# Windfoil: Workload-Independent Formation of Power and Thermal Islands for SoC Characterization

Michael Grieco
PI: Prof. Tanvir Ahmed Khan
May 15th, 2026

COLUMBIA | ENGINEERING
The Fu Foundation School of Engineering and Applied Science

# Power and thermal modeling

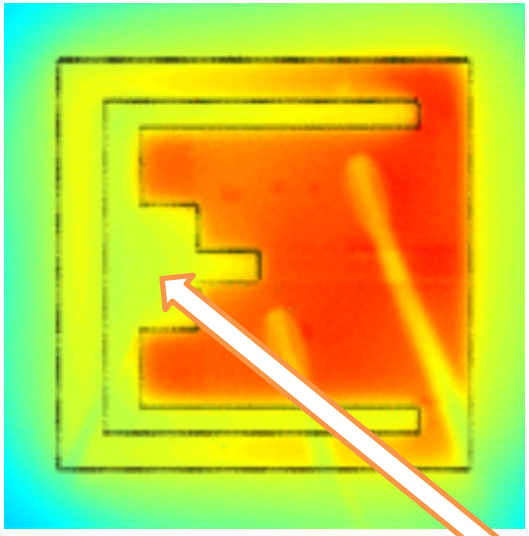

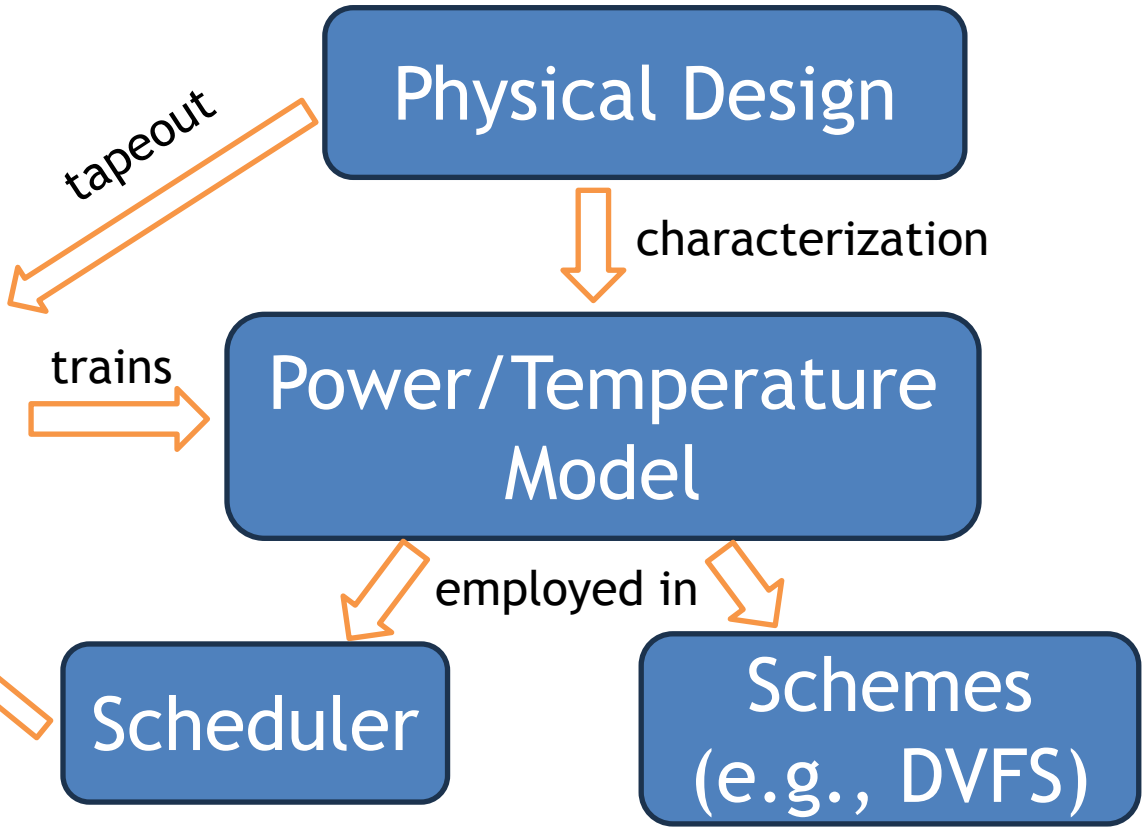


Management schemes require **accurately** modeling the power and temperature with sufficient **abstraction.**

# Power and thermal modeling

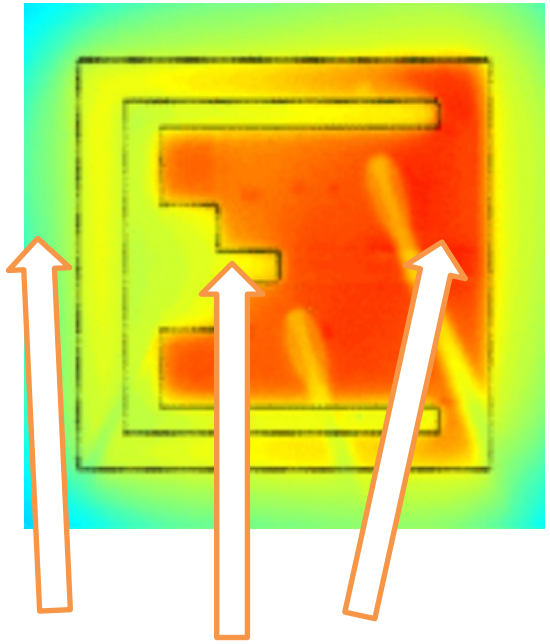

Performance events (PMU)

Compute power from activity using weights!

$$P = \sum \beta_i * \text{count}\{event_i\}$$

Train based on a "representative" workload

# Design-time approaches

Per-cycle power trace (not runtime)

Accuracy

*Ground truth*

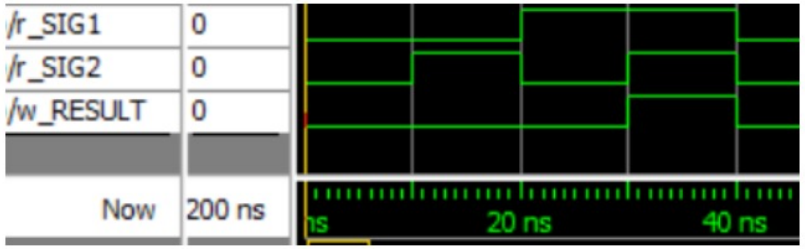


$$P = \sum \beta_i * \text{count}\{event_i\}$$

Power counter-based models

Performance counter-based models

Training/Inference Cost

# What we propose

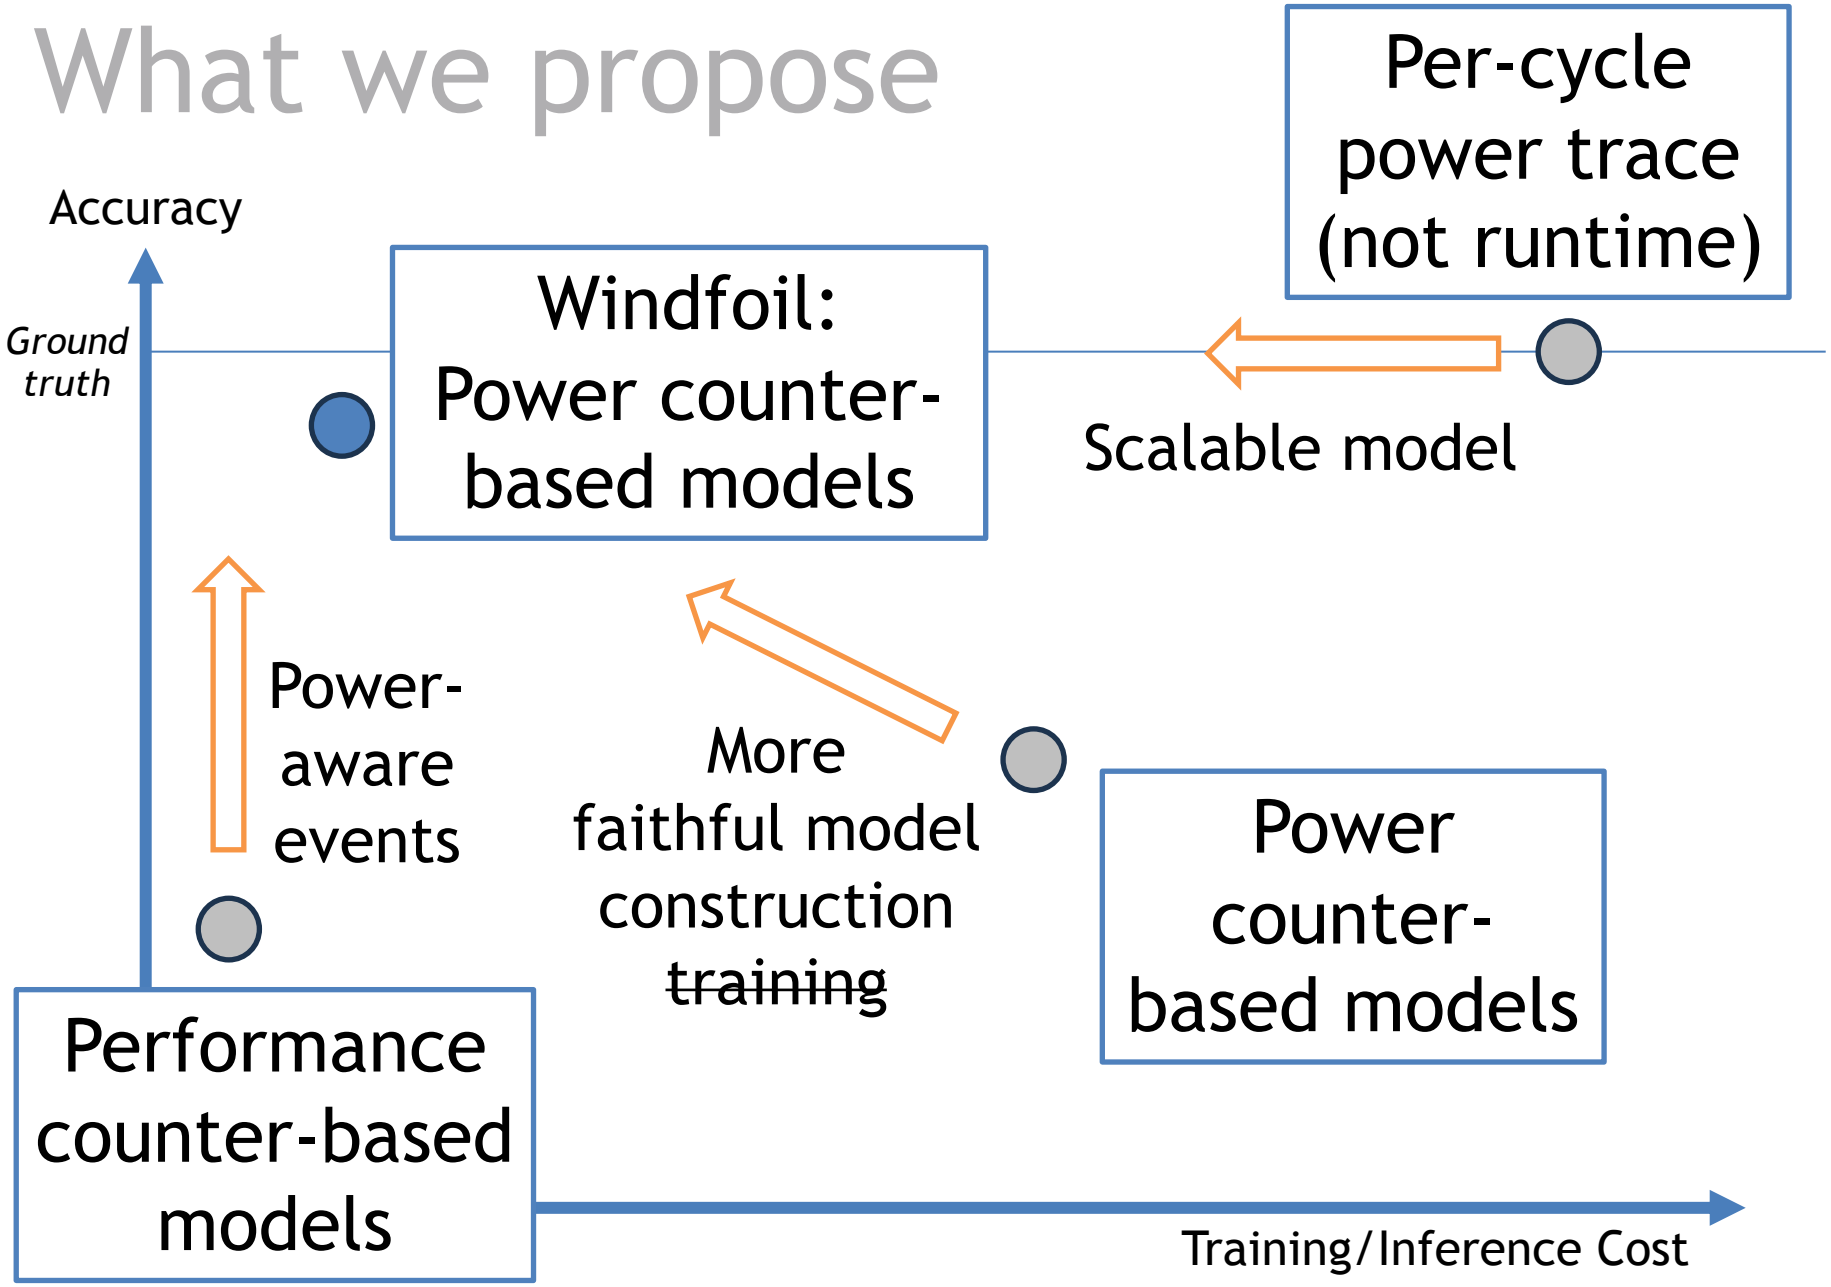

# Myth – event selection

Workload-based correlation makes coarse abstractions without semantics

Top-level modules ⟷ ? ⟷ Power distribution

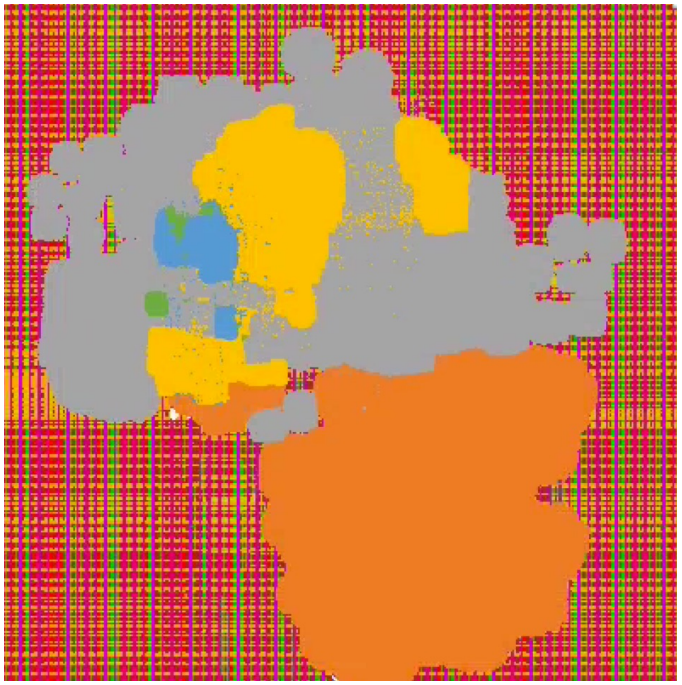

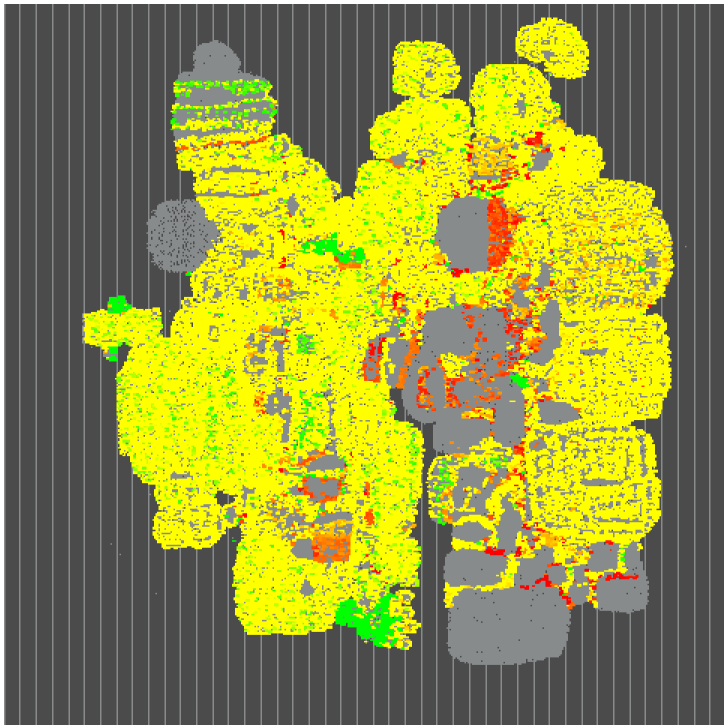

# Myth – model complexity

"Representative" workloads miss access patterns influence on power

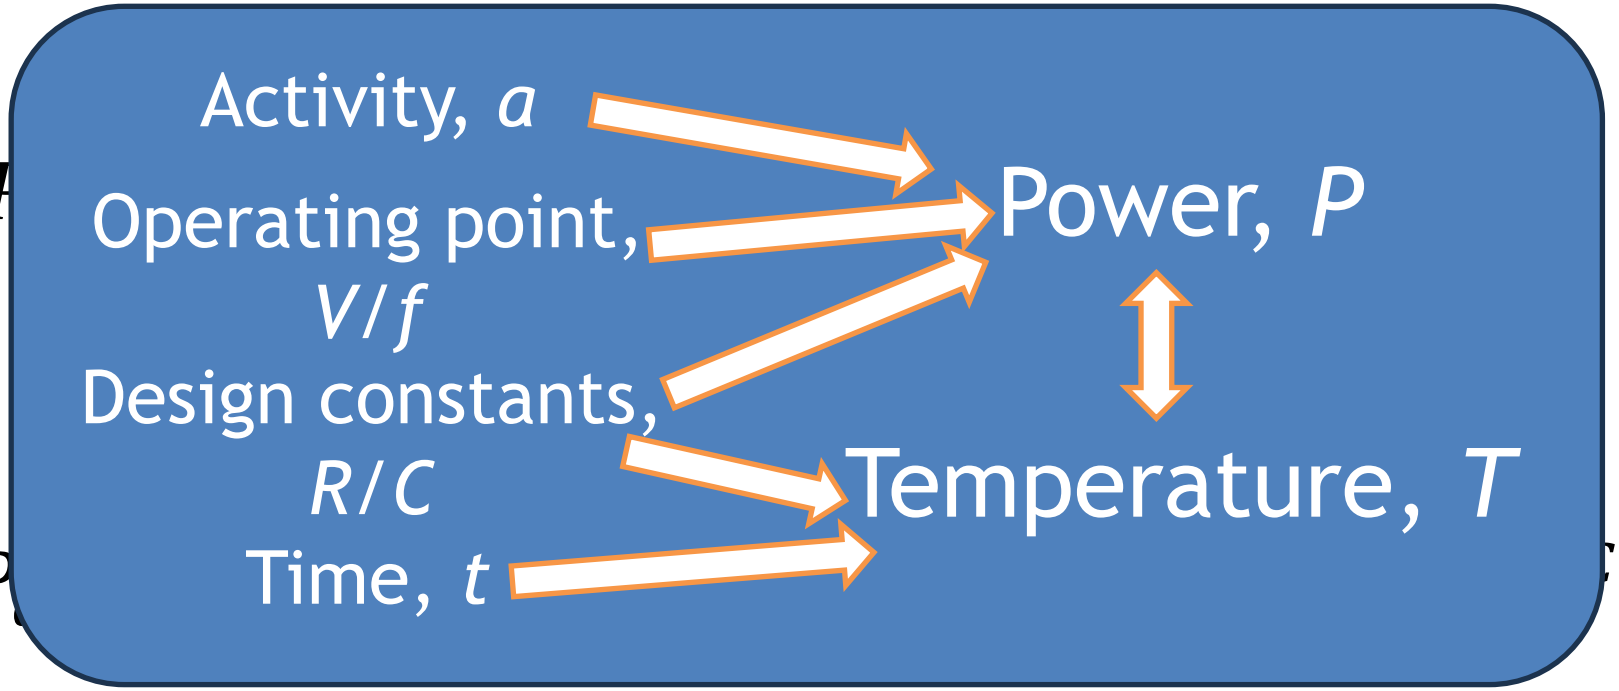

# Contributions

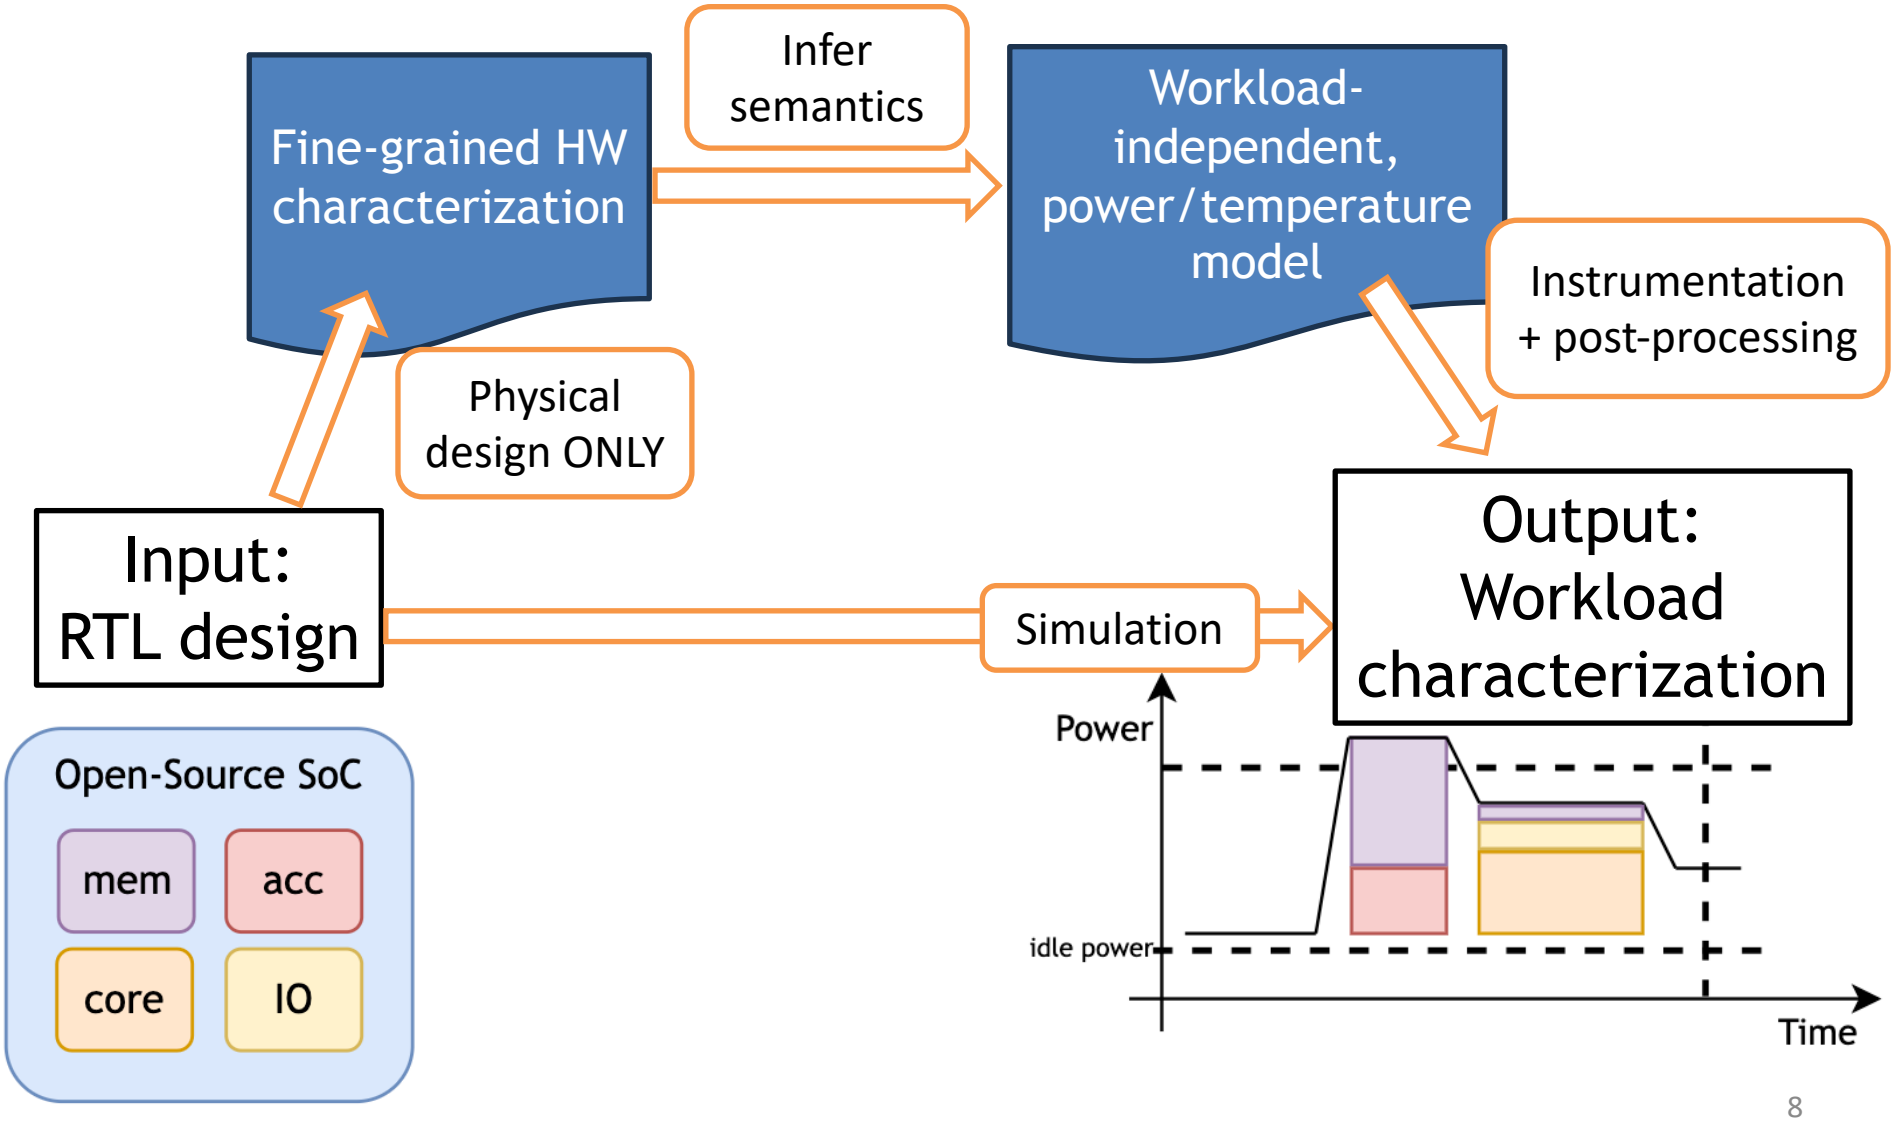

# Model – power

**Per-gate activity**

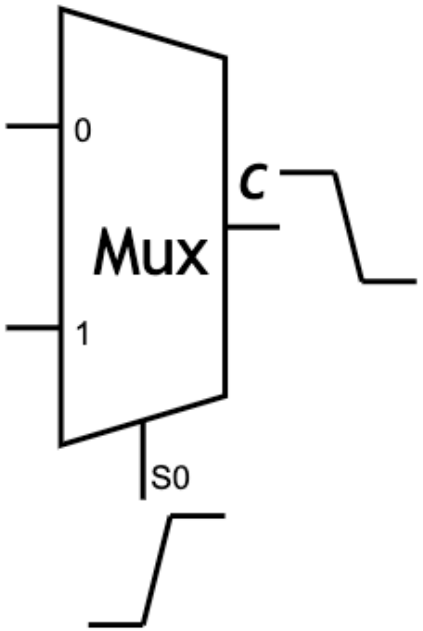


$P_c$

$\Pr(c|s_0)$

Composes into

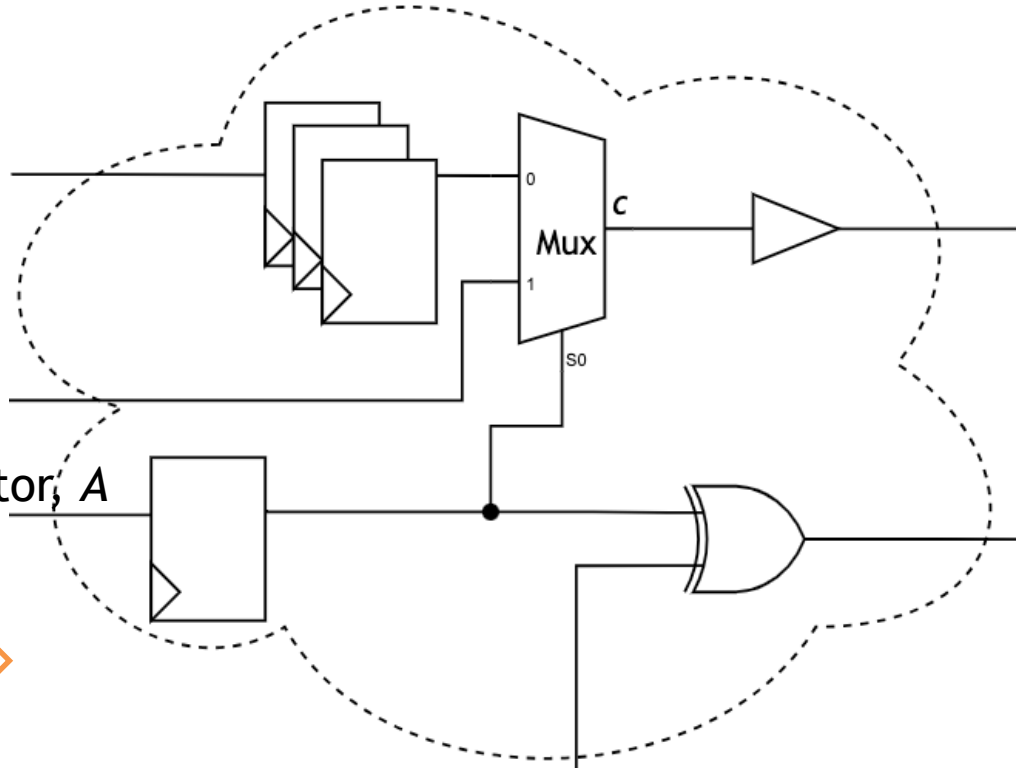


**Weighted "atomic" power**

$$P = P_A * 1 + P_c * \Pr(c|s_0) + \cdots$$

# Model – temperature

Degree of connection governs thermal transfer.

e.g., Degree = 5

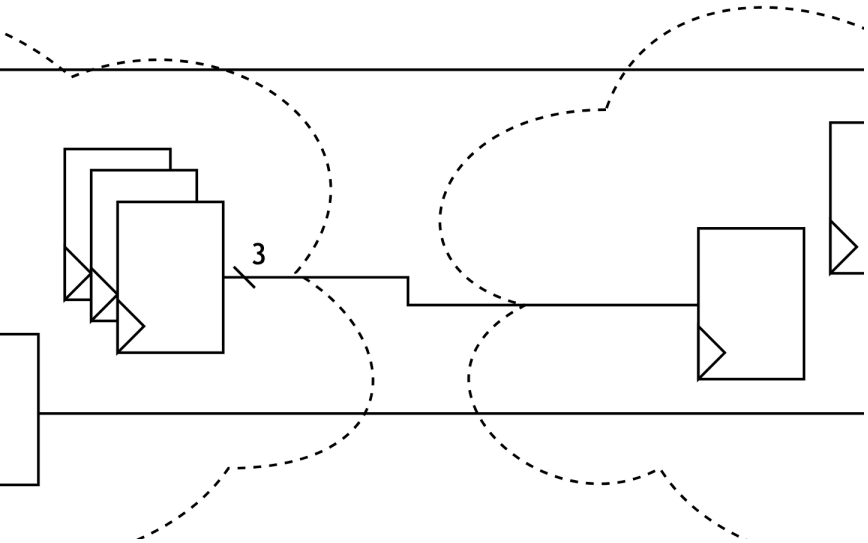


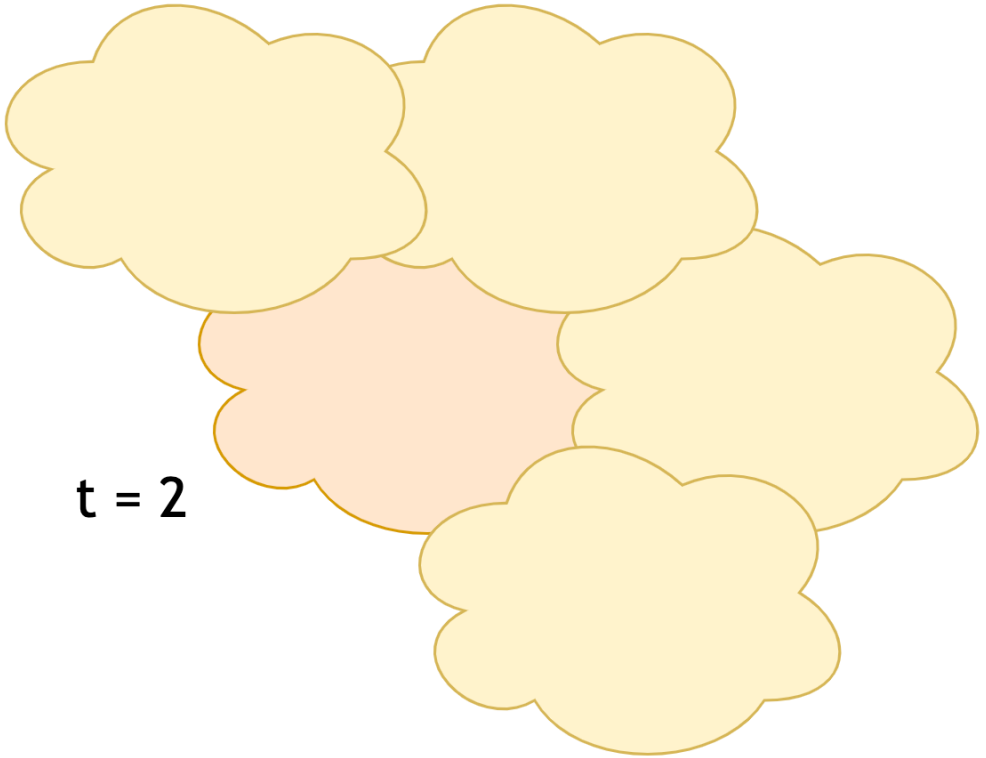

# Results – Validation

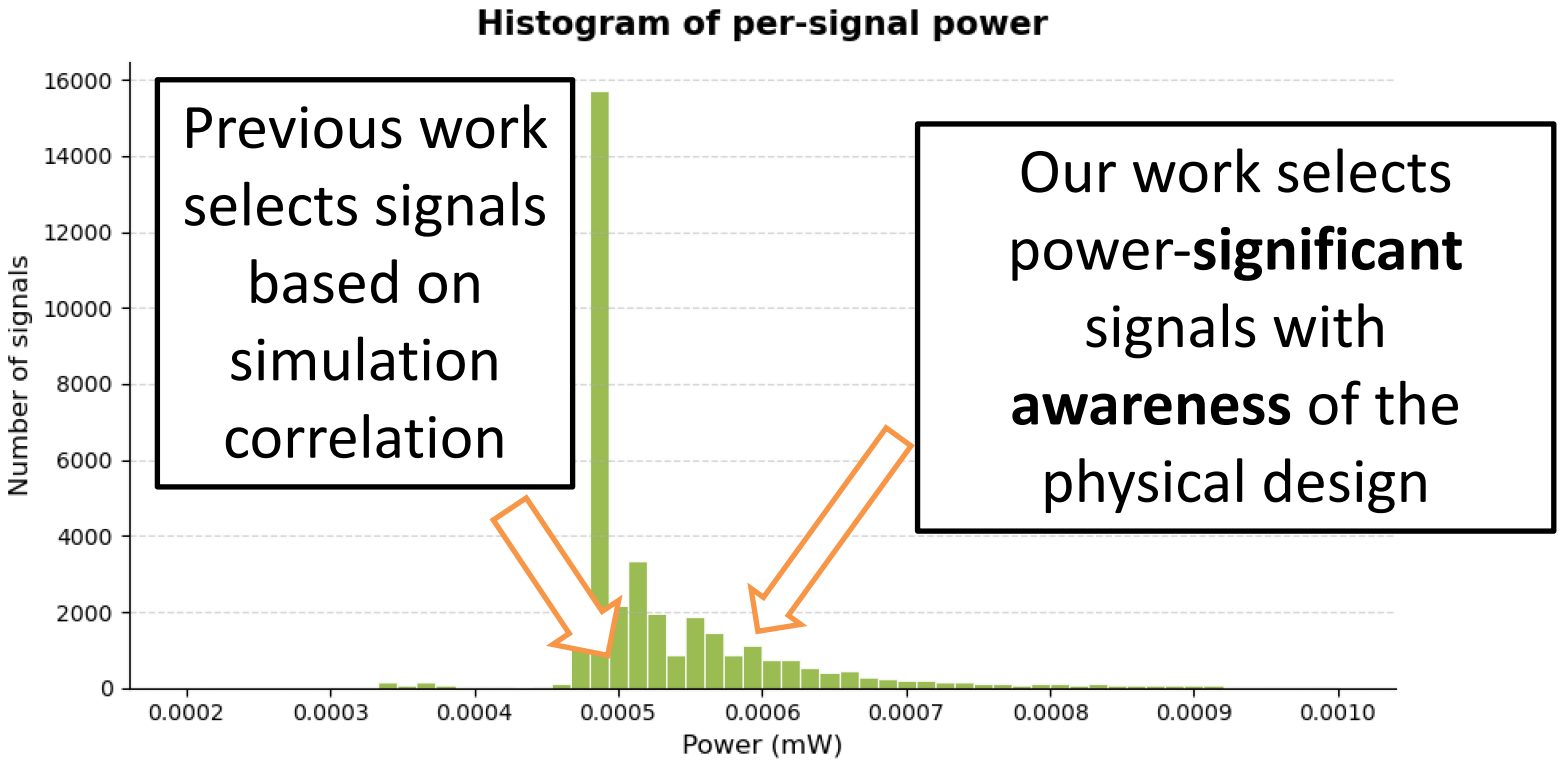


**Hypothesis**: our clustering provides a better abstraction of the power distribution

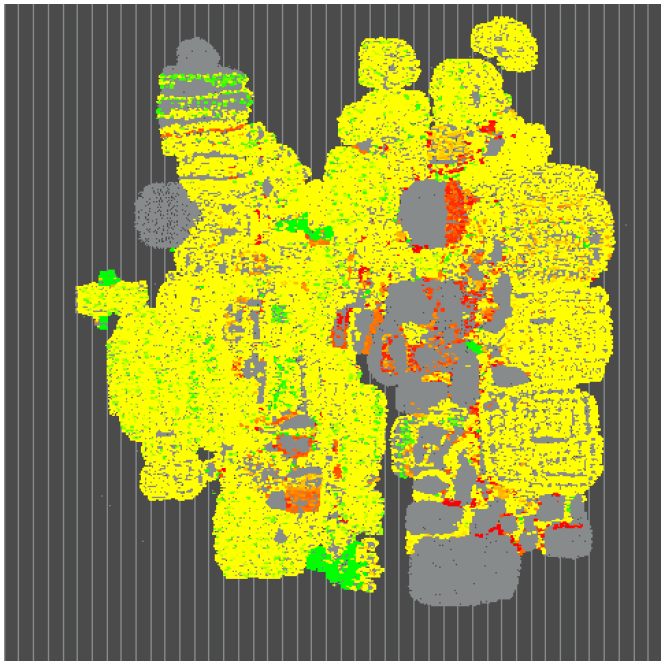

**Windfoil**: Workload-independent, physical design-aware, power and temperature model

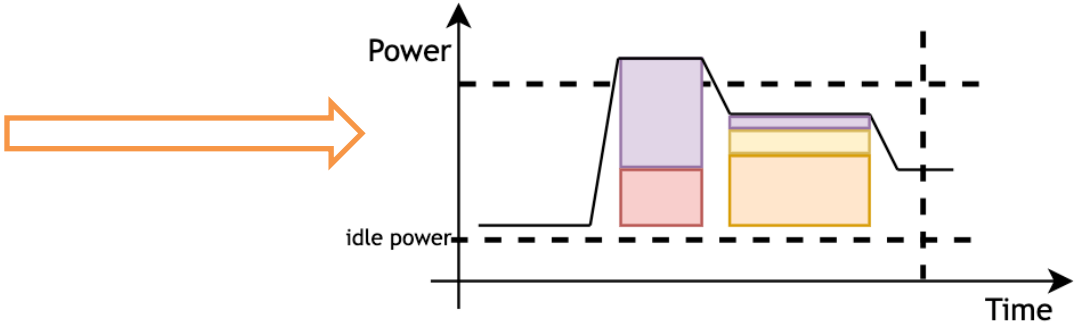


Open to discussion!

- Smarter clustering
- Stronger temperature model

Email: michael.grieco@columbia.edu
Github: michaelg29.github.io

# Backup slides

# Challenge – model accuracy/cost

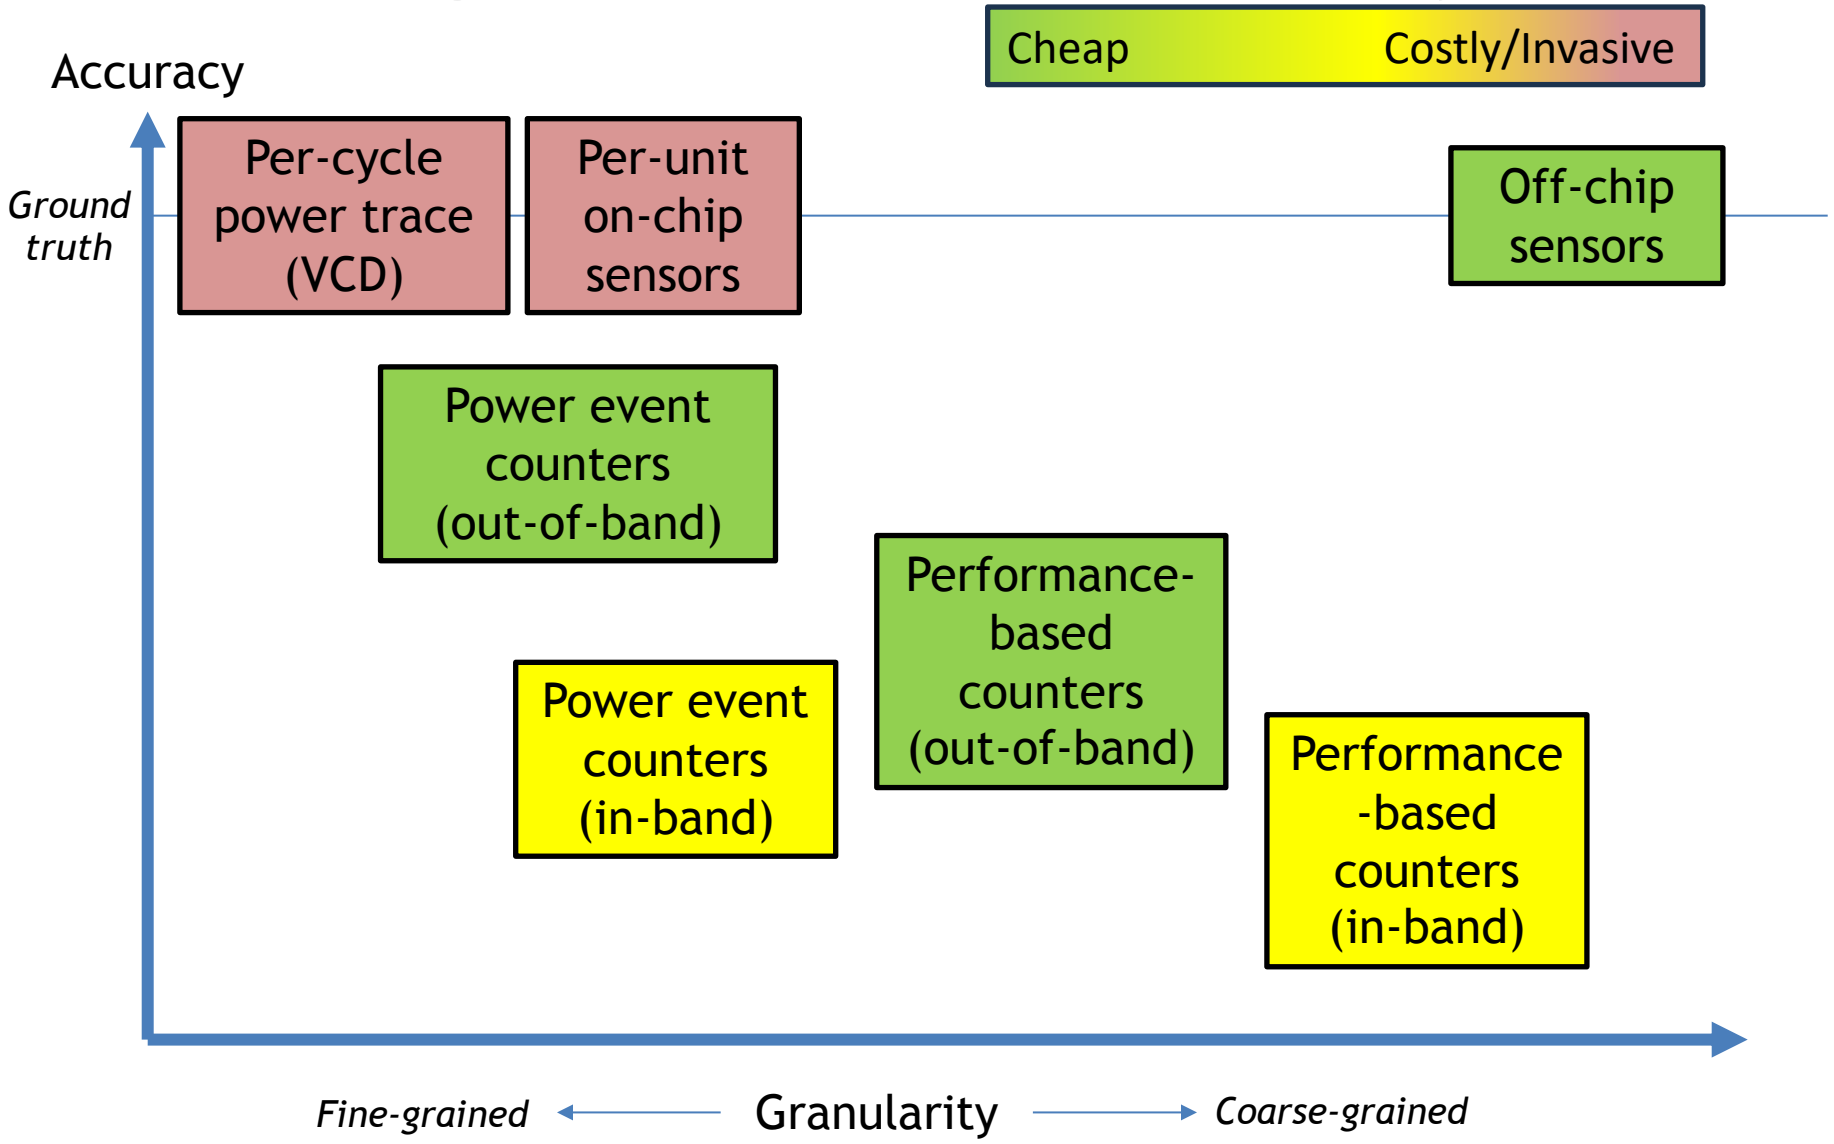

# Myth – instrumentation

Dependence on instrumentation

- "Sensors do not affect results"
- "Sampling rate is non-significant"

=> Heisenbugs

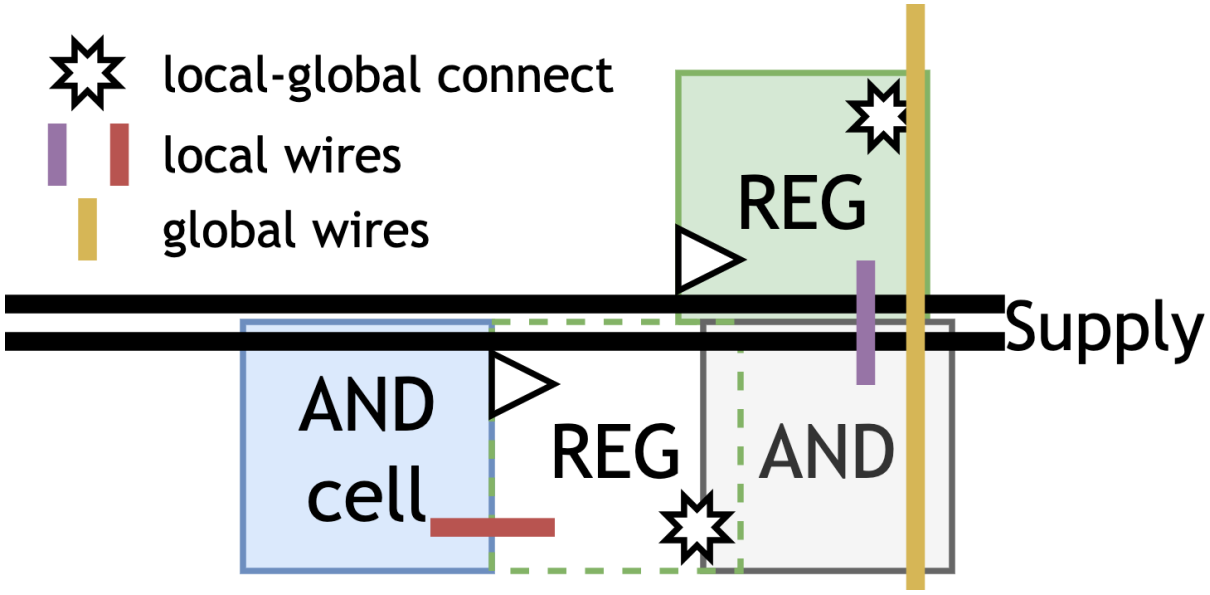


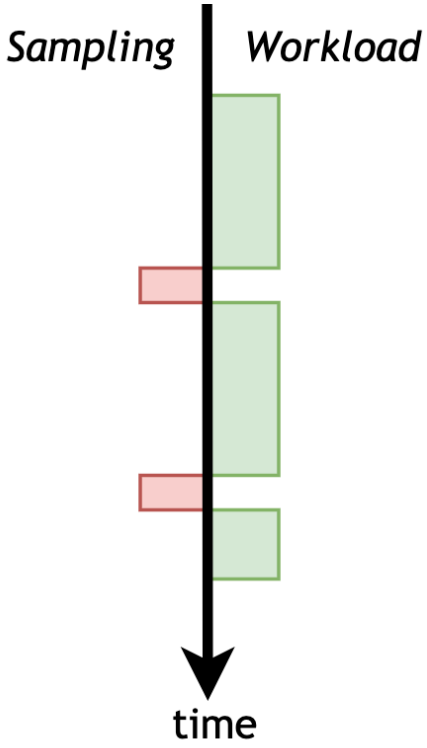

# Myth – model training

Dependence on training workload

- "Targeted microbenchmarks"
- "Siloed events"
- "Representative"
- "No interference"

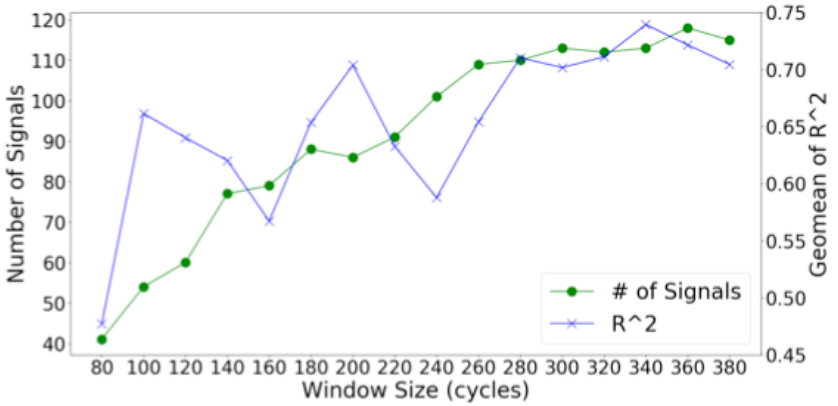


**Figure 7: The number of selected signals and the geometric mean of $R^2$ across module-level power models for different window sizes**

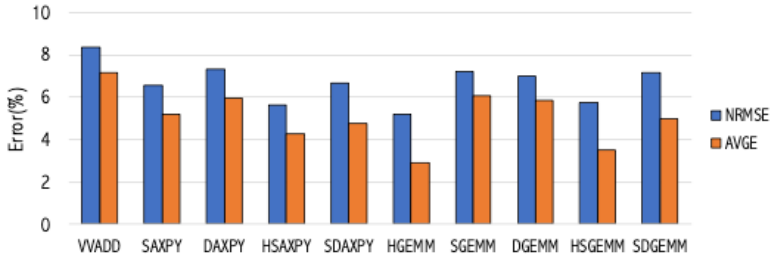


**Figure 8: Power prediction errors for microbenchmarks**

Kim et al., MICRO'19

# Implementation