

Welcome to

3th ANNIVERSARY

DESIGNCON[®] 2025
WHERE THE CHIP MEETS THE BOARD

Conference

January 28–30, 2025
Santa Clara Convention Center

Expo

January 29–30, 2025

DESIGNCON[®] 2025
WHERE THE CHIP MEETS THE BOARD

Information Classification: General



JAN. 28–30, 2025

#DesignCon

1

 **informa markets**

Low latency speculative error correction using simplified ML detector for 64Gbps wireline transceiver

Speaker: Ehud Nir

Authors: Ehud Nir, Mansi Mehrotra, Amin Karami, Chris Holdenried, Robert Wang (Cadence Design Systems Inc., Toronto, Canada)



SPEAKER



Ehud Nir

General Manager and Sr. Director of Engineering at Ethernovia Canada Inc.

udi.nir@ethernovia.com | <https://www.ethernovia.com>

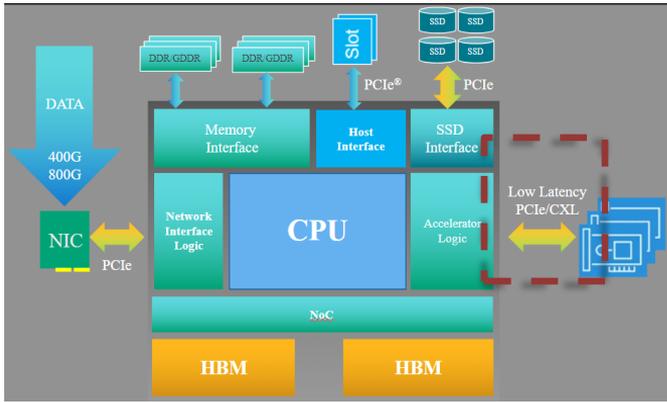
Ehud received the B.Sc. and M.Sc. degrees in Electronics Engineering from Tel Aviv University, Tel Aviv, Israel, in 2001 and 2013, respectively. He has over 25 years of experience in custom digital and DSP circuit design. In his previous position as Director of Digital Engineering at [Cadence Design Systems](#), he was developing high-speed digital circuits for wireline communication, including the circuits presented in this paper. He is also a husband, father and sailor.



MOTIVATION

- With the limited improvement in electrical links , architectural changes are required to deliver higher connectivity speed .
- MLSE is becoming a necessity to maintain reach and reduce retransmit rate.

Disaggregated Chiplet Architecture



- Due to its complexity, addition of MLSE to the PMA has significant impact on its latency, area and power.

Next generation (224G) optical link challenges

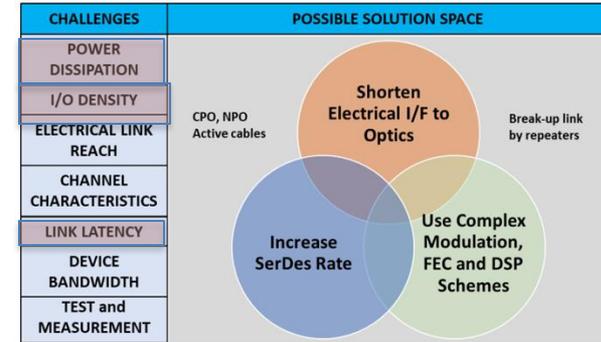


Figure 2 Next generation interconnect challenges



Figure 9 Block diagram of a DSP receiver

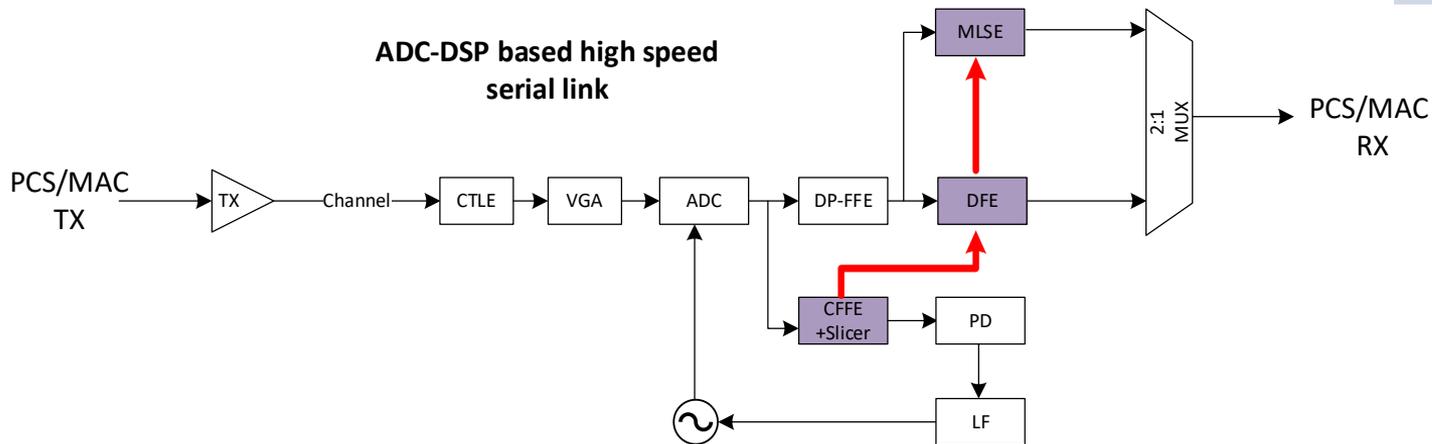
*Source: OIF-FD-CEI-224G



MOTIVATION

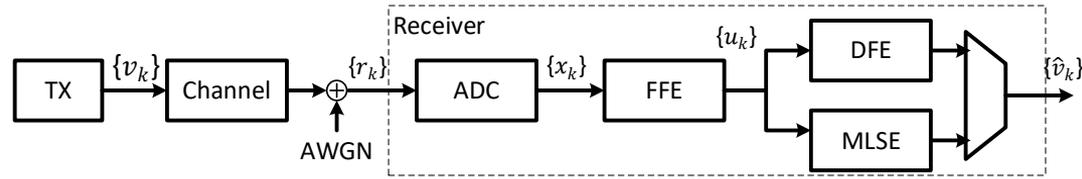
- Typical PAM4 DSP based SerDes includes 2-3 independent hard decision blocks, with increasing latency and decision accuracy
- Can soft information transferred between the blocks help to reduce overall latency while maintaining same BER?

Block	LR BER	Latency
CCFE	1e-4→1e-3	2ns
DFE	1e-7→1e-4	3ns
MLSE	1e-9→1e-6	9ns

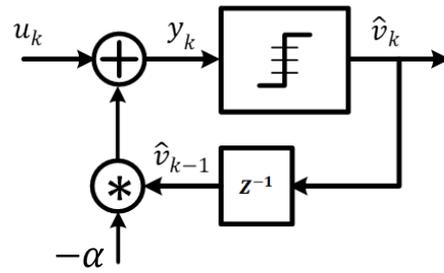


SERDES receiver equalization model

- ADC-DSP SerDes link model
- FFE output $\{u_k\}$ – all ISI removed beside 1st post cursor ($1 + \alpha D$ channel)
- DFE - Recovering transmitted symbol v_k using feedback from previous decision, \hat{v}_{k-1}
- ✓ In most cases the feedback will zero out the 1st post ISI, enabling correct slicer decision



$$u_k = v_k + \alpha v_{k-1} + W_k$$

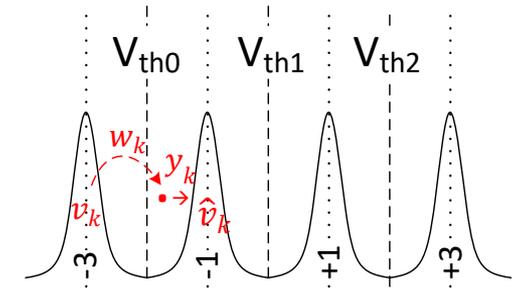


$$y_k = v_k + \alpha(v_{k-1} - \hat{v}_{k-1}) + W_k$$



DFE error propagation

- Assuming previous decision was correct ($v_{k-1} - \hat{v}_{k-1} = 0$), random error e_k will occur in the DFE when high noise w_k drives y_k beyond the slicer's threshold level
- In this case the feedback for next symbol no longer cancels the ISI
- Together with next symbol's noise (w_{k+1}), e_k can propagate to next symbol's decision and create an error burst $e_{k+1} = \pm 2 \rightarrow$ **main drawback of the DFE**



$$\hat{v}_k = v_k + e_k ; e_k = \pm 2$$

$$\rightarrow \alpha(v_k - \hat{v}_k) = \mp 2\alpha$$

$$\rightarrow y_{k+1} = v_{k+1} \mp 2\alpha + w_{k+1}$$

$$\rightarrow \hat{v}_{k+1} = v_{k+1} + e_{k+1}$$

$$\{v_k\} = \dots +1, -3, +1, -1, -3, -1, \dots$$

$$\{\hat{v}_k\} = \dots +1, -1, -1, +1, -3, -1, \dots$$

M. Emami Meybodi, H. Gomez, Y. -C. Lu, H. Shakiba and A. Sheikholeslami, "Design and Implementation of an On-Demand Maximum-Likelihood Sequence Estimation (MLSE)," in IEEE Open Journal of Circuits and Systems, vol. 3, pp. 97-108, 2022



MLSE

- MLSE stands for *Maximum Likelihood **Sequence** Estimation*
- Assuming i.i.d. symbols and AWGN channel , the most likely series of length K will have the smallest sum of Euclidean square distance from $\{u_k\}$ after convolution with the channel*:

$$[\hat{v}_0, \hat{v}_1, \dots, \hat{v}_K] = \arg \min \sum_{k=0}^K \|u_k - \hat{v}_k - \alpha * \hat{v}_{k-1}\|^2$$

- MLSE exhibits higher SNR by using current and past symbol's energy, in contrast with the symbol-by-symbol operation of the DFE.
- The MLSE implementation complexity is very high – it needs to compare 4^K variations of the sequence to select the most likely one

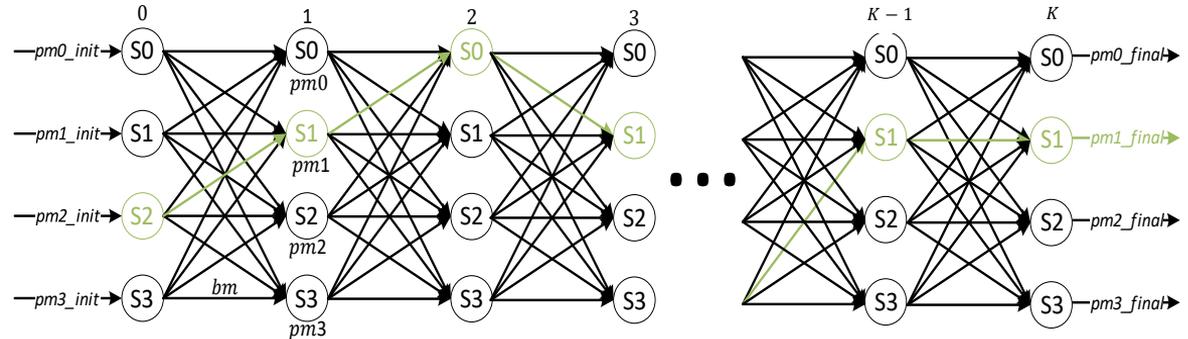
*Synchronization and traceback padding omitted

T. Xu , Z. Li, J. Peng, A. Tan, Y. Song, Y. Li, J. Chen and M. Wang., "Decoding of 10-G Optics-Based 50-Gb/s PAM-4 Signal Using Simplified MLSE," in IEEE Photonics Journal, vol. 10, no. 4, pp. 1-8, Aug. 2018



The Viterbi Algorithm

- VA is an efficient MLSE implementation using a sequential search in a trellis diagram
- The trellis consists of 4 states per unit interval (UI) for a PAM4 link over $1 + \alpha D$ channel
- The cost of traversing from one state to the next - the *branch metric* - is the Euclidean distance associated with these states: $\|u_k - \hat{v}_k - \alpha * \hat{v}_{k-1}\|^2$
- The state *path metric* is the total sum of the path's *branch metrics* leading to current state. A block called *Add-Compare-Select* is applied at each state/ UI to select the previous state that leads to minimum *path metric* for current state
- The smallest path metric in the final iteration for sequence of length K^* points to the **maximum likelihood sequence**
- By using VA the number of variations of the sequence looked at is reduced to $4^2 (K-1)$

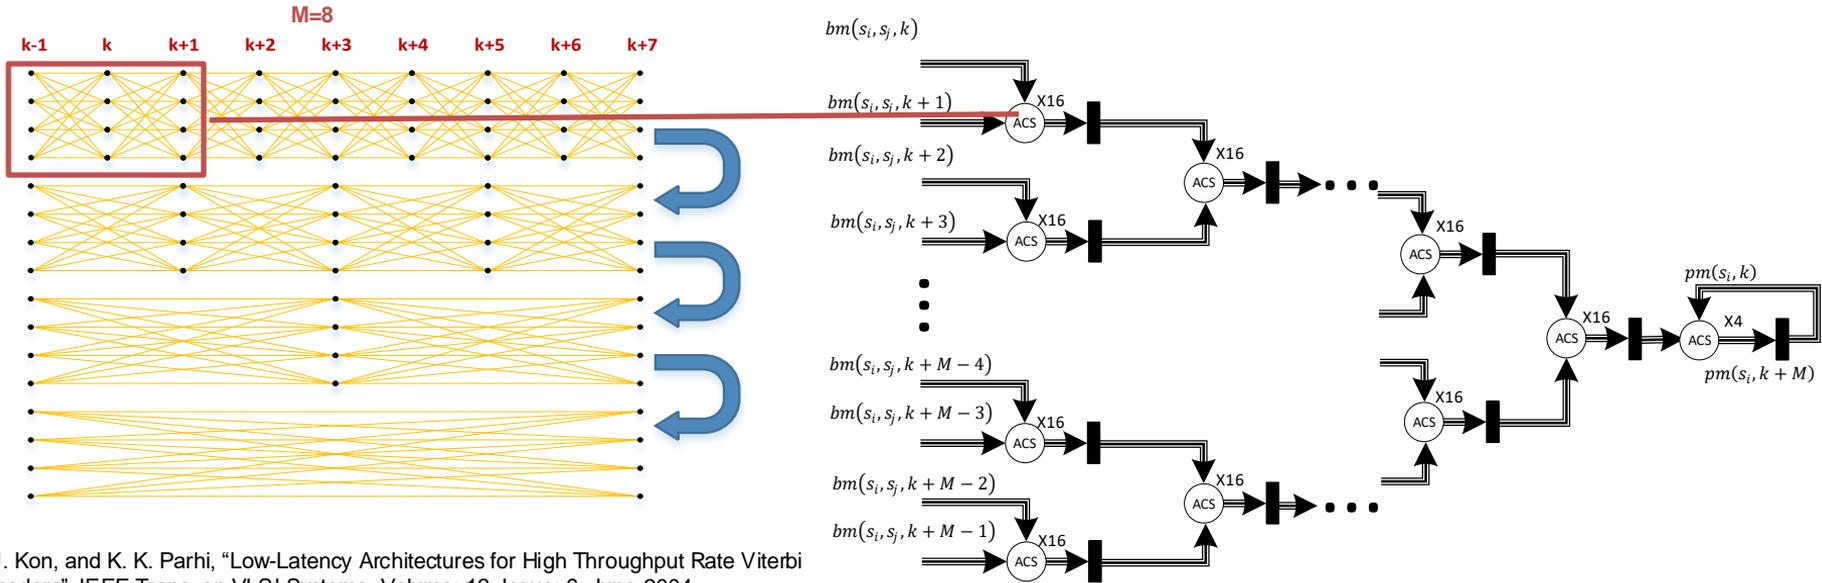


* Along with next cycle's traceback



Parallel implementation – Layered look ahead decoder

- Pipelined Viterbi decoder , processing a block of M symbols per cycle
- Intermediate states at time k are “looked ahead” by 16 ACS units which establish new minimum branches between time $k-1$ and $k+1$



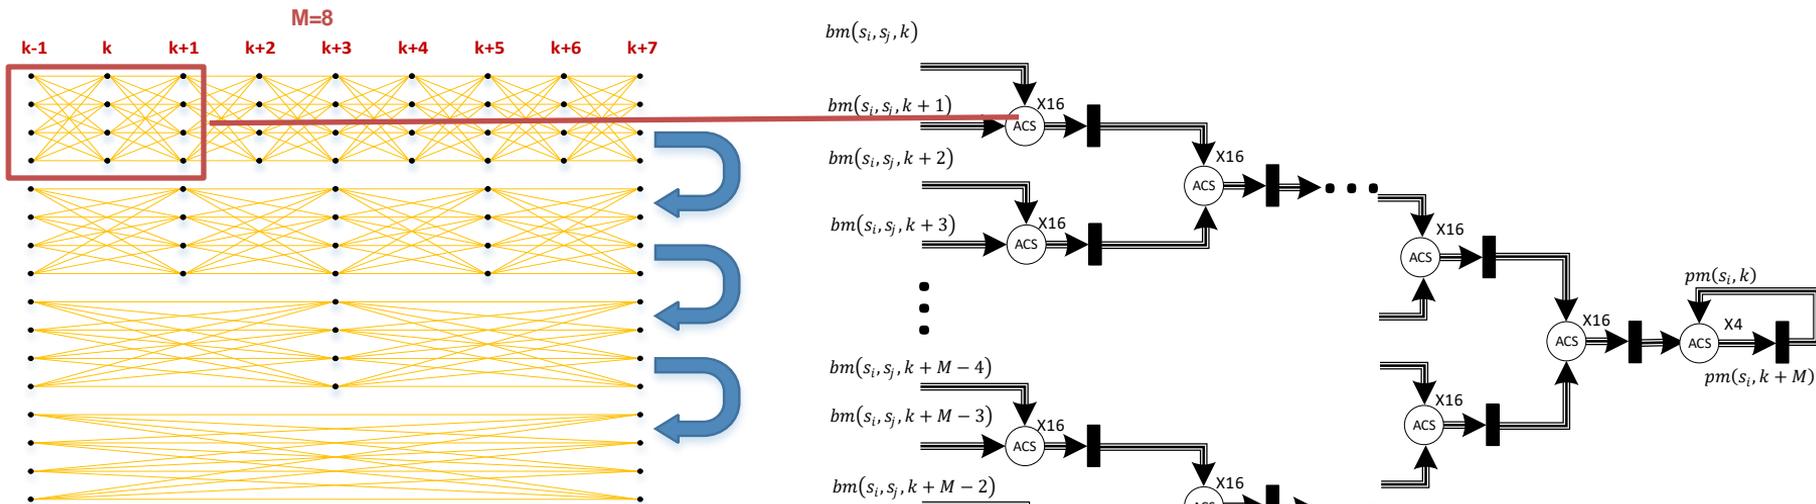
J. J. Kon, and K. K. Parhi, “Low-Latency Architectures for High Throughput Rate Viterbi Decoders”, IEEE Trans. on VLSI Systems, Volume: 12, Issue: 6, June-2004



Parallel implementation – Layered look ahead decoder

- The process repeats iteratively until 2 states are left
- Standard ACS calculates the 4 path metrics of the block*

*The minimum is determined by traceback from the next cycle



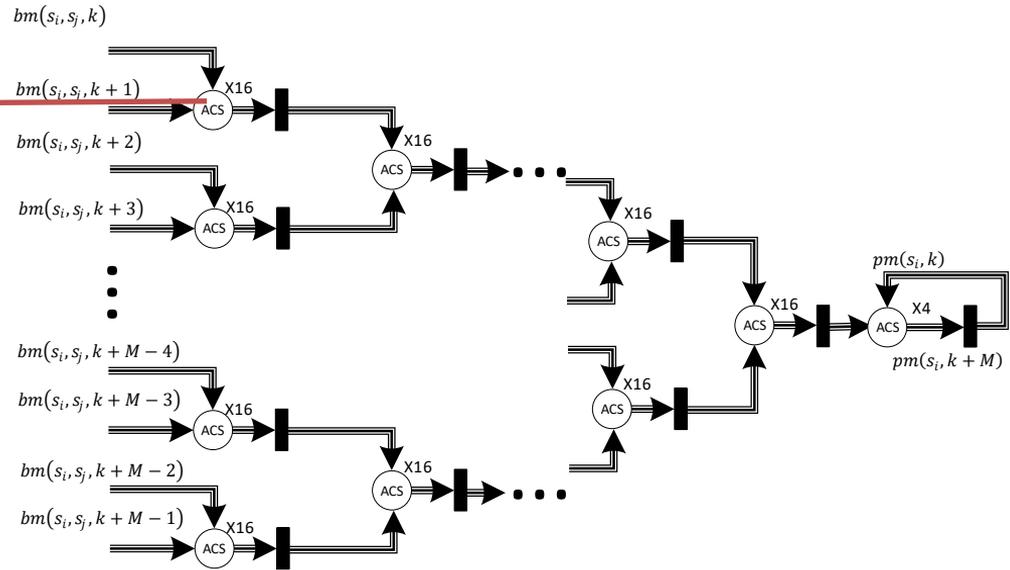
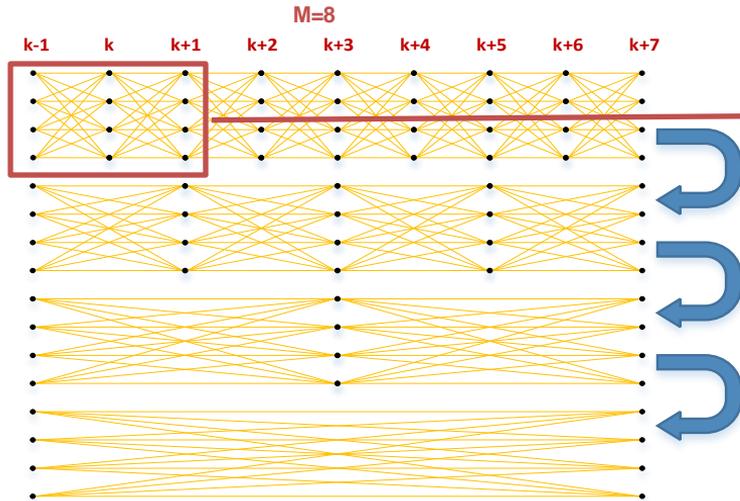
J. J. Kon, and K. K. Parhi, "Low-Latency Architectures for High Throughput Rate Viterbi Decoders", IEEE Trans. on VLSI Systems, Volume: 12, Issue: 6, June-2004



Parallel implementation – Layered look ahead decoder

- $Latency = 1 + \log_2 M \xrightarrow{M=32} 6$
- $Complexity$, measured by the number of 2 input adders/ comparators in the ACS tree:

$$7 * (16 * (M - 1) + 4) \xrightarrow{M=32} 3500$$

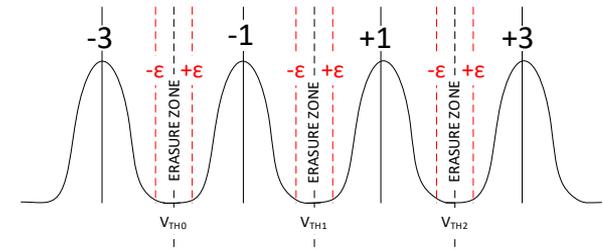
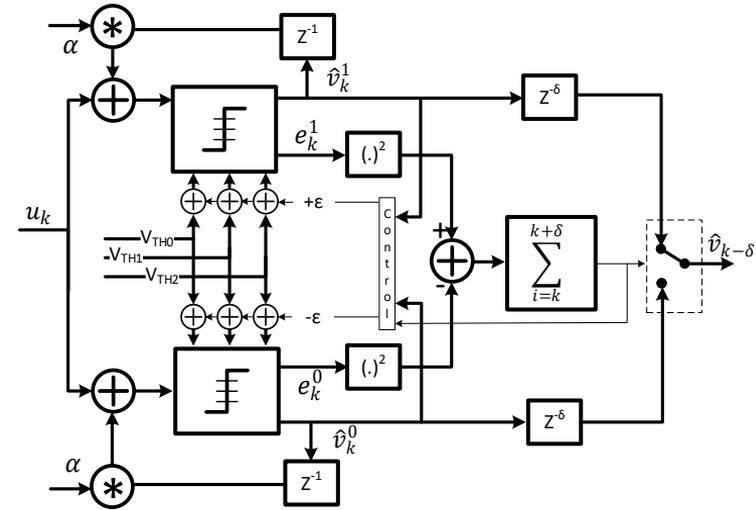


J. J. Kon, and K. K. Parhi, "Low-Latency Architectures for High Throughput Rate Viterbi Decoders", IEEE Trans. on VLSI Systems, Volume: 12, Issue: 6, June-2004



DFE Architecture Evolution – Dual Decision Feedback Equalizer

- Define *erasure zone* $[-\varepsilon, \varepsilon]$ around the thresholds
- Deploy two DFEs, one with thresholds skewed by $+\varepsilon$, and the other by $-\varepsilon$
- Incoming symbols with large noise will enter the erasure zone and trigger disagreement between the two DFEs
- When that happens, the skew is eliminated, and the error square (=Euclidian distance) is accumulated over a detection delay period δ
- The smallest sum at the end of the period points to the *more* likely decision for the triggering symbol
- 1.4dB improvement @ BER = $1e-5$ over standard DFE was reported for NRZ variant of DDFE



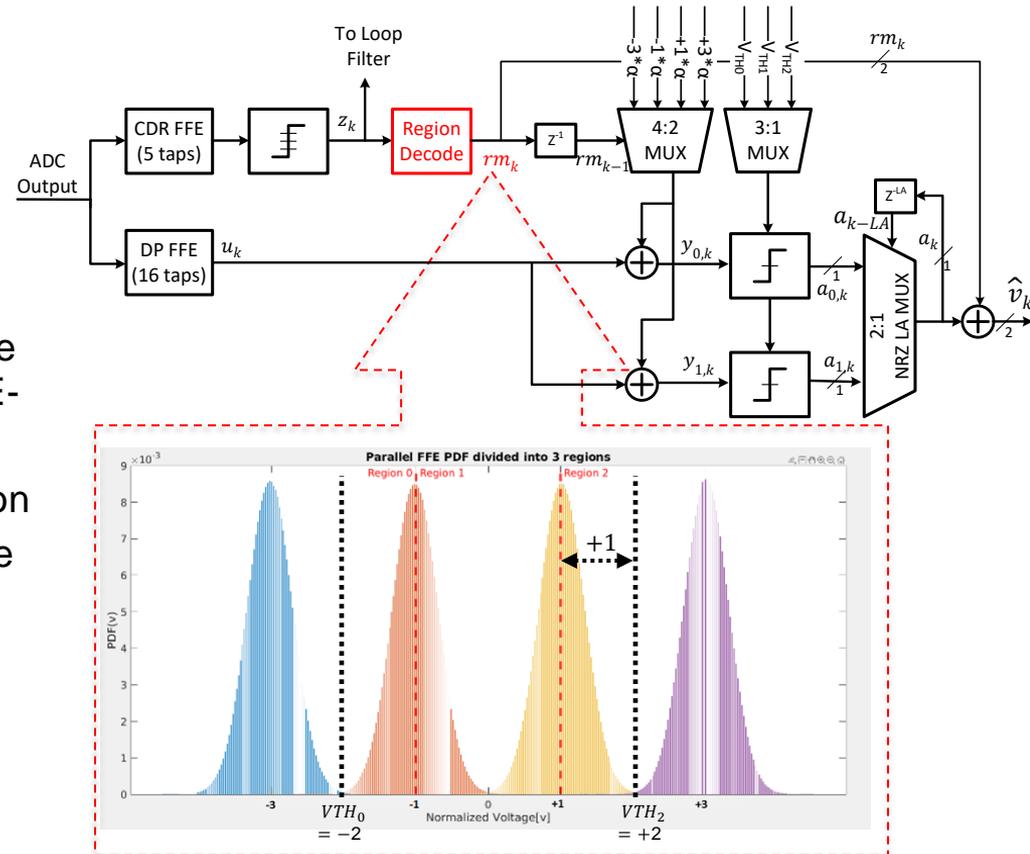
* ε range is $[0,1]$, where 1 is the $+1$ level amplitude

J.W.M. Bergmans, J.O. Voorman and H.W. Wong-Lam, "Dual decision feedback equalizer", IEEE Transactions on Communications, Volume: 45, Issue: 5, May 1997



Partially unrolled DFE

- A conventional 1 tap loop-unrolled DFE requires computation of all possible equalized values for the symbol at time k , based on 4 possibilities for the symbol at time $k - 1$.
 - PUDFE is eliminating 2 unlikely choices out of the 4, based on information from the faster CDR FFE-Slicer path
 - The CDR slicer output is decoded into likely region marker, rm_k . Probability of correct symbol to be out of the region is much lower than target BER
 - Candidates definition is modified to $rm_k + a_{0/1,k}$
- The PAM4 DFE is simplified into an NRZ like structure, using single threshold slicer and 2:1 muxing



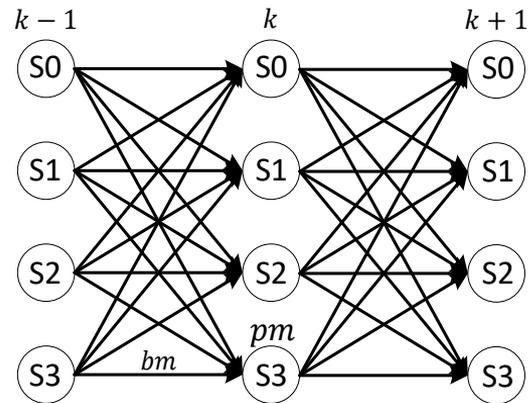
*LA=Look ahead

S. Kiran, S. Cai, Y. Luo, S. Hoyos, and S. Palermo, "A 32 Gb/s ADC-Based PAM-4 Receiver with 2-bit/Stage SAR ADC and Partially-Unrolled DFE", CICC 2018



Outline

- Motivation
- Receiver equalization and MLSE Introduction
- DFE Architecture Evolution
- **Speculative Error Correction for Partially Unrolled DFE**
- SEC validation in Silicon
- Summary and Conclusions



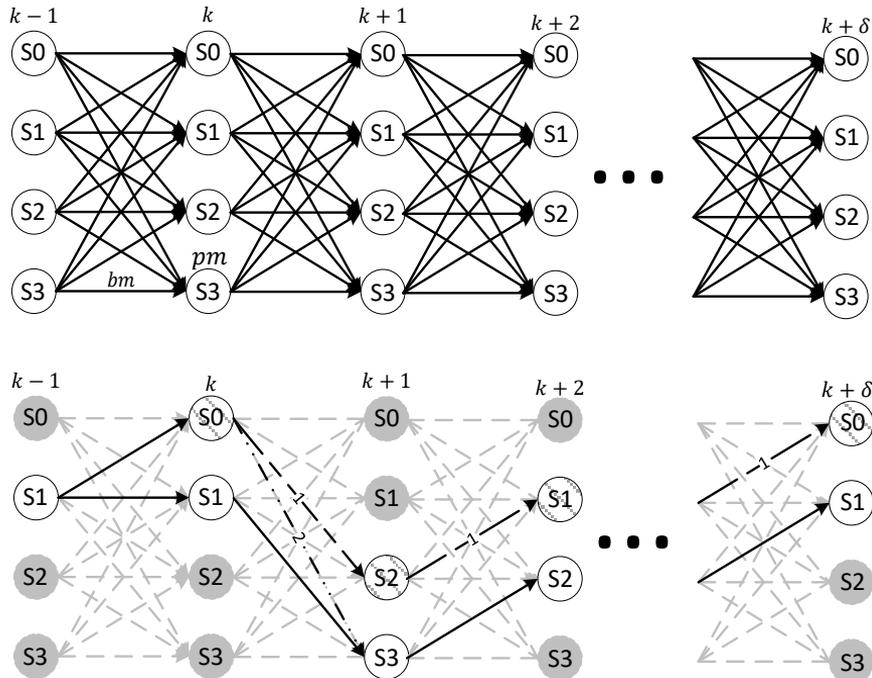
Speculative Error Correction rationale

- **Goal:** Improve the conventional DFE BER to MLSE levels, with limited impact on latency, area and power
- **Means:** Inspect the DFE loop unrolling candidates for errors using ML based detector; Correct the errors and eliminate the associated burst
- **Constraint :** To maintain the above goal, the detector must be simplified to select between the two most probable options, and complete operation within 1 cycle =1ns

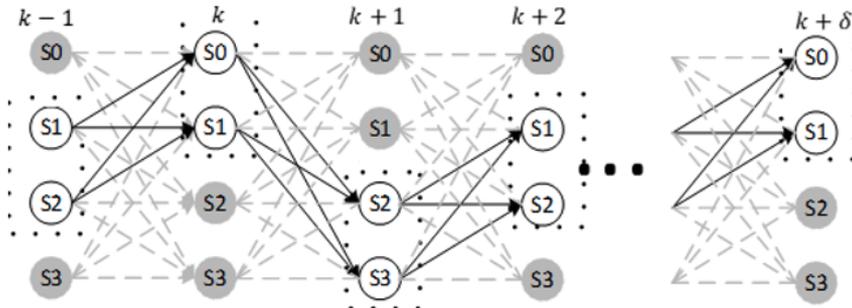


ML detector simplification

Our starting point is a VD searching for the ML sequence within a detection delay window of size $\delta + 1$: $4^2 * \delta$ trellis paths



Use PUDFE region marker rm_k to eliminate 2 unlikely states/ UI: $2^2 * \delta$ paths remain



Isolated error assumption in a candidate at time k : Multiple random error probability (burst excluded) within the detection window is $O(BER^2) \ll BER$. Only **2 paths** are left, with variable burst length (2 examples shown)

→ Simplified binary detector: $Vsum = \sum_{l=0}^{\delta} (bm_1(s_{k+l-1}^1, s_{k+l}^1, k) - bm_0(s_{k+l-1}^0, s_{k+l}^0, k))$



The SEC algorithm: Definitions

Inputs: unrolling candidate's slicer inputs at time k_0 :

unrolling candidate's slicer outputs at times $[k_0, k_0 + \delta]$:

PUDFE regions minimum at times $[k_0, k_0 + \delta]$:

FFE outputs at times $[k_0, k_0 + \delta]$:

PAM4 slicer thresholds:

post cursor ISI magnitude:

Output: PU candidate inversion (correction) signal at time k_0 :

Clocking: The algorithm is modeled at baud rate.

Integration: The SEC is introduced between the candidate slicers and the loop unrolling multiplexer

$$y_{i,k_0}, i \in [0, 1]$$

$$a_{i,k_0 \rightarrow k_0 + \delta}$$

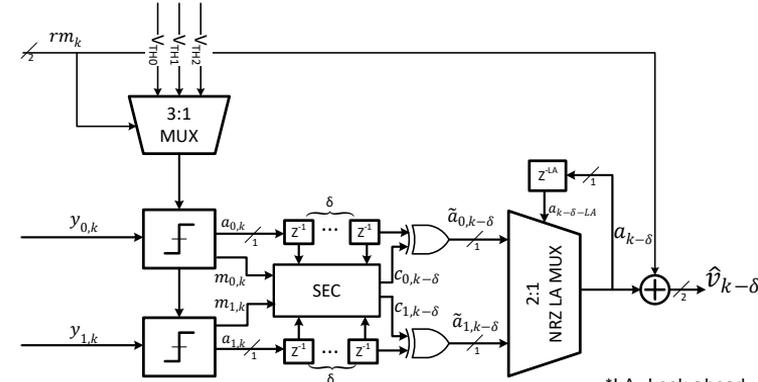
$$rm_{k_0 \rightarrow k_0 + \delta}$$

$$u_{k_0 \rightarrow k_0 + \delta}$$

$$vth_j, j \in [0, 2]$$

$$\alpha$$

$$c_{i,k_0}$$



*LA=Look ahead



The SEC algorithm: Initializations

for $i=0$ to 1 **do** // Loop through the two candidates

Initializations

$m_{i,k_0} = 0$ // SEC marker

$c_{i,k_0} = 0$ // SEC correction signal

$v_{sum} = 0$ // ML detector

if $|y_{i,k_0} - v_{th_{rm_{k_0}}}| < \epsilon$ **do** // Erasure zone marking (inside slicer)

$m_{i,k_0} = 1$

end if

if $m_{i,k_0} == 1$ **do** //Launch ML detector

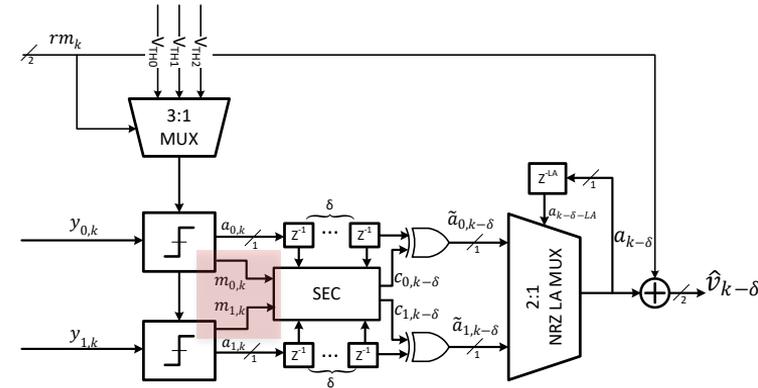
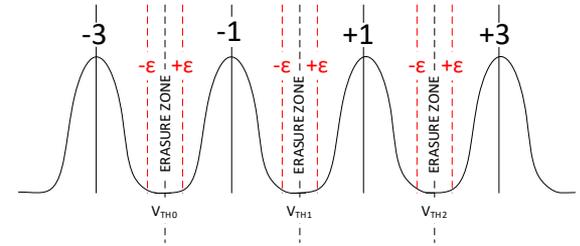
Create two sequences of length $\delta + 2$: $q_{l=0.. \delta+1}^0, q_{l=0.. \delta+1}^1$

$q_{l=0}^0 = i$ // Initialize sequences 0 location with prev. symbol

$q_{l=0}^1 = i$

$q_{l=1}^0 = a_{i,k_0}$ //Assign the original k_0 candidate

$q_{l=1}^1 = \sim a_{i,k_0}$ //Assign the spec. k_0 candidate correction



The SEC algorithm: ML detector

for $l=1$ to $\delta+1$ **do** //Apply simplified ML across detection delay candidate span

$$bm_{0,l} = (u_{0,k_0+l-1} - (rm_{i,k_0+l-1} + q_l^0) - \alpha * (rm_{i,k_0+l} + q_{l-1}^0))^2$$

$$bm_{1,l} = (u_{1,k_0+l-1} - (rm_{i,k_0+l-1} + q_l^1) - \alpha * (rm_{i,k_0+l} + q_{l-1}^1))^2$$

$$Vsum = Vsum + bm_{1,l} - bm_{0,l}$$

$$Vsum = \sum_{l=0}^{\delta} (bm_1(s_{k+l-1}^1, s_{k+l}^1, k) - bm_0(s_{k+l-1}^0, s_{k+l}^0, k))$$

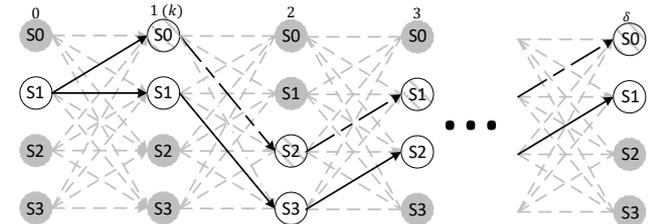
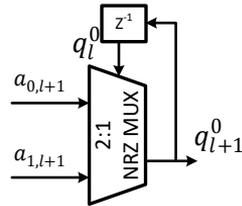
if $l \leq \delta$ **do** // Apply DFE mux to extract next symbol sliced values

$$q_{l+1}^0 = a_{i=q_l^0, l+1}$$

$$q_{l+1}^1 = a_{i=q_l^1, l+1}$$

end if

end for //detection delay for



if $Vsum < 0$ **do** // If ML detector is negative error corrected candidate is more likely

$$c_{i,k} = 1$$

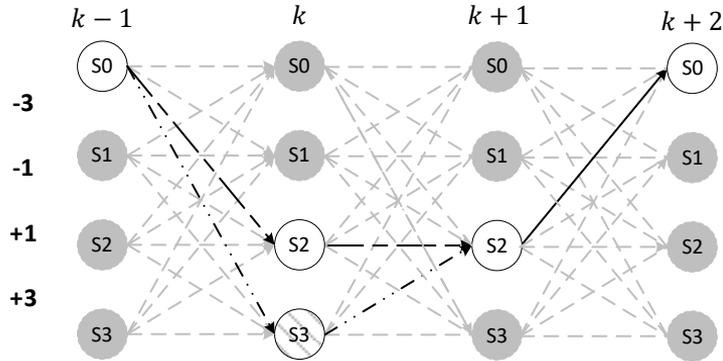
end if

end if

end for //Candidates for



SEC numerical example



DFE

Transmitter: $v_k = +3$
 DFE slicer input: $y_k = v_k + w_k = 1.9$
 DFE slicer output: $\hat{v}_k \xrightarrow{V_{TH2}=+2} +1$



Current symbol energy – noise drives correct BM slightly higher than erroneous. DFE only uses this information

PUDFE-SEC

SEC marker check: $m_k \xrightarrow{y_k - V_{TH2} = 0.1 < \epsilon} 1$

Original sequence: -3,+1,+1,-3

Alt sequence: -3,+3,+1,-3

Original candidate BMs sum: $\sum (0.9^2, 0.82^2, (-0.22)^2) = 1.53$

Alternative candidate BMs sum: $\sum (1.1^2, 0.13^2, (-0.22)^2) = 1.28$

Prev symbol energy (αD), has lower noise. Correct BM is much lower than erroneous. Used by SEC/MLSE only

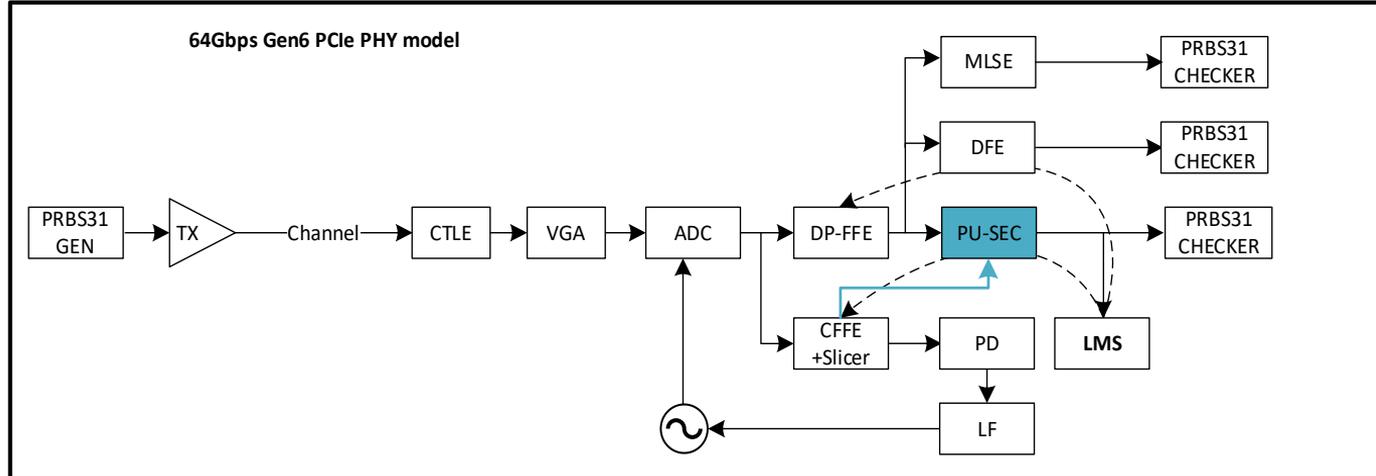
SEC output: $VSUM = -0.25 < 0 \rightarrow c_k = 1$

PUDFE-SEC output: $\hat{v}_k \xrightarrow{c_k=1} +3$



Gen6 PCIe PHY SEC Simulation Environment

- The SEC algorithm was implemented within a Gen6 PCIe PHY Simulink model
- It was integrated with the partially unrolled DFE, driven by the Datapath FFE and CDR FFE-Slicer

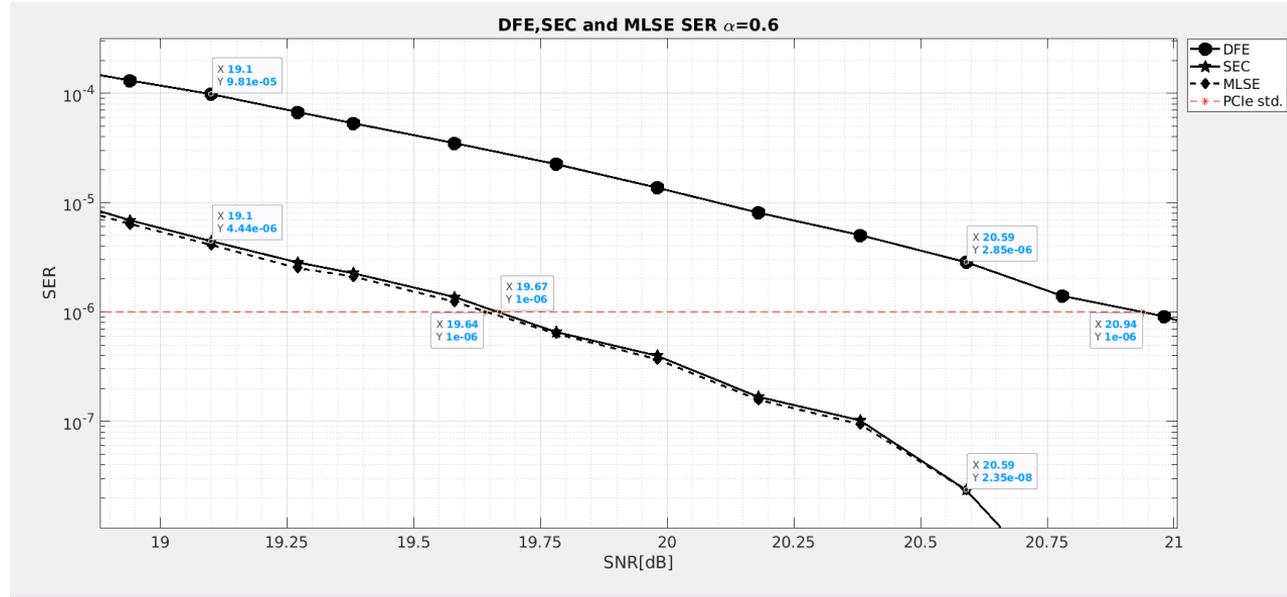


MATLAB and Simulink are registered trademarks of The MathWorks, Inc



SEC Simulation Results

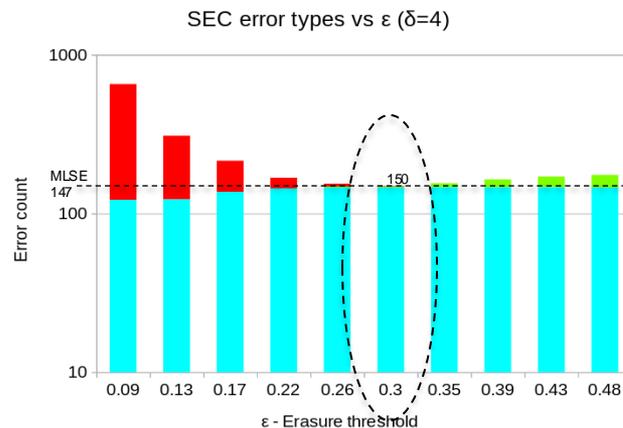
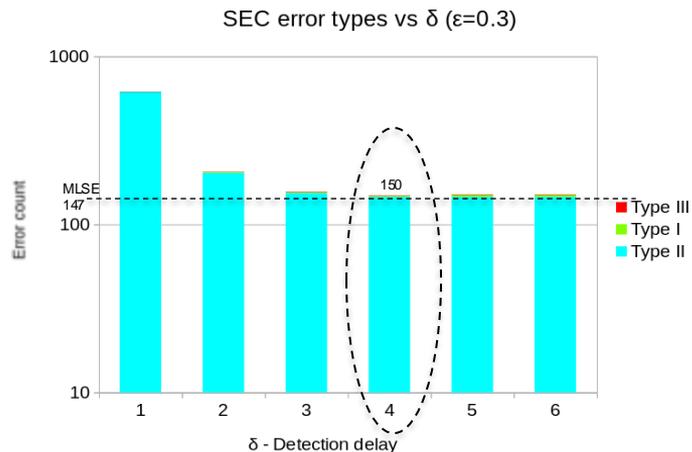
- The DFE, MLSE and SEC paths were simulated with >100M symbols for each datapoint
- The MLSE and SEC exhibit 1.3dB and 1.27dB of SNR gain respectively at PCIe SER standard of 1e-06.
- The SEC-PUDFE is providing 15x to >100x SER improvement over the conventional DFE



SEC error types and dependency

Error Type	Description	Dependency
I	Error within the detection delay , in symbol other than the marked candidate	Dominant when ϵ and δ are high - probability to include more errors within the detection delay increases
II	SEC ML detector decision is wrong	Dominant when detection delay δ is too short. p_{II} converges to the MLSE error probability as δ is increased.
III	SEC was not triggered	Dominant when ϵ is low.

- In a 15M symbol simulation ($\alpha=0.6$) , minimum was reached with $\epsilon=0.3$ and $\delta=4$.
- The SEC SER was $1e-5$ and error count was 150 - 2% above the MLSE (147).

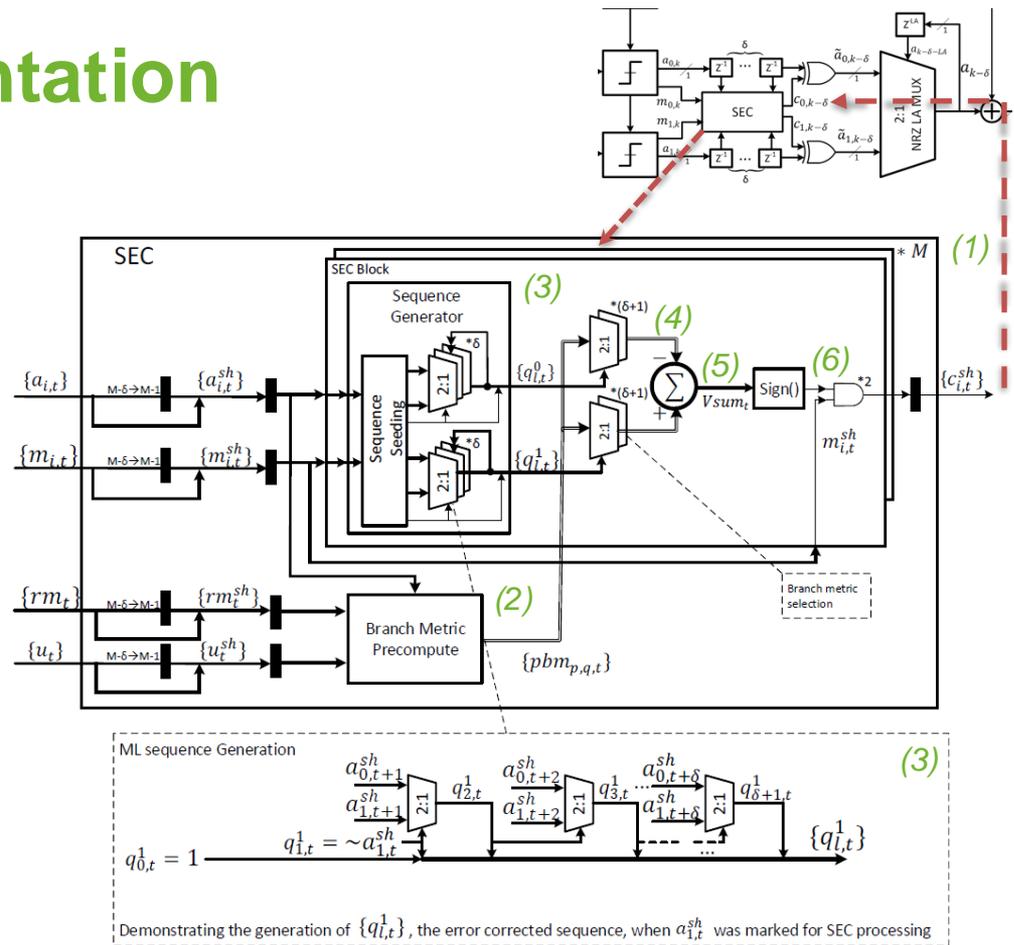


SEC circuit implementation

- The circuit is clocked with 1GHz clock (Concurrent SEC operation over $M=32$ UIs/ cycle) (1)
- Branch metrics are precomputed to ease timing (2)
- Original and error corrected sequences are generated (3) and select respective BMs (4)
- Selected BMs are accumulated with respective sign (5) and correction signal is extracted from $Vsum\ sign$ (6)

- **Latency** $= 1 + \frac{\delta}{M} \xrightarrow{M=32, \delta=4} = 1.13$

- **Complexity** $= M * 2 * (\delta + 1) \xrightarrow{M=32, \delta=4} = 320$



SEC Latency and PPA comparison

	Layered look ahead VD **	This work: SEC with $\delta = 4$	Fully unrolled DFE
Technology	-	5nm	5nm
Data Rate	64Gbps	64Gbps	64Gbps
Modulation	PAM-4	PAM-4	PAM-4
Parallel Block Size (M)	32	32	32
Clock Frequency [GHz]	-	1	-
Added latency [cycles]	6	1.1 (18%)	-
Complexity metric*	3500	320 (9%)	-
5nm FinFET Technology Synthesis gate area [μm^2]	-	4.5K (SEC) 10K (PU-SEC)	4.5K
SNR [dB] @ SER=1e-6	19.64	19.67	20.94

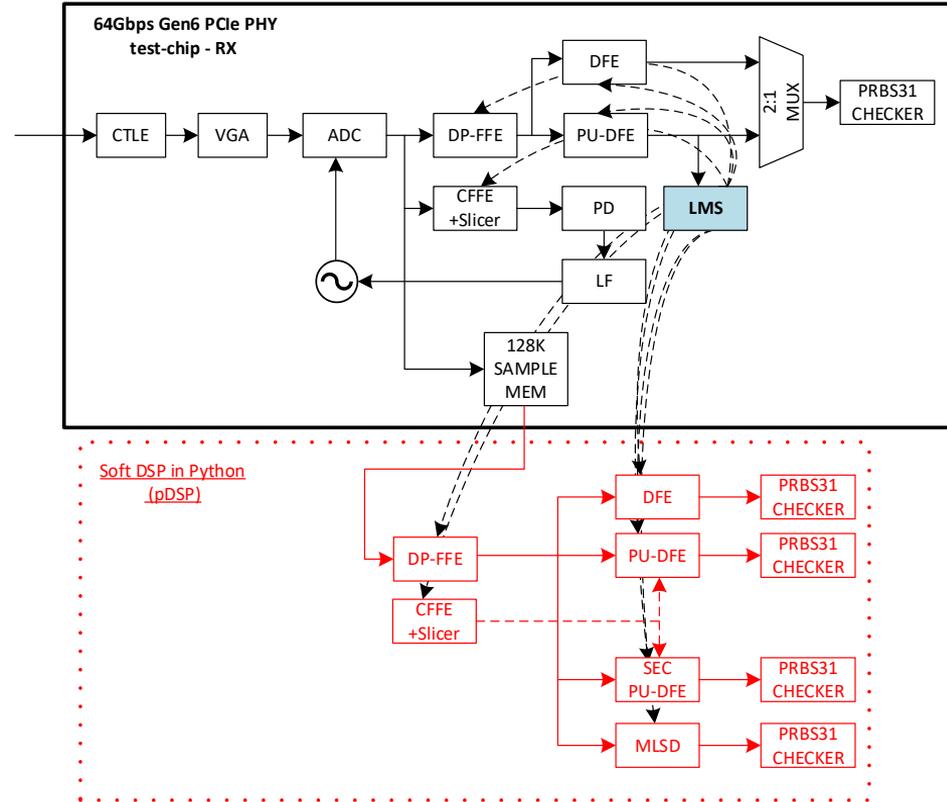
* Number of 2 input adders used in the ACS unit

**Based on J. J. Kon, and K. K. Parhi, "Low-Latency Architectures for High Throughput Rate Viterbi Decoders", IEEE Trans. on VLSI Systems, Volume: 12, Issue: 6, June-2004



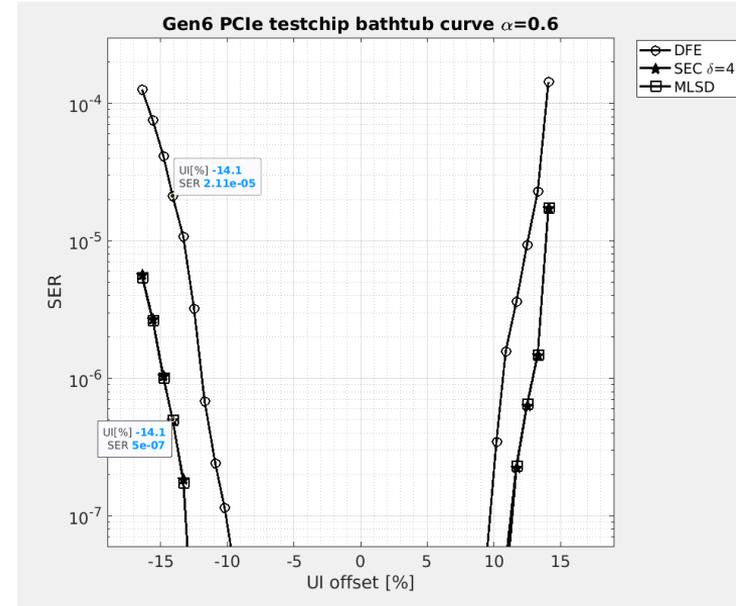
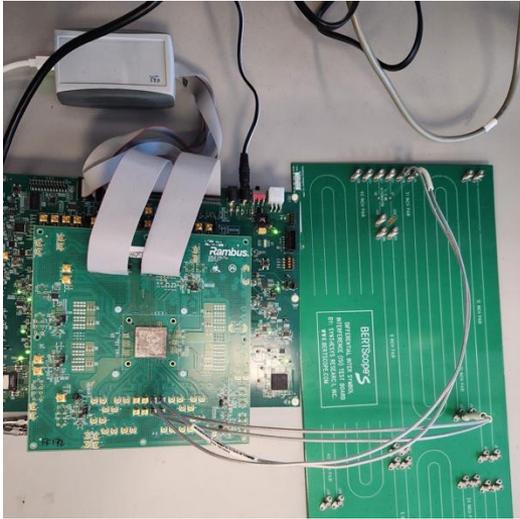
Gen6 PCIe PHY SEC silicon validation setup

- The SEC was validated in silicon by post processing samples from PCIe Gen6 (64Gbps) wireline transceiver test-chip fabricated in 5nm FinFET technology.
- The Receiver AFE includes a Continuous Time Linear Equalizer (CTLE) to equalizes the incoming signal, followed by a Variable Gain Amplifier (VGA) to adjusts the signal power to the ADC full scale.
- A 32 GS/s time interleaved ADC with programmable full scale digitizes the VGA output and the data is then processed in the DSP.
- The DSP is processing 32 ADC samples every 1GHz cycle. The data path consists of a 16-tap FFE and a 1-tap DFE.
- In the timing recovery path (DTL), a 7-tap FFE and slicer are followed by baud-rate phase detector which controls a Phase Interpolator (PI). The PI adjusts the quadrature sampling clock phase and frequency
- The SEC was implemented in python (pDSP) and driven with samples from the on-chip ADC memory buffer.
- The Receiver LMS engine was used to train the pDSP coefficients.



Horizontal SER Bathtub measurement

- The SEC-PUDFE bathtub was measured with a 36 dB overall channel loss
- It is providing 15x to 60x SER improvement over the conventional DFE, when DFE SER $1E-4$.
- The number of errors at the SEC-DFE output are within $\pm 5\%$ compared with the MLSE.



* To reduce the run time to 10 hours or less per datapoint, the measured MLSE/SEC SER was limited to values higher than $1E-7$



Conclusions

- The SEC circuit presented in this paper, is an improvement to partially unrolled DFE resulting in SER that approaches that of full MLSE implementation (<0.03dB) @ PCIe target BER of 1e-6.
- A simplified and very short ML detector is applied on loop unrolling candidates that exhibits high noise levels. If it is more likely that a candidate has an error, the detector will correct it.
- By correcting the random errors, any associated bursts are eliminated.
- The algorithm is considered speculative because most candidates corrected by it will eventually not be selected by the loop unrolling multiplexer.
- For block size of 32 samples, the added latency of the circuit is ~1ns, and the synthesized area is 4.5K μm^2 in 5nm FinFet technology.
- These values are much lower than those required for full MLSE.
- The SEC was verified in simulation and validated by post processing ADC samples read from a PCIe Gen6 (64Gbps) wireline transceiver test-chip fabricated in 5nm FinFet technology .



Thank you!



QUESTIONS?

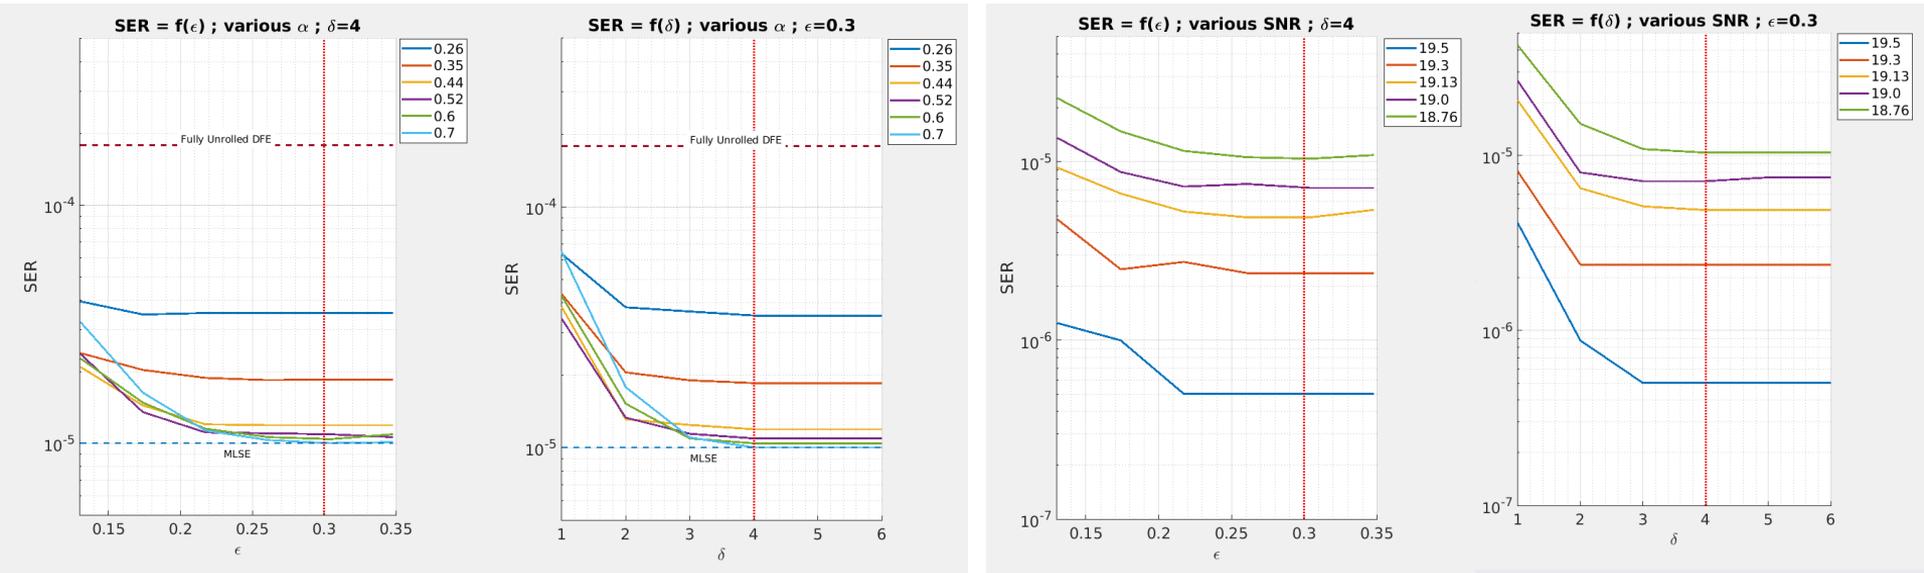
Ehud Nir

Director of Digital Engineering, Cadence Design Systems Inc.

enir@cadence.com | www.cadence.com



Optimum ϵ , δ for various α , SNR



- No significant variability of SER when ϵ , δ are increased above previously defined optimum

