 0.2

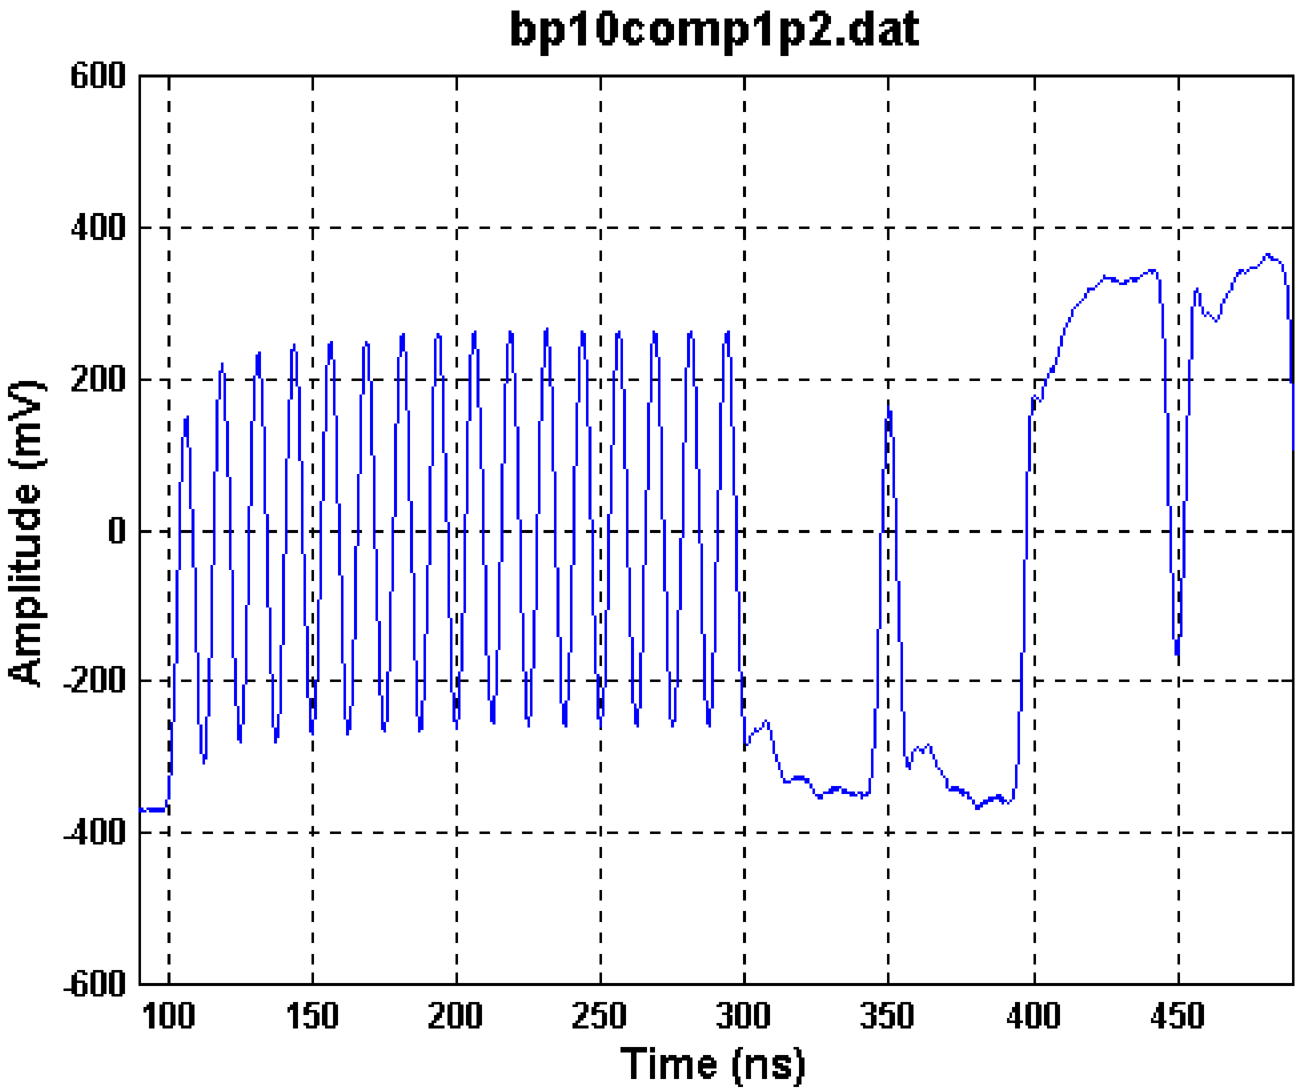

Eye Diagram: bp10, ΔA = 0.2

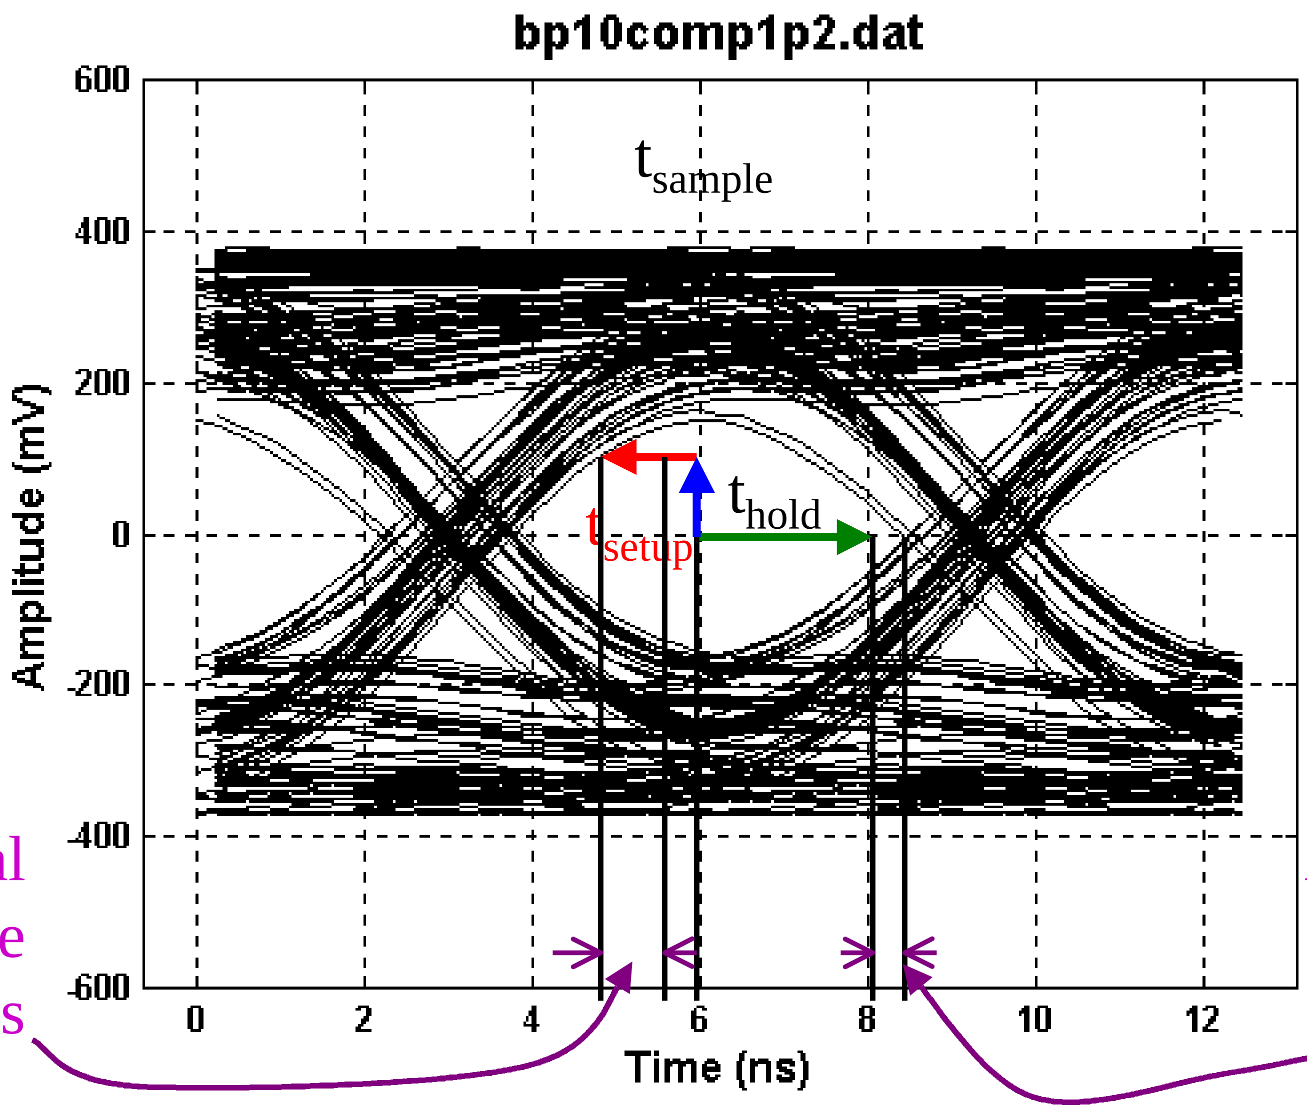

Additional set-up time ~800ps

Additional hold time ~400ps

Waveform data: bp10, ΔA = 0.5

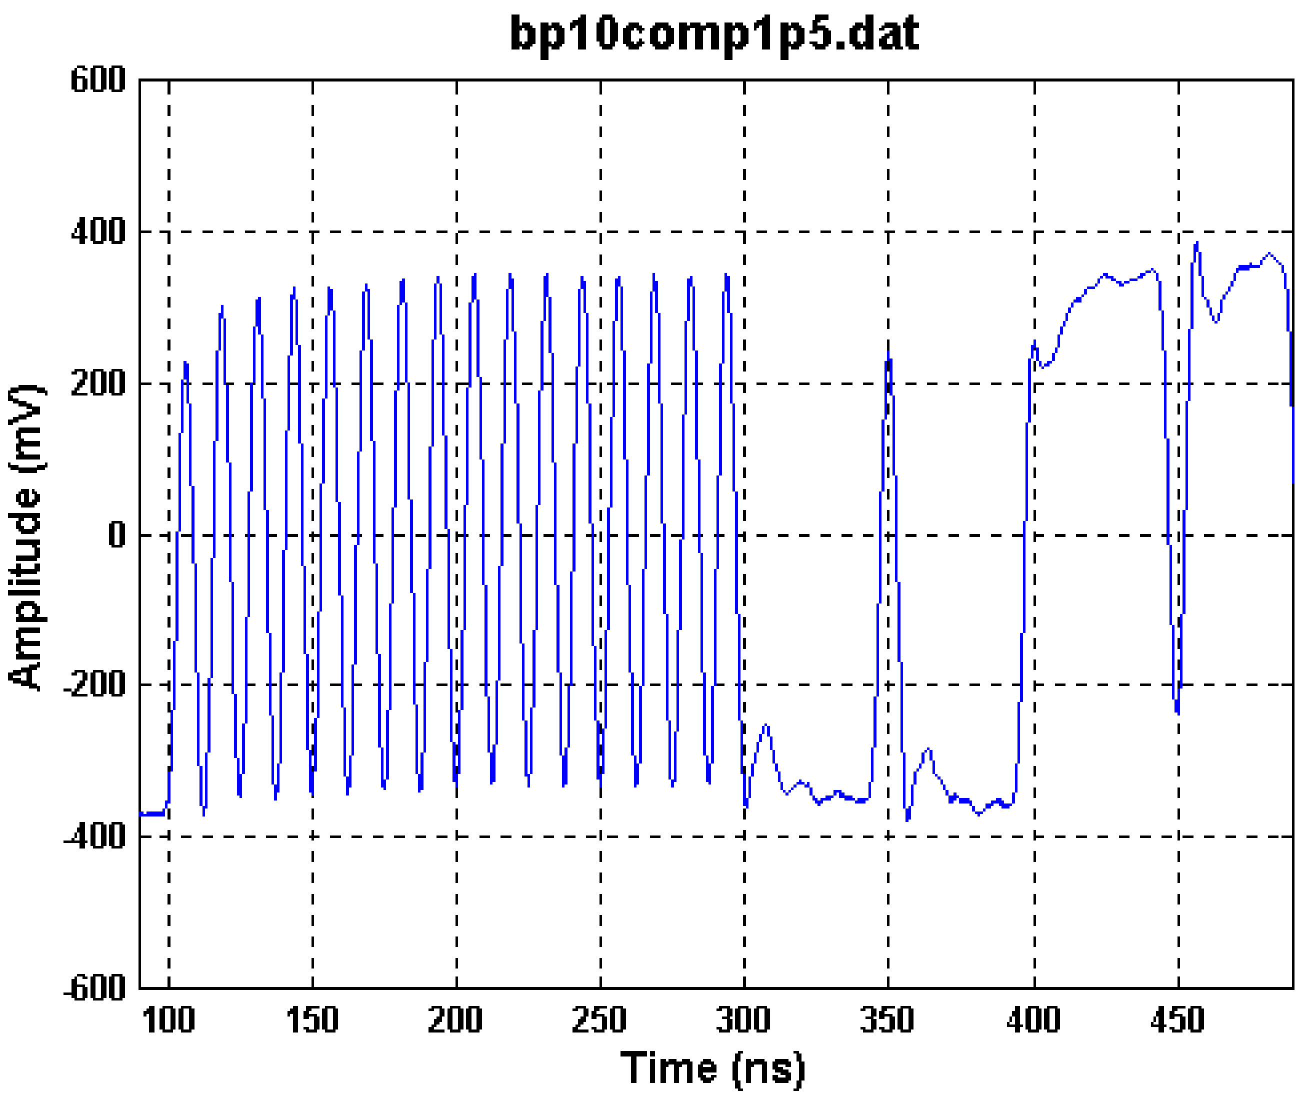

Eye Diagram: bp10, ΔA = 0.5

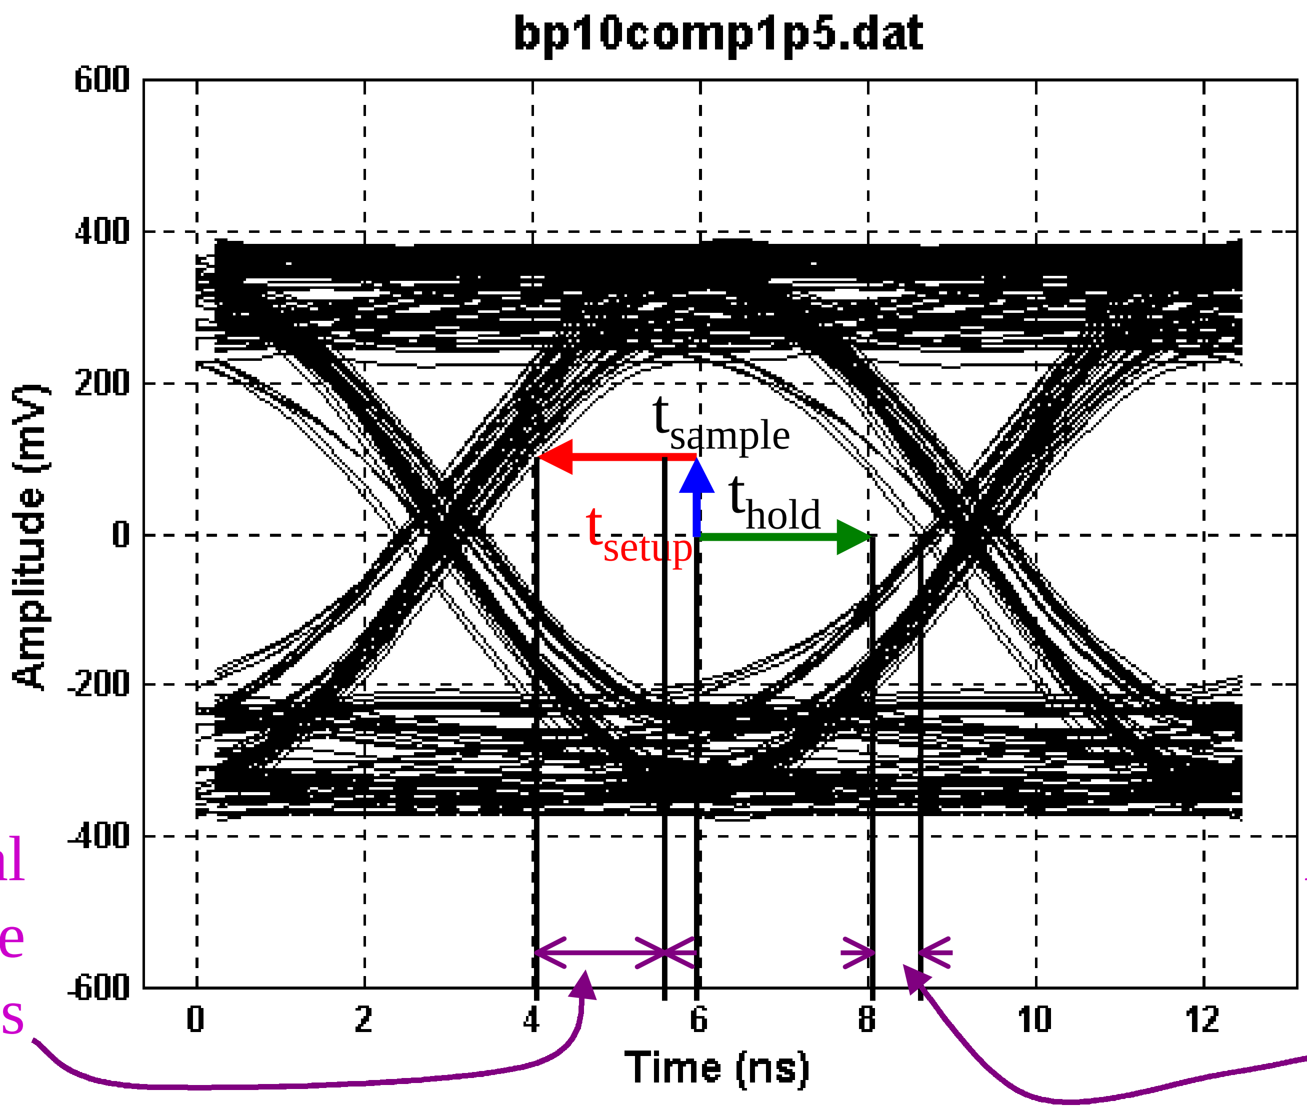

Additional set-up time ~1.6ns

Additional hold time ~600ps

Waveform data: bp10, ΔA = 0.8

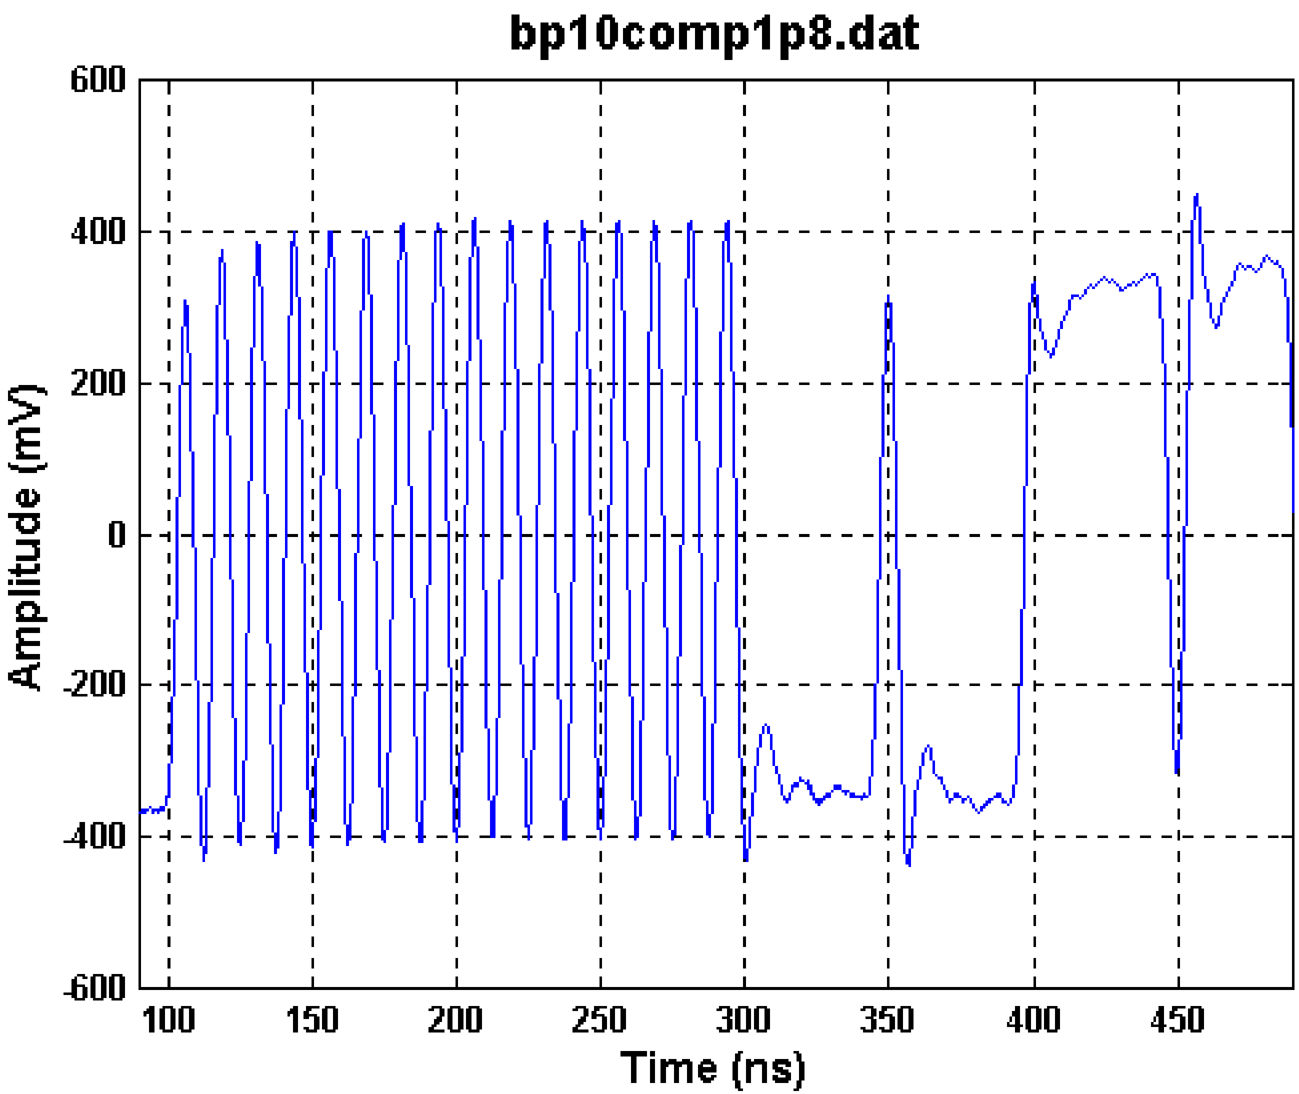

Eye Diagram: bp10, ΔA = 0.8

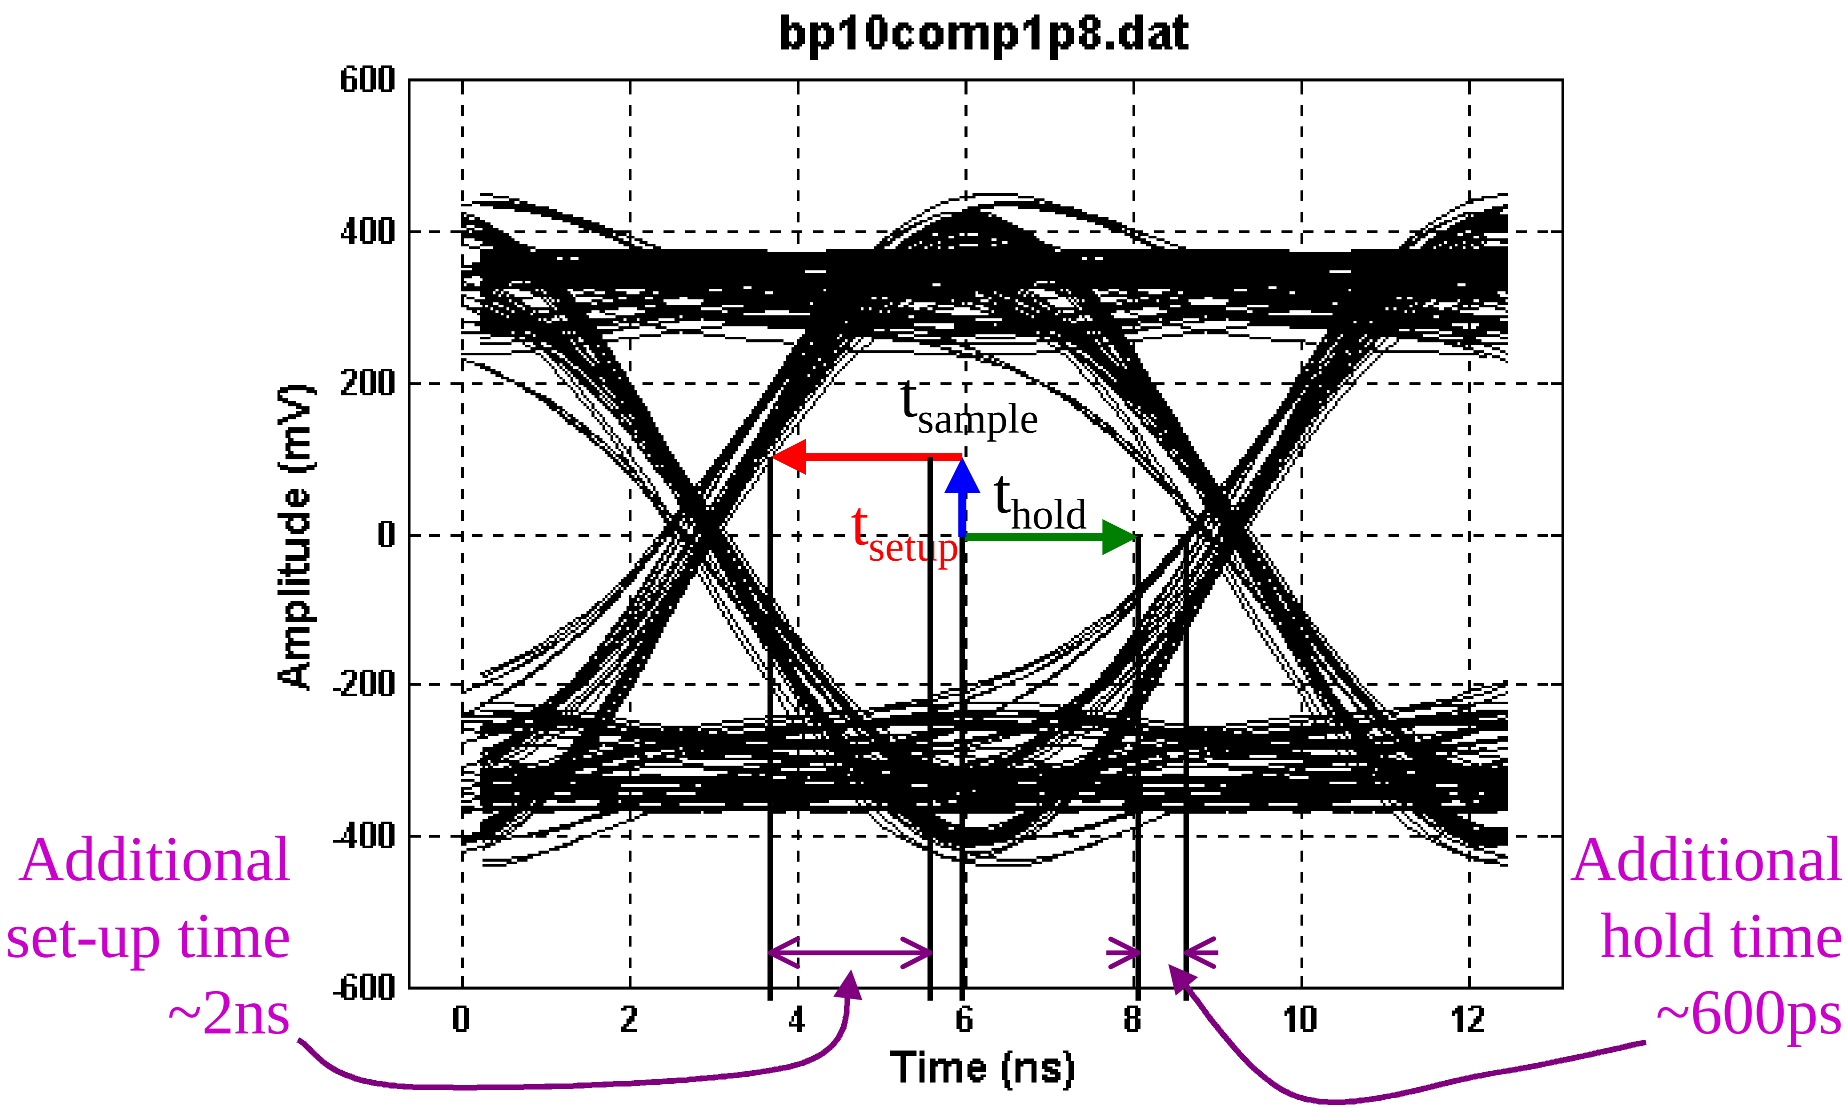

Amplitude Pre-comp Vs Setup Time

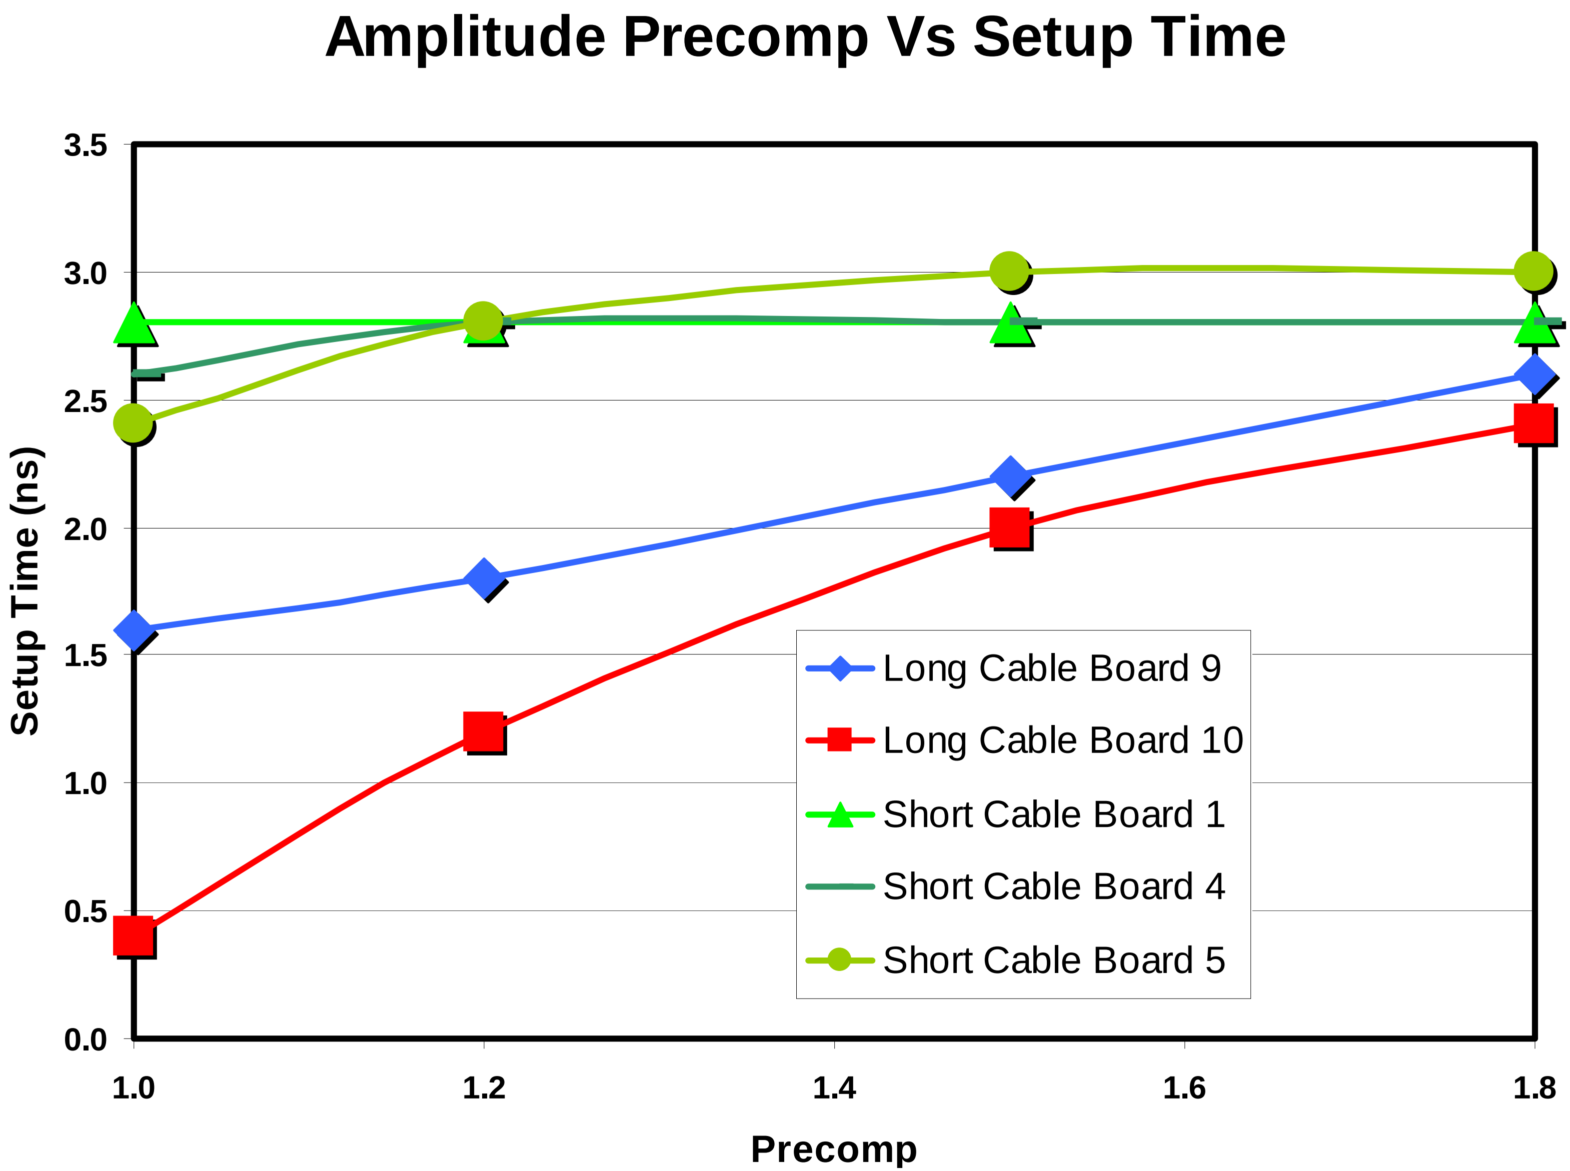

Amplitude Pre-comp Vs Hold Time

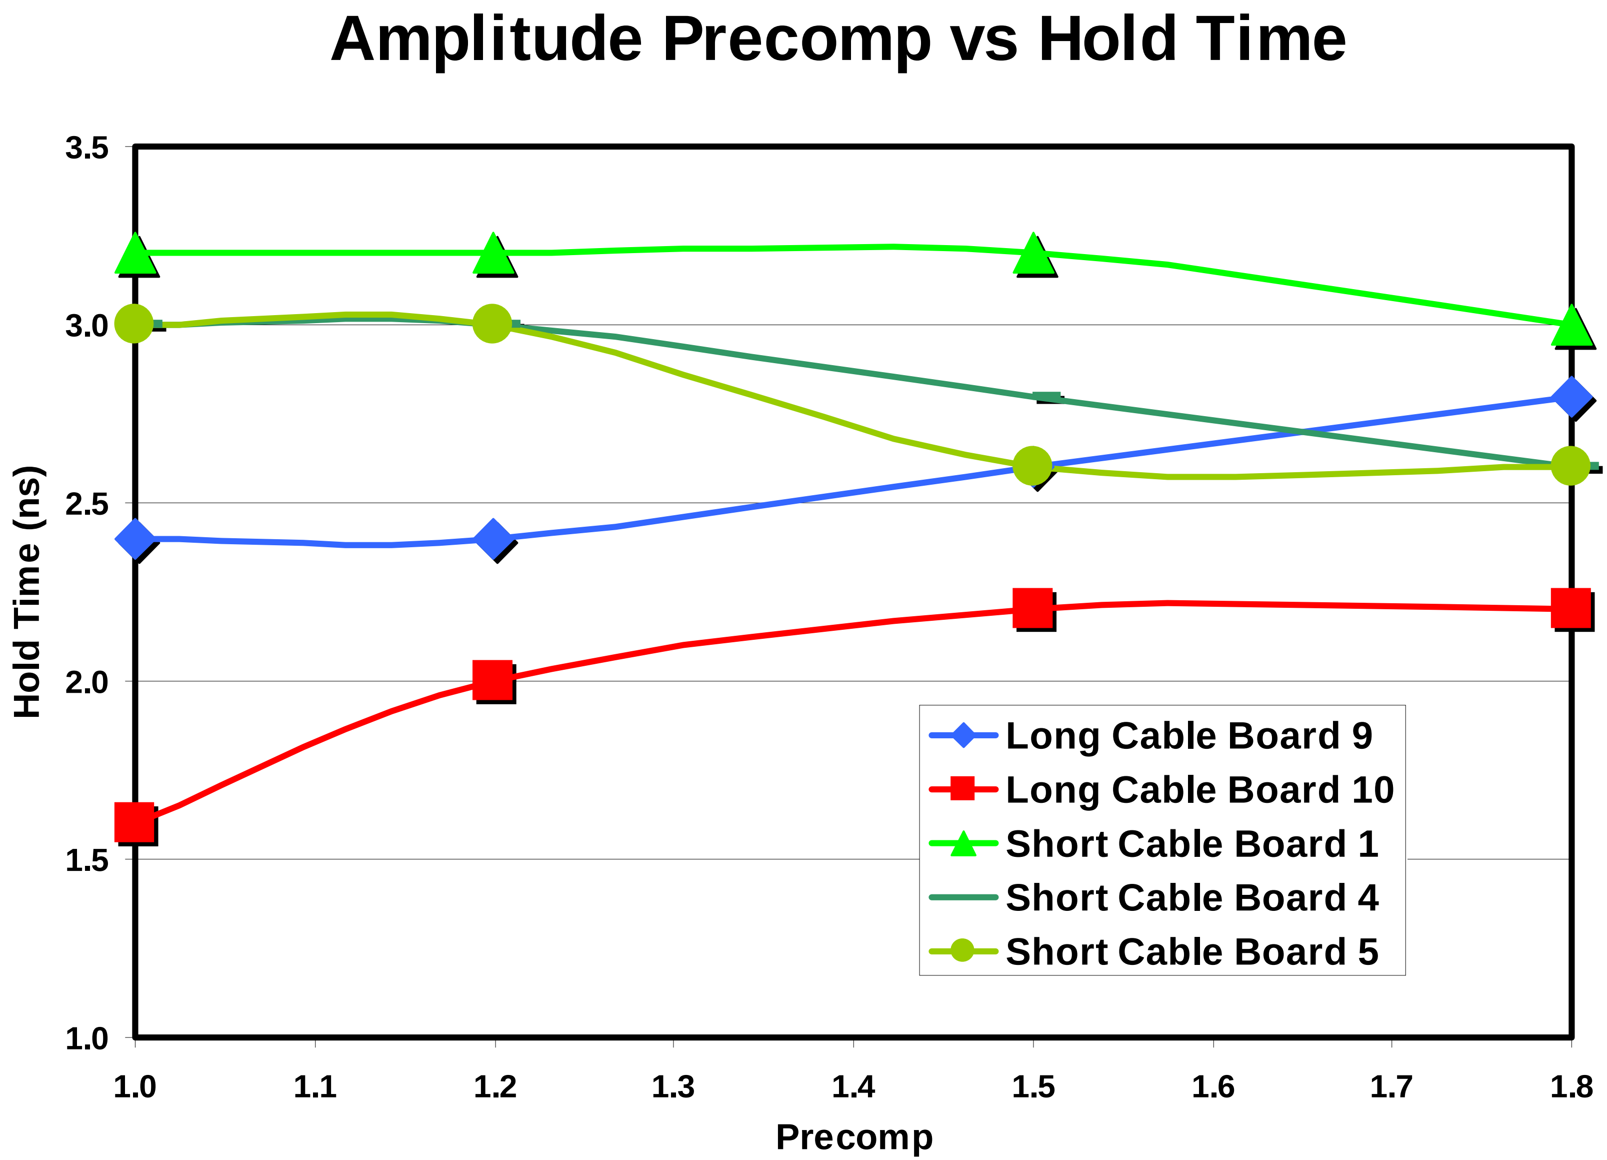

Quantum™

LSI SPI-4 Timing Budget: 99-261r1

Timing Budget (ns)	**.**
Basic Period	**12.50**
DT-Period	**6.250**
Period Tolerance	**0.700**
Deterministic Errors	**.**
Silicon TX Driver Skew	**1.000**
Package Skew (Initiator)	**0.065**
PCB Layout Skew (Initiator)	**0.200**
Cable Skew (@25ps/ft)	**2.500**
PCB Layout Skew (Target)	**0.200**
Package Skew (Target)	**0.065**
Silicon Rx Routing Skew	**1.000**
HL vs LH Matching	**0.500**
Non-Deterministic Errors	**.**
Low Vt vs Substrate Noise	**0.200**
PLL Jitter	**0.250**
Cross Talk Induced Jitter	**0.500**
Cable Period Distortion ISI	**2.000**
Input Slew Rate Dependent Skew	**0.200**
Receiver Amplitude Dependant Skew	**0.200**
Self Cal Accuracy (± 100ps)	**0.200**
Data Valid Window	**1.850**
Data Setup/Hold	**0.925**

Conclusions on Amplitude Pre-Comp

- Amplitude pre-comp appears to work for long cables and for closely-spaced loads (i.e., set-up time improves, amplitude margin improves).
  - Without it (or something else) certain configurations will fail: zero set-up time, no amplitude margin.
  - Optimum value for pre-comp amplitude is between 1.5 and 1.8.
  - This value depends on the configuration.
- Amplitude pre-comp does not improve signal integrity for short cables:
  - Extra signal amplitude contributes to ringing and overshoot;
  - Amplitude pre-comp can even slightly decrease available hold time.
- Unsolved issues with EMI, Power, cross-talk, and capacitance.

Quantum™ Comments on Amplitude Pre-Comp

- Data presented is for various pre-comp boost values over a 'low frequency' drive level of 400mV peak differential
- Driver and terminators are ideal
  - Other errors due to driver/terminator mismatch and interconnect resistance, requires an addition of approximately 70mV in the receiver margin above the 100mV in Figs 47,48 of SPI-3
- Only 1 line is driven, no cross-talk components are included
- Setup & hold measurements must include a timing error for residual skew.
- Transmit pre-comp appears able to meet the present interconnect configuration range, but:
  - fixed value of pre-comp is non-optimum at many drops
  - meeting the 12m / 15 drop configuration stresses available driver output signal capabilities
  - little or no range left to accommodate speed increase to 640

Quantum™

Comments on Amplitude Pre-Comp

- Driver Issues for Amplitude Pre-comp
  - Data shows that a Pre-comp boost of 1.5 to 1.8 is required to handle 12m cable / 15 load case
  - Expect bus DC errors similar to U160
  - Expect cross-talk and reflection problems with pre-comp boost to be worse than with U160
    - **There is no reason to expect that a minimum (un-boosted) transmit signal lower than the 360mV min in SPI-3 is workable**.

  - Taking the SPI-3 360mV minimum, and +/-20% driver tolerance → 400mV nominal un-boosted level
  - 1.8x boost → 720mV nominal boosted level
  - + 20% driver tolerance → 864mV peak differential boosted signal

    - ***These signal ranges exceed the present SPI-3 min/max limits and will present driver design difficulties and increased power dissipation.***

Next Steps

- We will collect more data on different configurations:
  - heavily and lightly loaded busses
  - typical and atypical
  - point-to-point
- We will investigate:
  - can lower amplitudes be used to address large chip power requirements?
  - how bad will common-mode degrade for large amplitudes?
  - would receiver compensation work?
  - how much capacitance will be added by larger drivers?
  - how much capacitance is acceptable?
  - could we use a different terminator scheme?