

We3G-2

Equalization Tuning of the PCIe Physical Layer by Using Machine Learning in Industrial Post-silicon Validation

F. E. Rangel-Patiño^{1,2}, A. Viveros-Wacher², C. Rajyaguru³, E. A. Vega-Ochoa²,
S. D. Rodriguez-Saenz², J. L. Silva-Cortes², H. Shival³, and J. E. Rayas-Sánchez¹

¹Department of Electronics, Systems, and Informatics, ITESO – The Jesuit University of Guadalajara,
Tlaquepaque, Jalisco, 45604 Mexico

²Intel Corp. Zapopan, Jalisco, 45019 Mexico

³Intel Corp. Folsom, CA, 95630 USA

- Introduction to PCIe PHY tuning
- A machine learning (ML) proposal
- Unsupervised and supervised ML
- PHY tuning and optimization
- Clustering and GPR modeling results
- PHY optimization formulation and results
- Conclusions

The Journey of PCIe

Continuous improvement in data rates and usage models



Servers



Storage




Communications



Embedded

I/O BANDWIDTH DOUBLES EVERY 3 YEARS



Usage of PCIe technology continues to grow driven by BW demand

PCIe Tuning Complexity


- PHY tuning is very time consuming
- Typically based on exhaustive search methods to find the “best” Tx and Rx EQ settings

1 Identify tuning “knobs”

PHY tuning may involve hundreds of Tx/Rx knobs to be adjusted


2 Define the FOM (i.e Performance metric)

Define a metric to qualify the recipe ingredients, which will be used by the optimization algorithms




3 Knobs optimization to maximize FOM

Find the best knobs values combination that optimizes the FOM



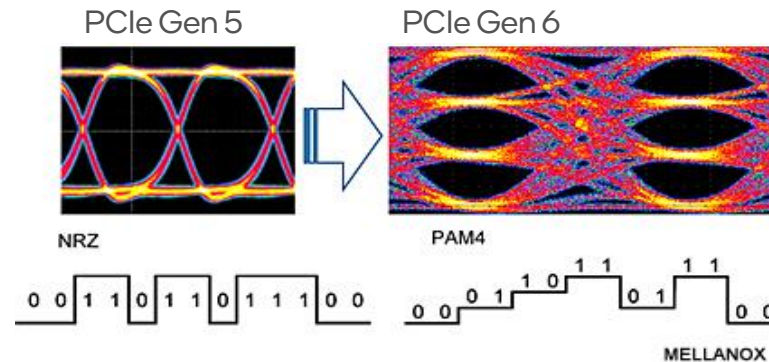
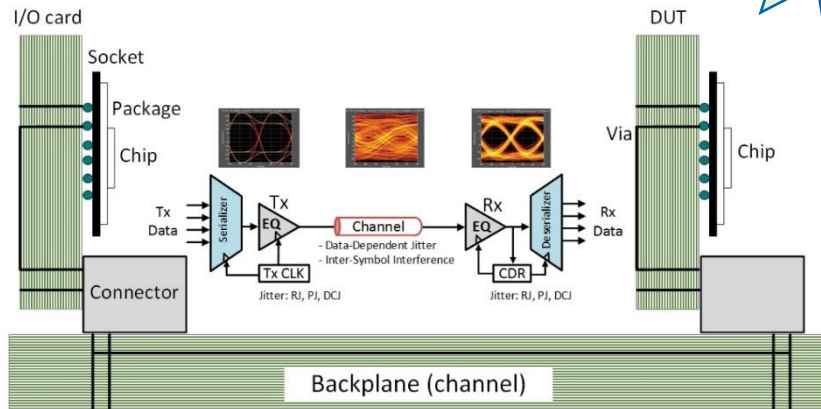
4 Confirm runs to verify results



5 Recipe is done




Challenge: make PHY tuning inexpensive by reducing the number of lab measurements




Challenge: provide higher margins on a wide variety of devices and channels, along with process, voltage and temperature (PVT) variation


Many EQs maps, obtained from lab measurements, are typically employed for different lanes, channels, PVT, etc.



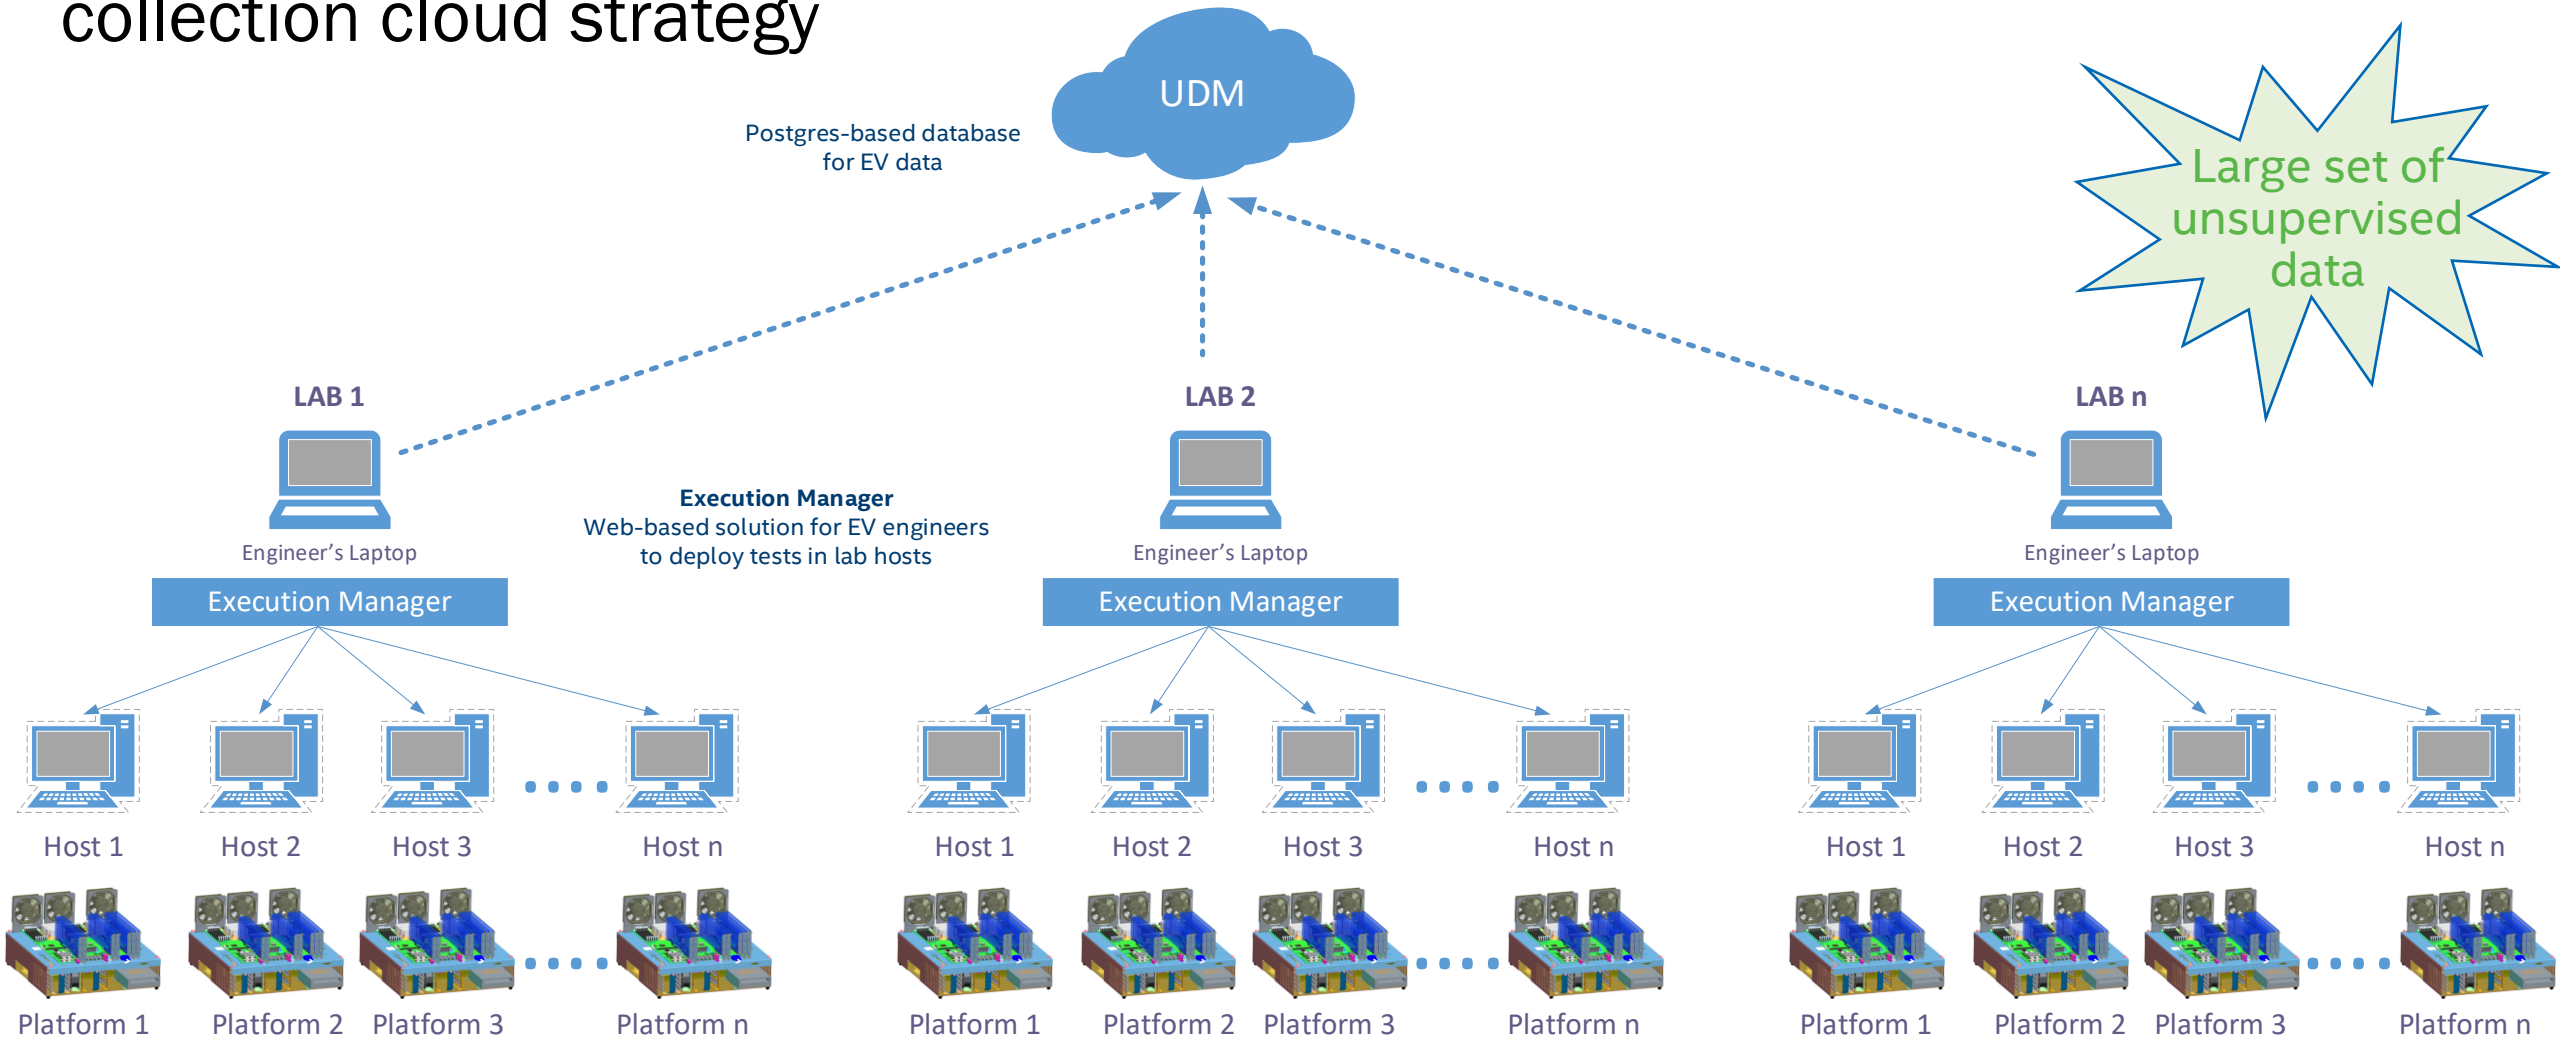
Short to mid range interconnect
chip to chip within a PCB



Mid to long range interconnect chip to
chip across a backplane/midplane or a
cable



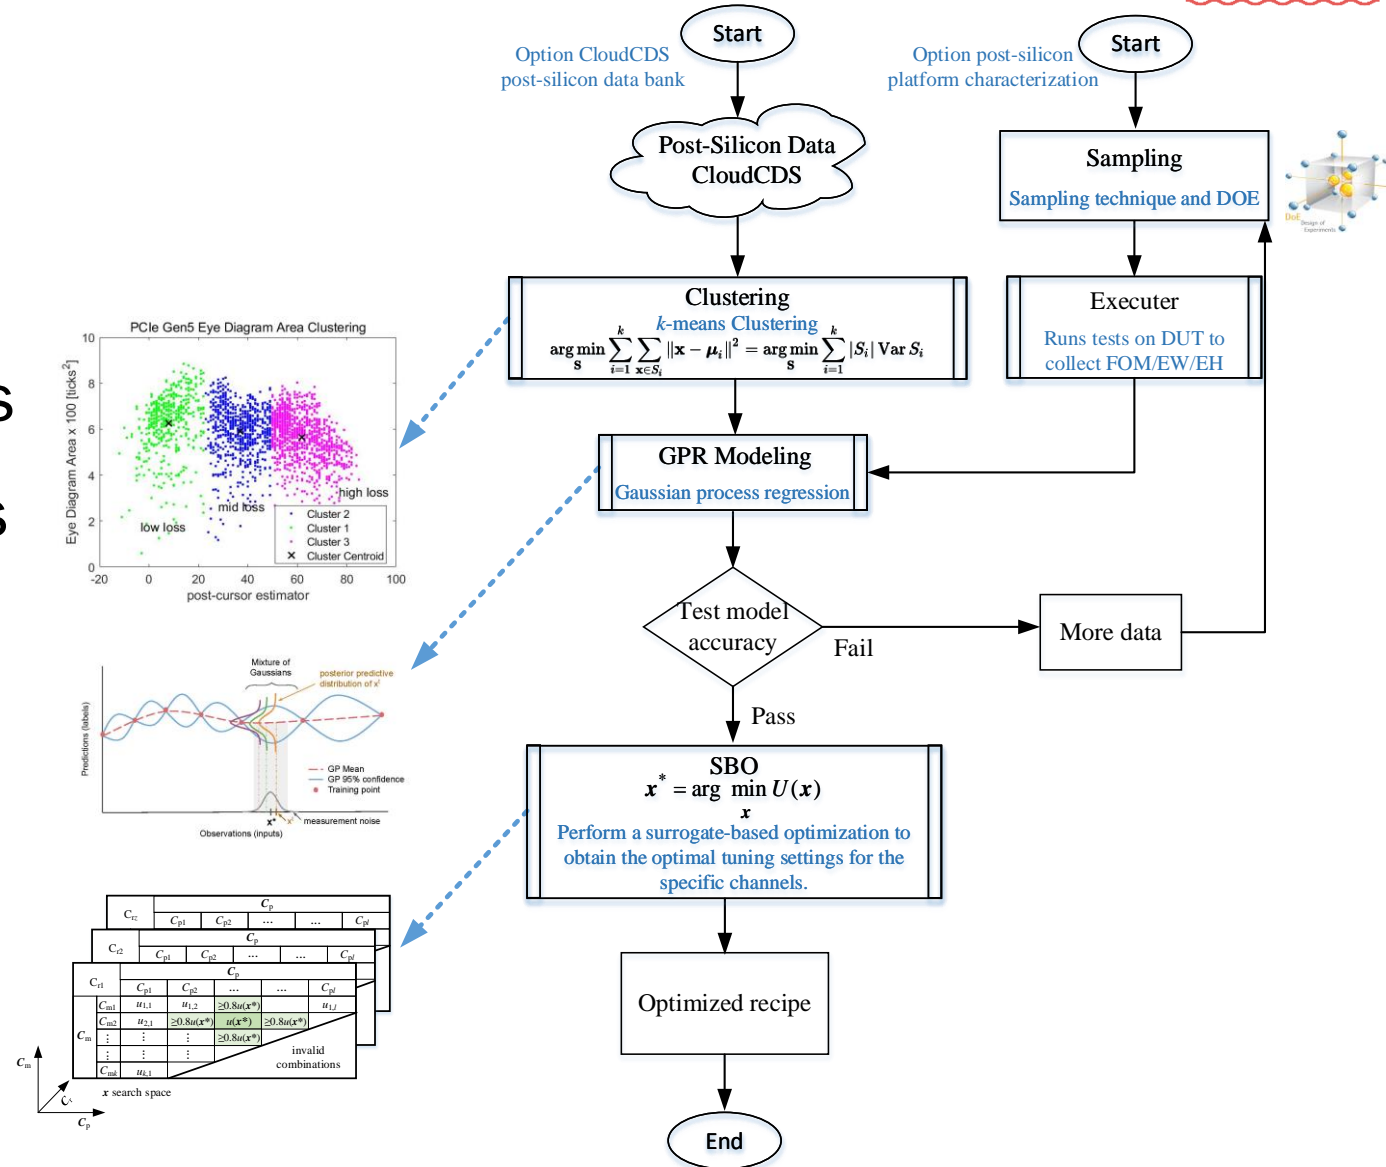
To maximize resources globally, electrical validation is moving to a data-collection cloud strategy



A Machine Learning Proposal


- We propose a combination of unsupervised and supervised ML techniques to tune high-speed PCIe5 NRZ and PCIe6 PAM4 designs on a real post-Si platform
- Our proposal helps to:
 - Eliminate manual PHY tuning efforts
 - Reduce engineering and debugging costs
 - Accelerate post-silicon validation
- We aim at increasing accuracy and robustness of the solution by using massive and diverse post-Si validation cloud data

- We use unsupervised ML modeling to cluster prior SMV and EQ data
- The clustered data segments are used to train supervised ML models
- These models approximate margins for a given combination of PHY parameters
- The models are used to find optimal PHY tuning parameters through a surrogate-based optimization (SBO)



Unsupervised ML: Clustering

- Unsupervised ML algorithms learn patterns from untagged data
- Clustering is the task of grouping similar data points
- We use the k-means clustering algorithm that partitions n observations into k clusters




Given a set of observations (x_1, x_2, \dots, x_n) , where each observation is a d -dimensional real vector, k-means clustering aims to partition the n observations into k ($\leq n$) sets $S = \{S_1, S_2, \dots, S_k\}$ to minimize the within-cluster variance

We solve:

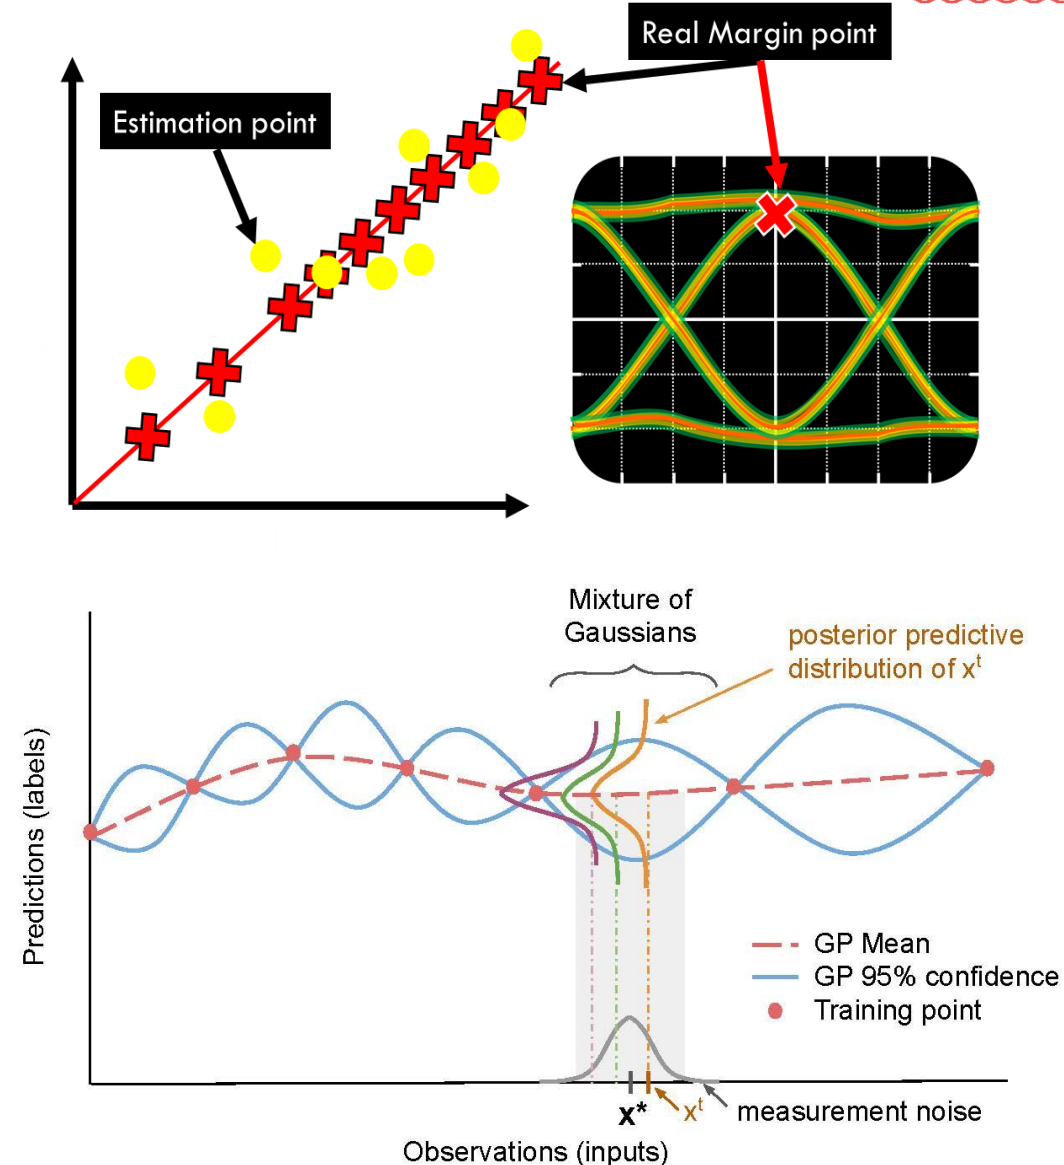
$$\arg \min_{\mathbf{S}} \sum_{i=1}^k \sum_{\mathbf{x} \in S_i} \|\mathbf{x} - \boldsymbol{\mu}_i\|^2 = \arg \min_{\mathbf{S}} \sum_{i=1}^k |S_i| \text{Var } S_i$$

where $\boldsymbol{\mu}_i$ is the mean of points in S_i




By Chire - Own work, CC BY-SA 4.0,
<https://commons.wikimedia.org/w/index.php?curid=59409335>


- We use supervised machine learning regression to model eye margins
- Given the large statistical fluctuations in post-Si electrical validation measurements, we use Gaussian process regression (GPR)
- GPR aims to predict not only the outputs based on inputs, but also their variability (probability distribution)




Clustering Results






*K-means
Clustering*



GPR Modeling Results




GPR
Modeling




Equalization: PCIe PHY Tuning

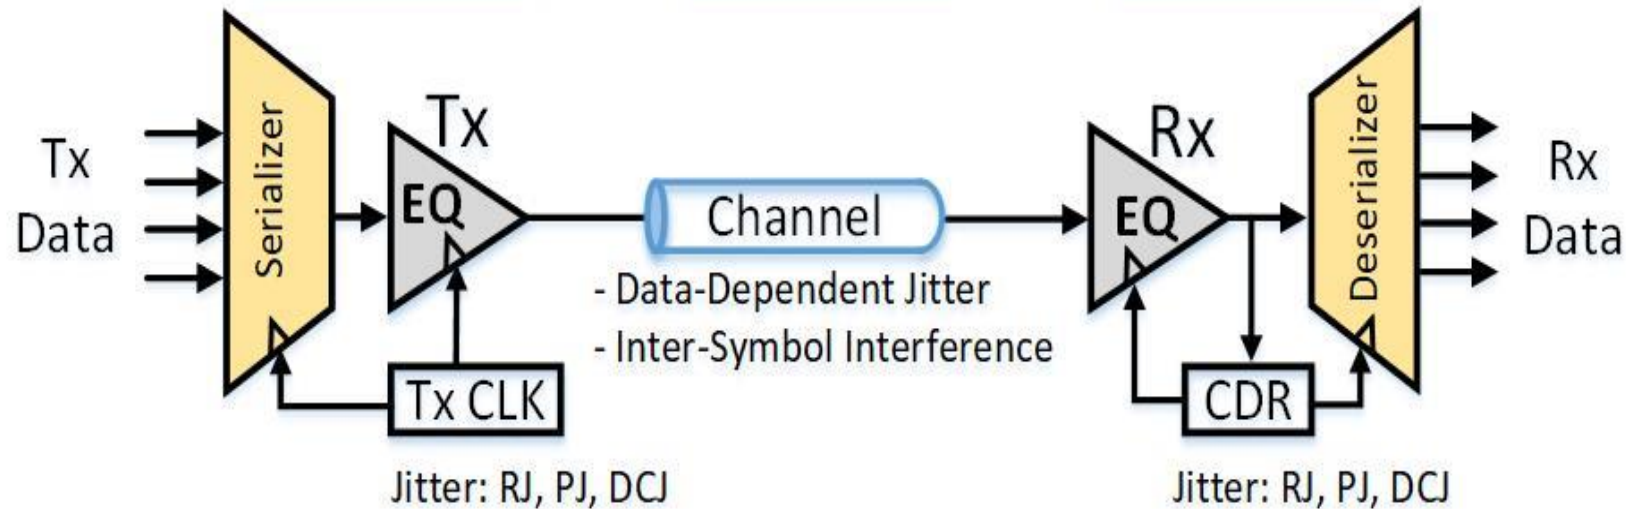

An “open eye” at
the Transmitter



A “closed eye”
at the Receiver



Equalized



Equalization settings need to be optimized for best link performance



- $\mathbf{R}_m \in \mathbb{R}^2$: GPR response using functional margins
 $\mathbf{R}_m = \mathbf{R}_m(\mathbf{x}, \boldsymbol{\psi}, \boldsymbol{\delta}) = [e_w(\mathbf{x}, \boldsymbol{\psi}, \boldsymbol{\delta}) \quad e_h(\mathbf{x}, \boldsymbol{\psi}, \boldsymbol{\delta})]^T$ →
- We solve by direct optimization
 $\mathbf{x}^* = \arg \min_{\mathbf{x}} U(\mathbf{x})$ where \mathbf{x} has the EQ settings
- Our unconstrained objective function is

$$U(\mathbf{x}) = -[e_w(\mathbf{x})][e_h(\mathbf{x})] + L(\mathbf{x}) \left[\frac{|u(\mathbf{x}^{(0)})|}{\max\{\mathbf{l}(\mathbf{x}^{(0)})\}} \right]$$

where $\mathbf{x}^{(0)}$ is the starting point and $L(\mathbf{x})$ is a corner limits penalty function defined as

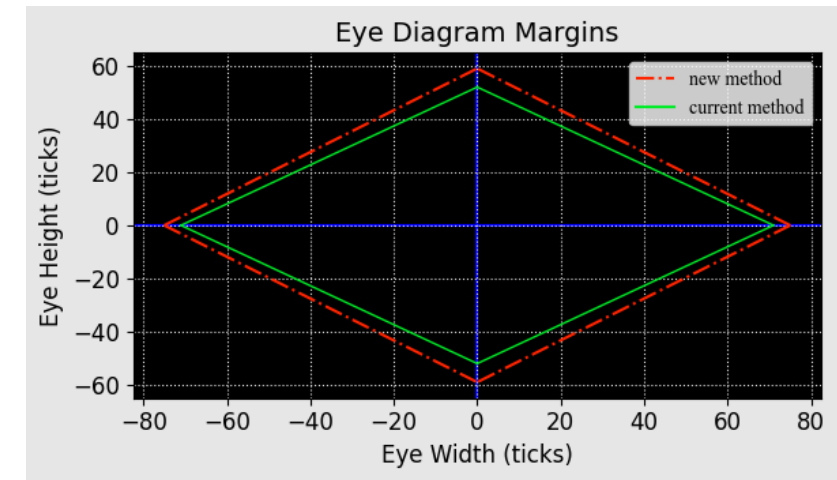

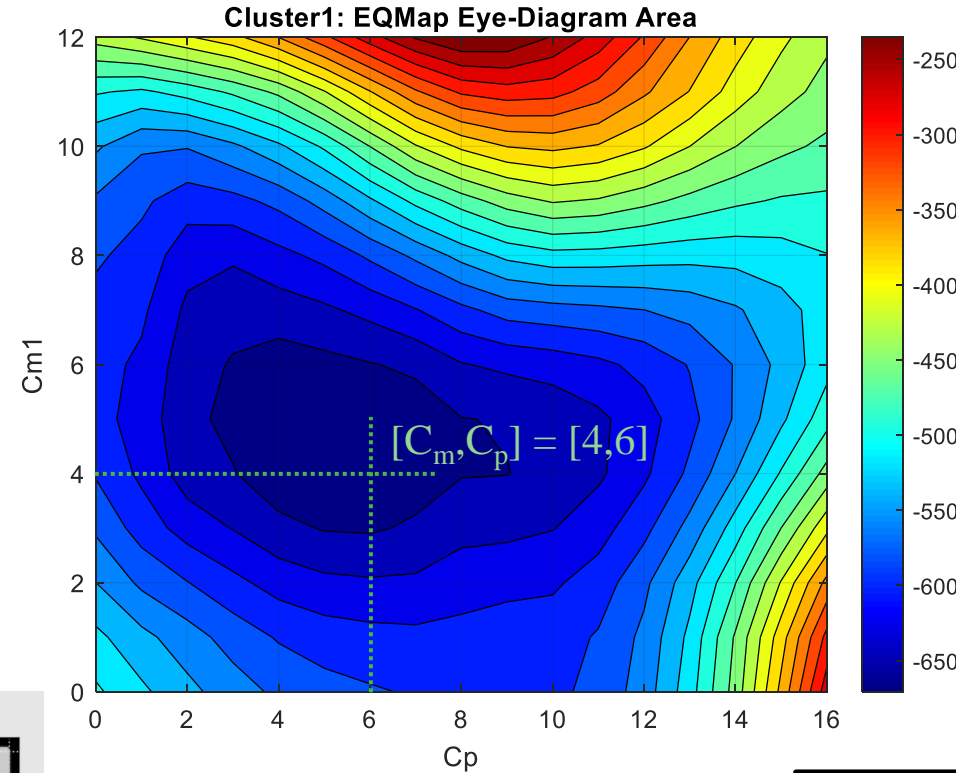

$$L(\mathbf{x}) = \max\{0, \max\{\mathbf{l}(\mathbf{x})\}\}$$

$$\text{with } \mathbf{l}(\mathbf{x}) = 0.8u(\mathbf{x}) \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} - \begin{bmatrix} u(\mathbf{x} - [1 \ 0]^T) & u(\mathbf{x} + [1 \ 0]^T) \\ u(\mathbf{x} - [0 \ 1]^T) & u(\mathbf{x} + [0 \ 1]^T) \end{bmatrix}$$



lower bounds of $0.8u(\mathbf{x}^*)$ in the vicinity of \mathbf{x}^*

SBO PCIe PHY Tuning Results



- 20% increase in eye diagram area (measured)
- Current method takes days, while new method takes a few hours

PCIe Gen5 – FIR EQ Opt. Summary

Cluster 1: $\mathbf{x}^* = [4 \ 6]^T$, $U(\mathbf{x}^*) = -672.77$
 Cluster 2: $\mathbf{x}^* = [4 \ 8]^T$, $U(\mathbf{x}^*) = -719.35$
 Cluster 3: $\mathbf{x}^* = [5 \ 4]^T$, $U(\mathbf{x}^*) = -688.73$

- We propose a new methodology for PCIe link equalization
- Our methodology can be applied to other interfaces for PHY optimization
- We use ML techniques to cluster post-silicon data from different channels and feed those clusters to a GPR-based metamodel for each channel
- We use SBO to find the optimal PHY settings
- A significant increase in eye diagram margins is achieved, accelerating the PHY tuning process