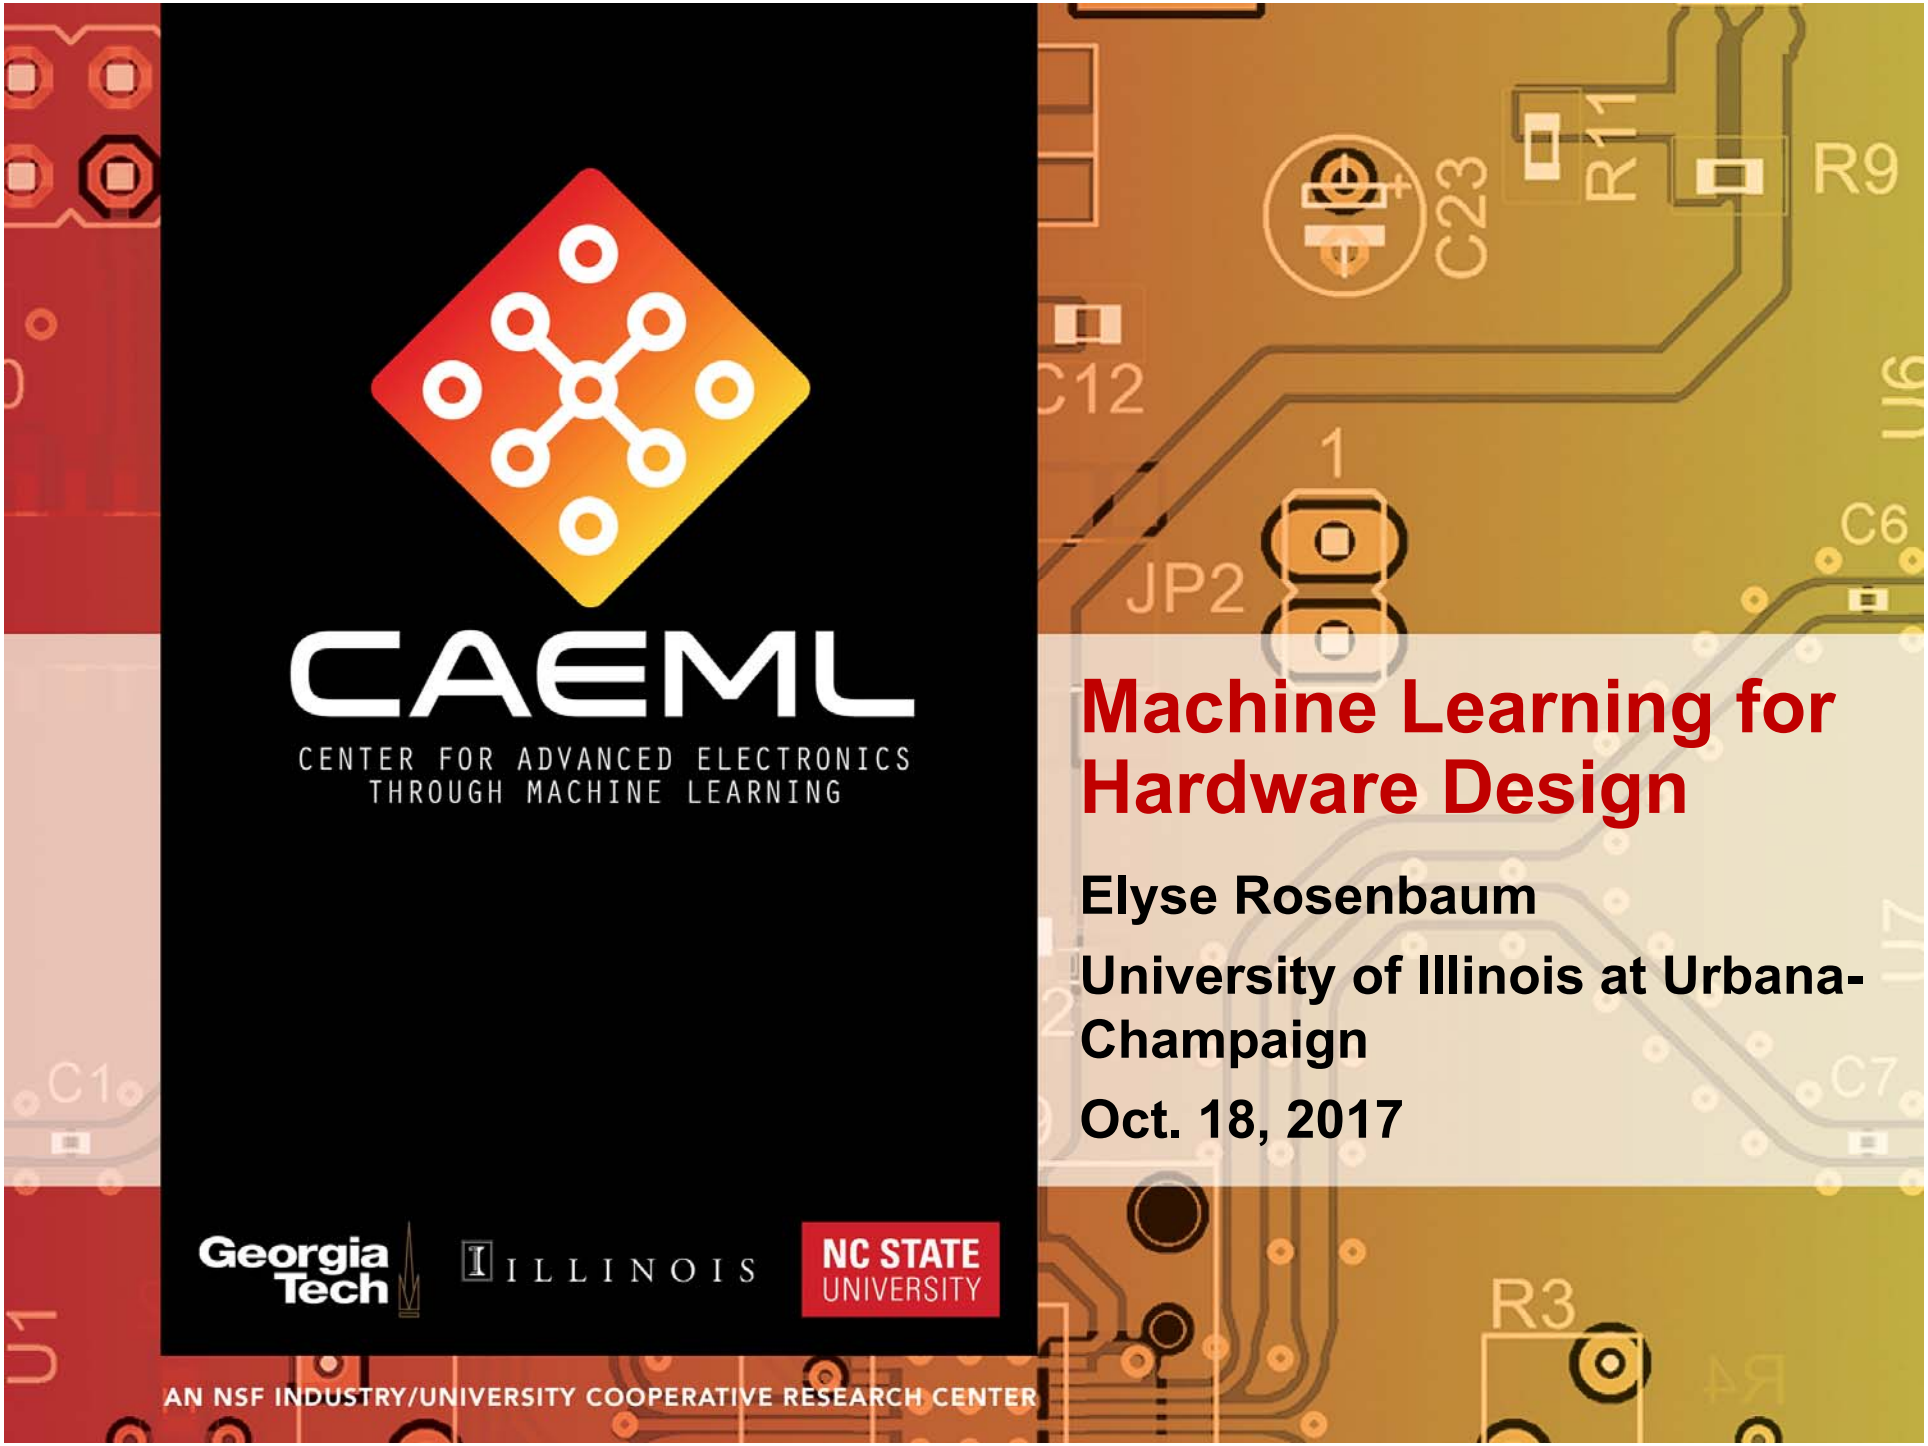

Machine Learning for Hardware Design

Elyse Rosenbaum

University of Illinois at Urbana-Champaign

Oct. 18, 2017

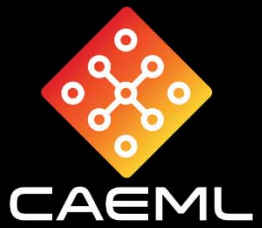

Questions, Questions, Questions

1. How can design productivity be improved?
2. What is machine learning?
3. How do hardware designers utilize machine learning?
4. What is CAEML?

To Improve Productivity…

- Avoid respins
  - o You're not working on the next product while you are busy fixing the current one
  - o Short-term and long-term financial cost
  - o Achieved by comprehensive (yet quick) simulation-based design verification
- Execute an optimal design rather than a "good enough" design
  - o Metrics may include manufacturing cost, power dissipation, performance, reliability
  - o Achieved through model-based design

… Better Models are Required

- Many of the observed failures during qualification testing are the direct result of an insufficient modeling capability
  - Sources of such failures include mistuned analog circuits, signal timing errors, reliability problems, and crosstalk [1]
  - Variability cannot be modeled in a manner that is both accurate and computationally efficient
- Simulation-based design optimization has had only limited success
  - Simulation "in-the-design-loop" often too slow and leads to impractical designs
- Proposal: use machine learning algorithms to generate behavioral models with the needed accuracy and efficiency

[1] Harry Foster, "2012 Wilson Research Group Functional Verification Study," http://www.mentor.com/products/fv/multimedia/the-2012-wilson-research-group-functional-verification-studyview

A Definition of Machine Learning

The application of ***statistical learning theory*** to construct accurate predictors (*f*: inputs→outputs) from data

- Today, this is feasible even with a large number of inputs ("features") due to the availability of powerful computing machines, and
  - Optimization solvers
  - Parallel programming

Who Generates the Models?

Hardware Designer

Computer Scientist

Not!

ML requires *domain expertise.*

Domain Expertise

- For electronics modeling, hardware designer understands better than anyone else
  - How the component fits into the larger system
  - How the model will be used (e.g., analysis, simulator)
  - What constitutes a good model
  - Available training data
  - Physics underlying the component's behavior
- Analogy: Within the field of deep learning
  - Convolutional neural network used for image recognition
  - Long short-term memory network (LSTM) used for natural language processing

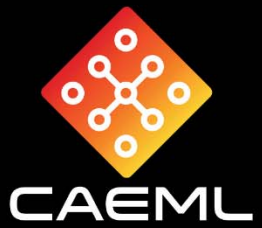

Using ML for Hardware Design

- Here: brief, introductory examples
  - o Bayesian optimization to find region of design space in which cost function is minimum
  - o IP-obscuring models for non-linear circuits
  - o Model-based design with surrogate models
- Eight technical presentations will follow

Thermal Design Optimization for 3D-IC

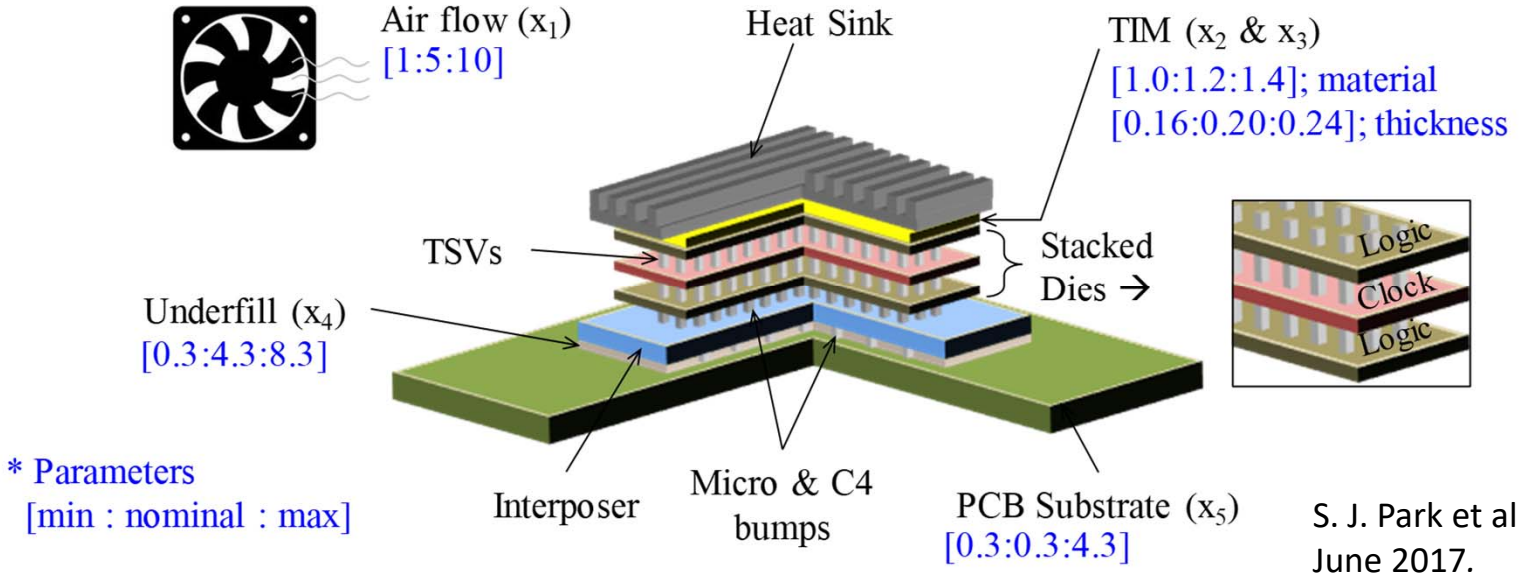

S. J. Park et al., *IEEE Trans. VLSI,* June 2017.

- Clock Skew is affected by Temperature (magnitude and gradient)
- Temperature is controlled by FIVE features, which have constrained values
- Objective is therefore to TUNE the feature values to minimize Skew
- More generally, seek to find $X_{opt} = \underbrace{argmin}_{X}((f(X))$. Accurate modeling of f(X) needed only near minimun
- Use ML-based Bayesian Optimization

Full Design Space Exploration too Costly

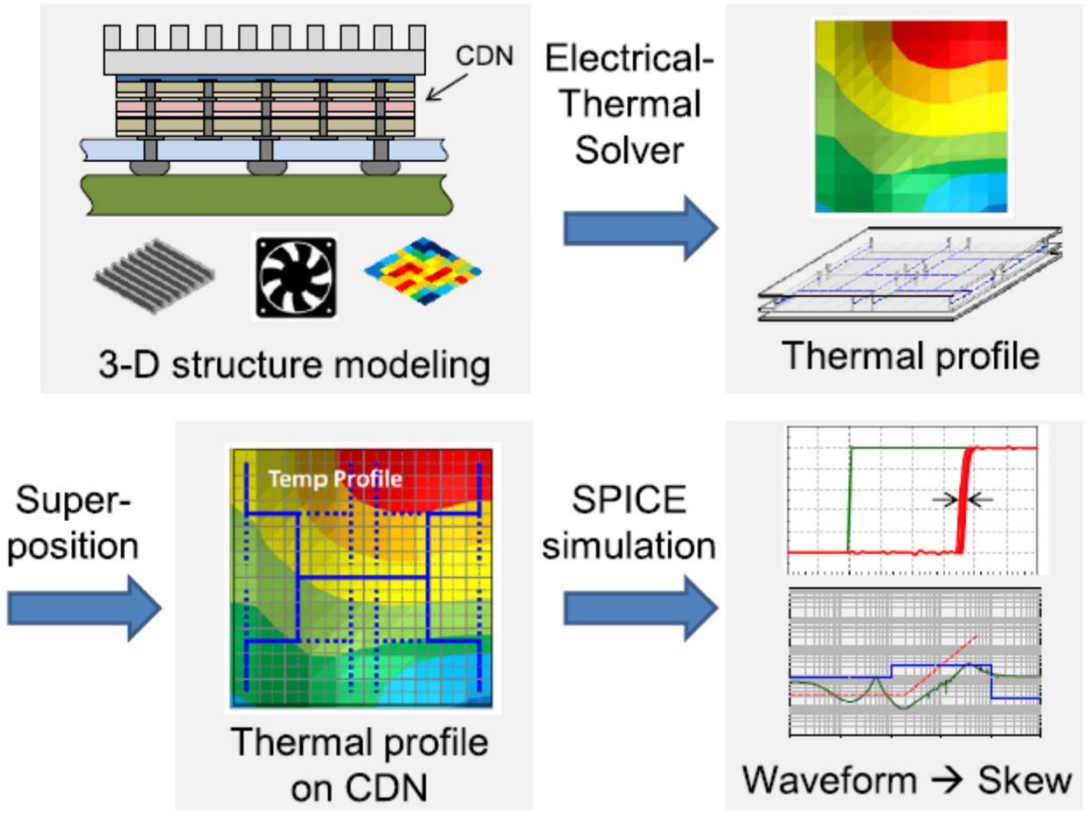

- 3D (finite volume) simulations + SPICE-type circuit simulation
- Need to limit the number of designs that are simulated

IP-obscuring Model for Transient Simulation

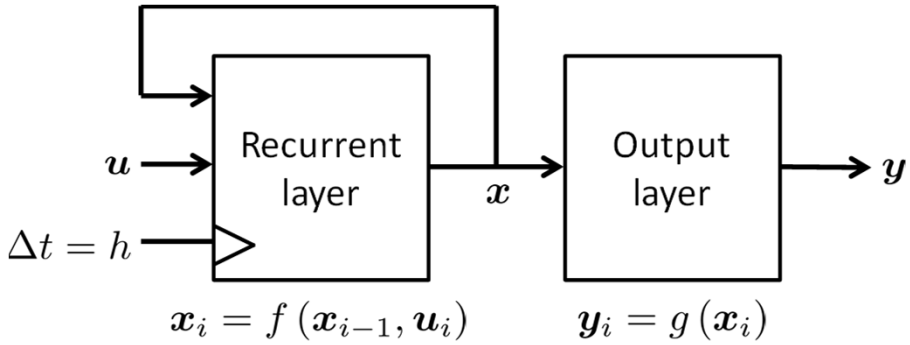

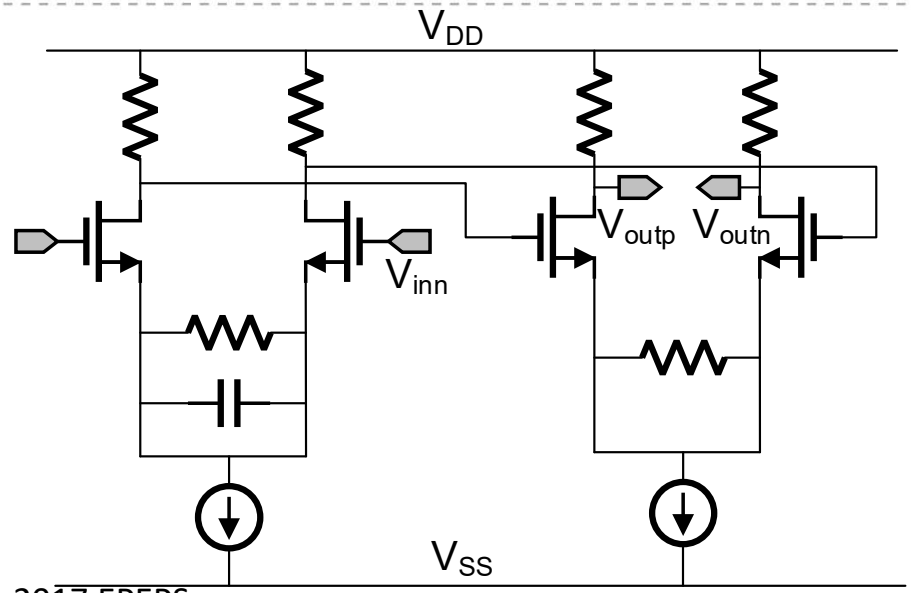

Z. Chen et al., 2017 EPEPS.

- RNN known to estimate non-linear systems that can be represented by a state-space model
- Developed a continuous-time RNN for circuit simulation

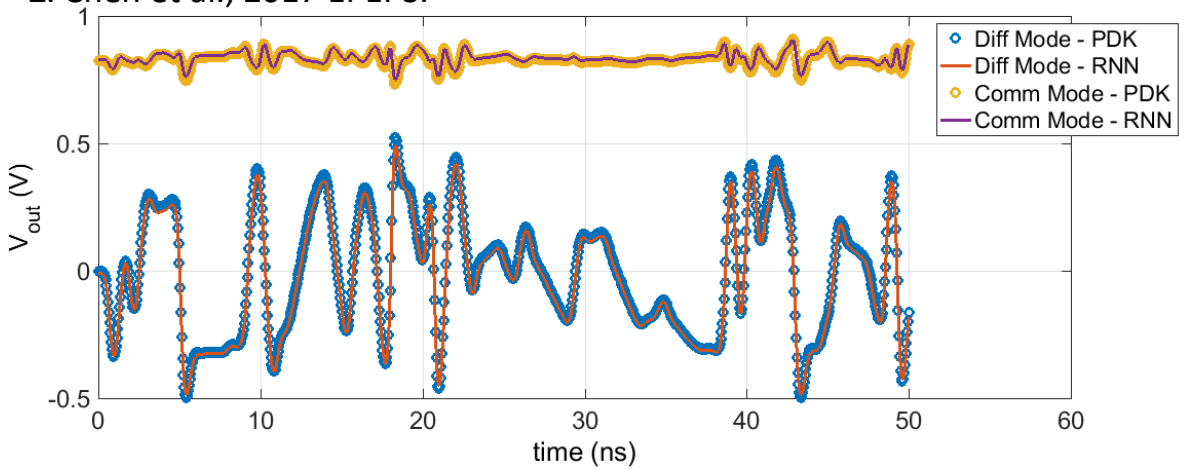

Surrogate Model Based Circuit Design

Goal: Identify the optimal set of calibration knobs for this design

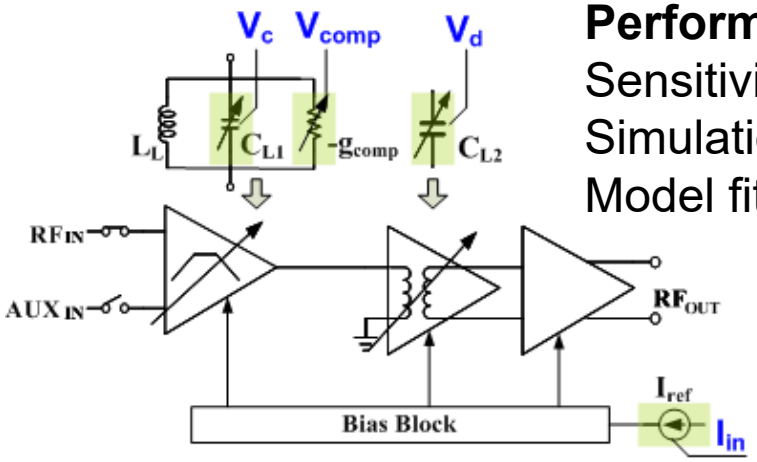

Perform:
Sensitivity Analysis
Simulation-based sampling
Model fitting

Multi-dimensional response surfaces

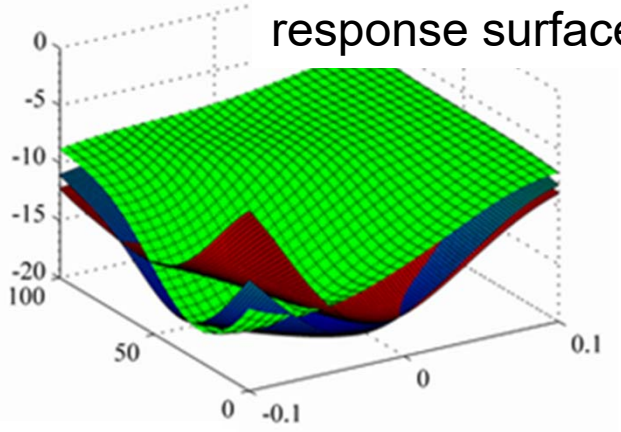

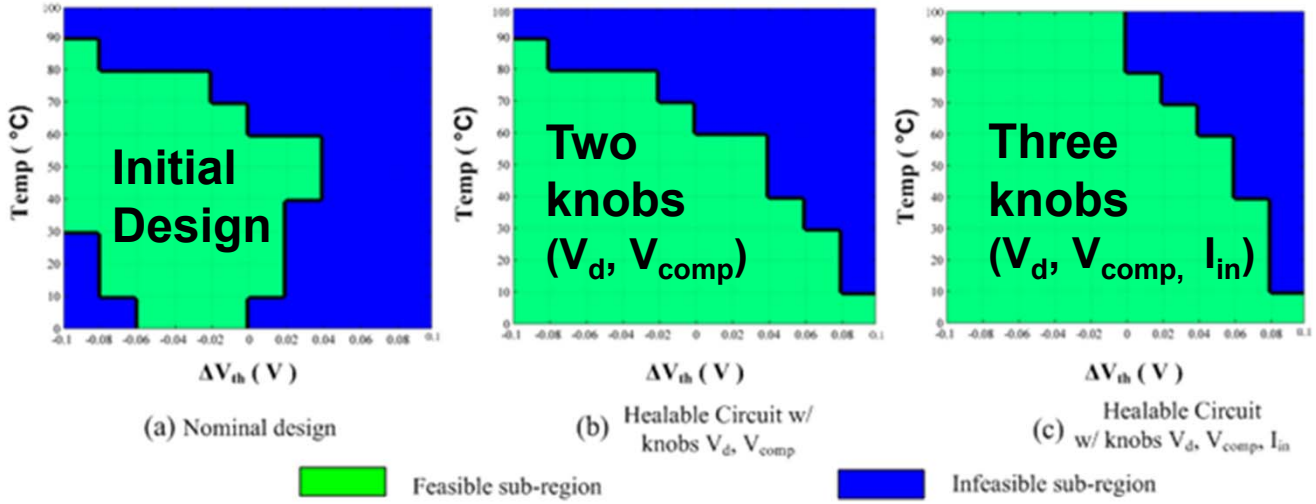

7D : 3 responses, 4 "knobs"

Ref: P. Franzon (CAEML / NCSU)

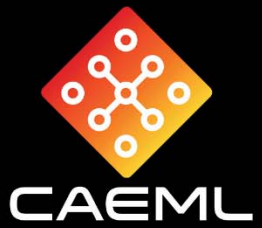

Center for Advanced Electronics through Machine Learning (CAEML)

- Mission:
  - To enable fast, accurate design and verification of microelectronic circuits and systems by creating machine-learning algorithms to derive models used for electronic design automation
- An NSF I/UCRC
  - Industry/University Cooperative Research Center
  - Three sites
    - University of Illinois at Urbana-Champaign (lead site)
    - Georgia Tech
    - North Carolina State University
  - Site directors: Rosenbaum (UIUC), Franzon (NCSU), Swaminathan (GT)
  - Other faculty: Cangellaris (UIUC), Kiyavash (UIUC), Raginsky (UIUC), Schutt-Aine (UIUC), Davis (NCSU), Floyd (NCSU), Ji (GT), Lim (GT), Raychowdhury (GT)

Core Elements of I/UCRC Program

- Precompetitive research
- Membership fee for participation
  - Industry funds are pooled and used to support research
    - Extremely low overhead charge (10%)
  - NSF pays administrative expenses
  - Members have rights to IP
- Semi-annual meetings for stakeholders
- Industry Advisory Board (IAB)
  - Selects projects via voting and consensus
- Research results are shared with all members
- Workforce development

Current Members

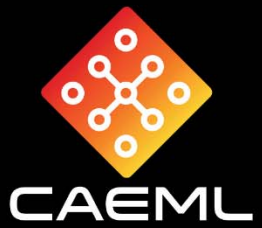

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

LOCKHEED MARTIN

XILINX®

CISCO

cadence™

QUALCOMM®

IBM

SAMSUNG

Sandia National Laboratories

SYNOPSYS®

Hewlett Packard Enterprise

NVIDIA®

Research Portfolio

- Theory, Devices and Systems
  - Modular machine learning
  - High-speed links
  - Power delivery
  - System-level ESD
  - IP reuse
  - Design rule checking
- New in 2018: Trusted platform design; FPGA compilation strategy; Early detection of hardware failure