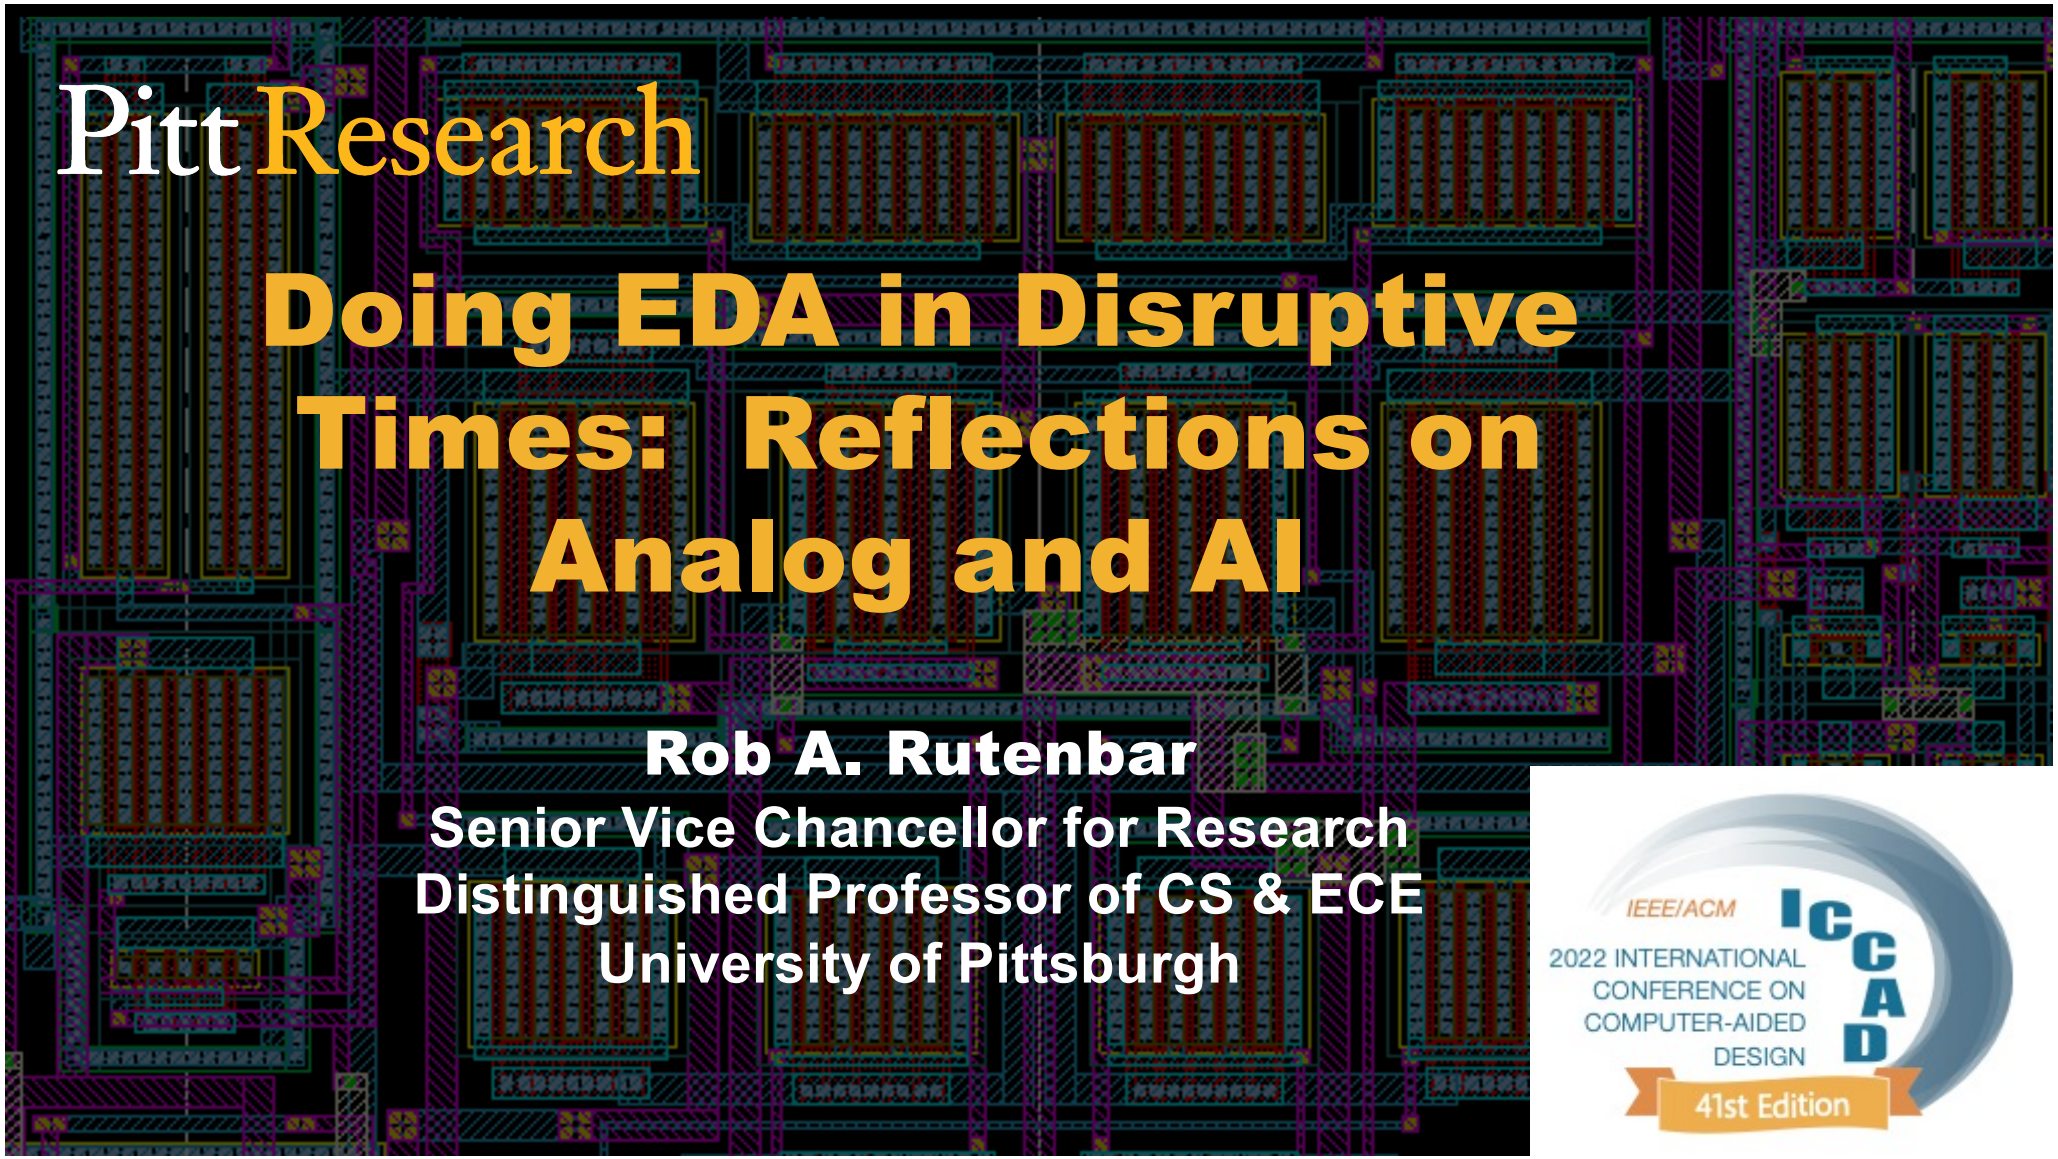
Pitt Research
Doing EDA in Disruptive Times: Reflections on Analog and AI
Rob A. Rutenbar
Senior Vice Chancellor for Research
Distinguished Professor of CS & ECE
University of Pittsburgh
IEEE/ACM
ICCAD
2022 INTERNATIONAL CONFERENCE ON COMPUTER-AIDED DESIGN
41st Edition

A Talk in 3 Disruptions (3 Chips)

1

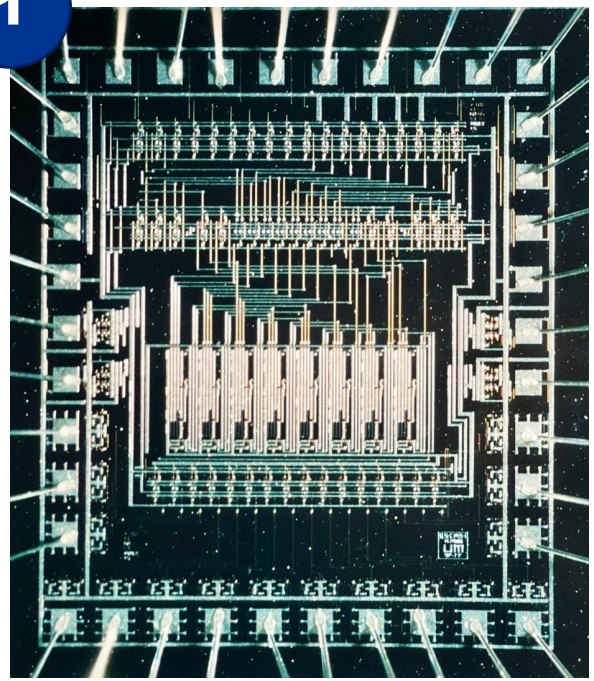

2

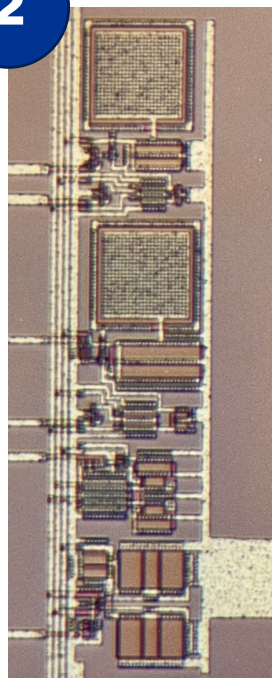

3

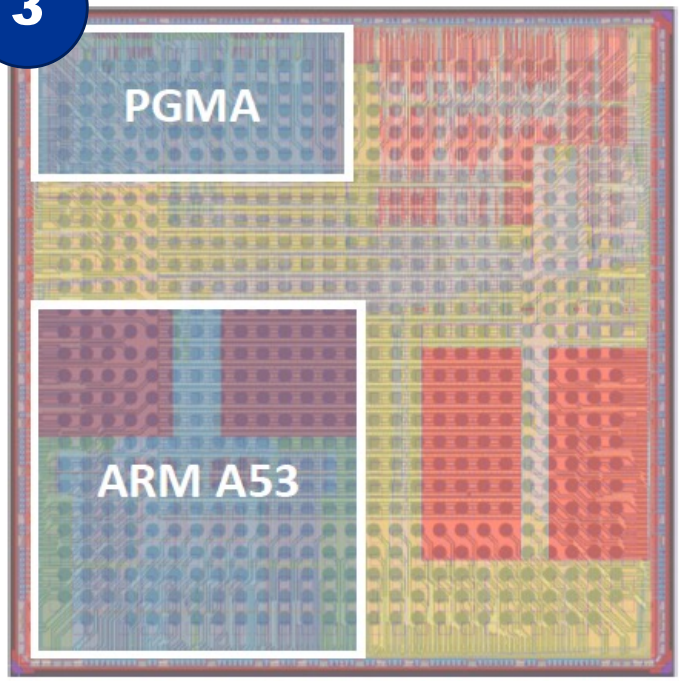

Pitt Research

Disruption #1

1

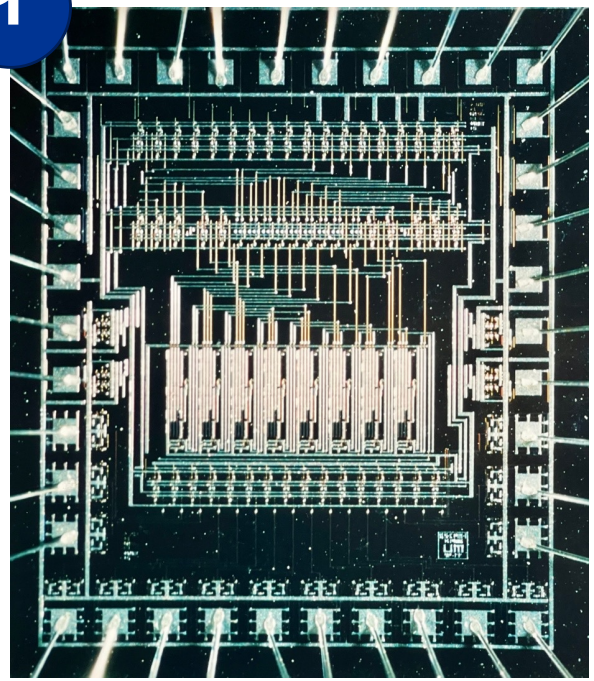

~1980

- λ = 4μm
- Depletion-mode NMOS
- Single metal
- 503 transistors
- 200KHz
- Fabbed by General Motors

Pitt Research

Disruption #1

- Multiply-Add Circuit

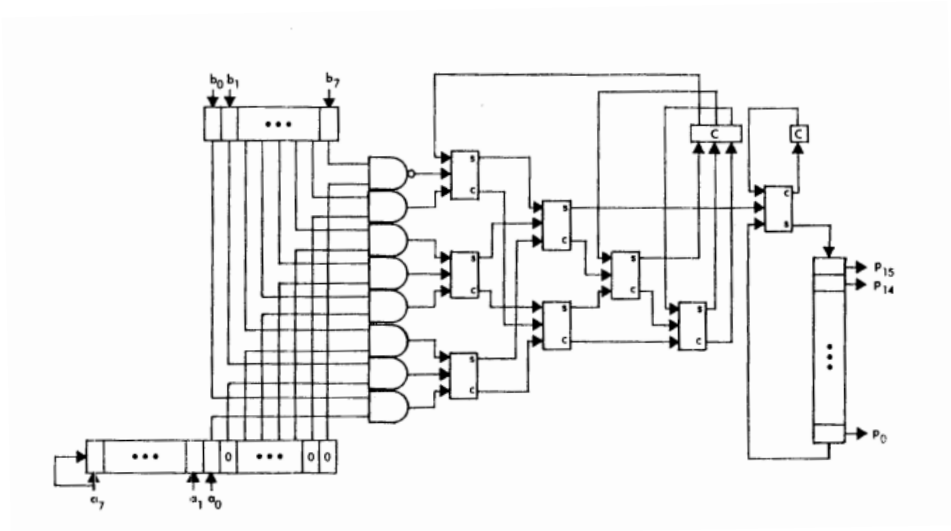

< 100 gates

- Technology

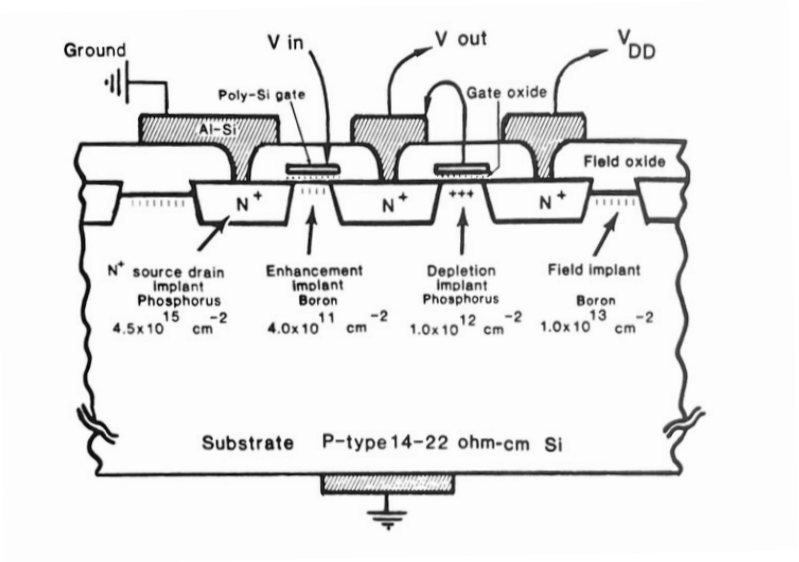

4000nm node (~½-CMOS)

Pitt Research

Disruption #1: Mead & Conway

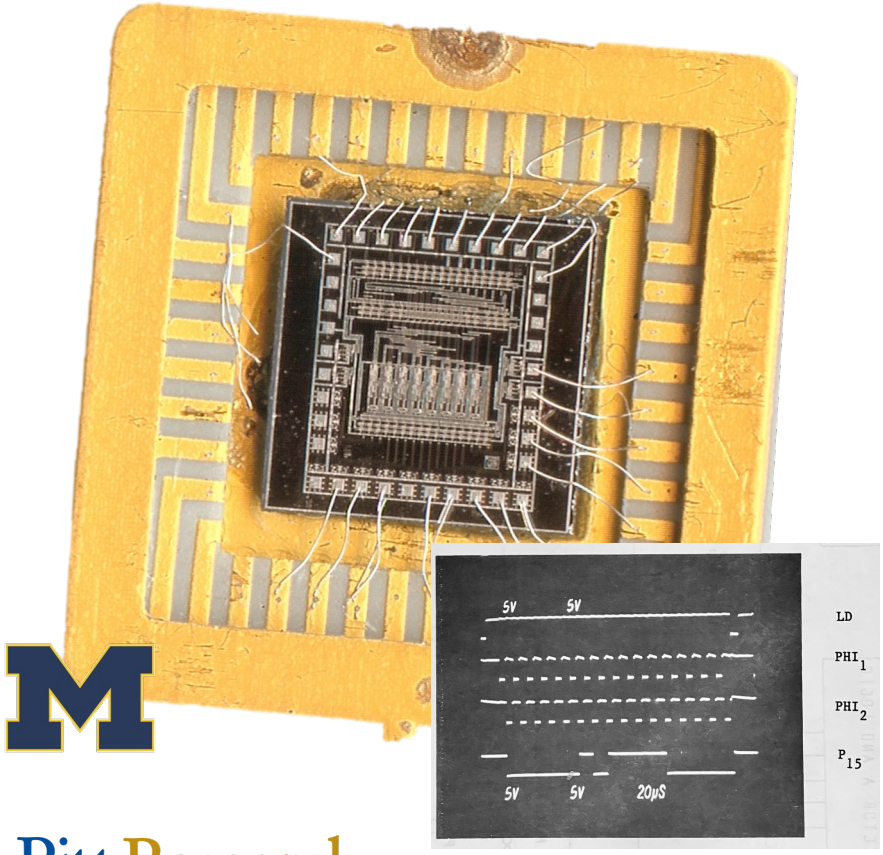

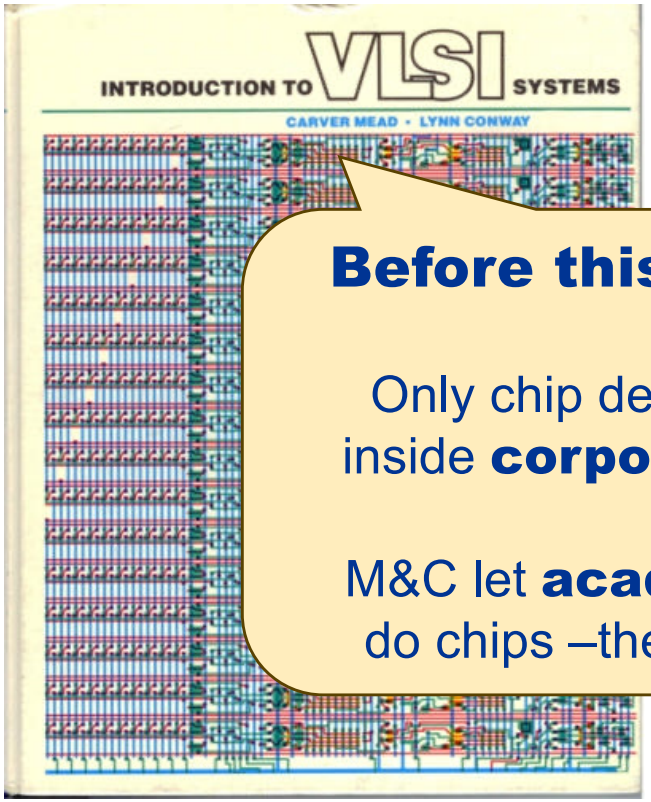

Before this book:

Only chip designers inside **corporations**

M&C let **academics** do chips –then **EDA**

Pitt Research

~My 1980 Design Methodology & Toolchain

- Ladder
- Line-printer
- Masking tape
- Sharpie
- DRC checker
- BSD Unix ”vi” editor

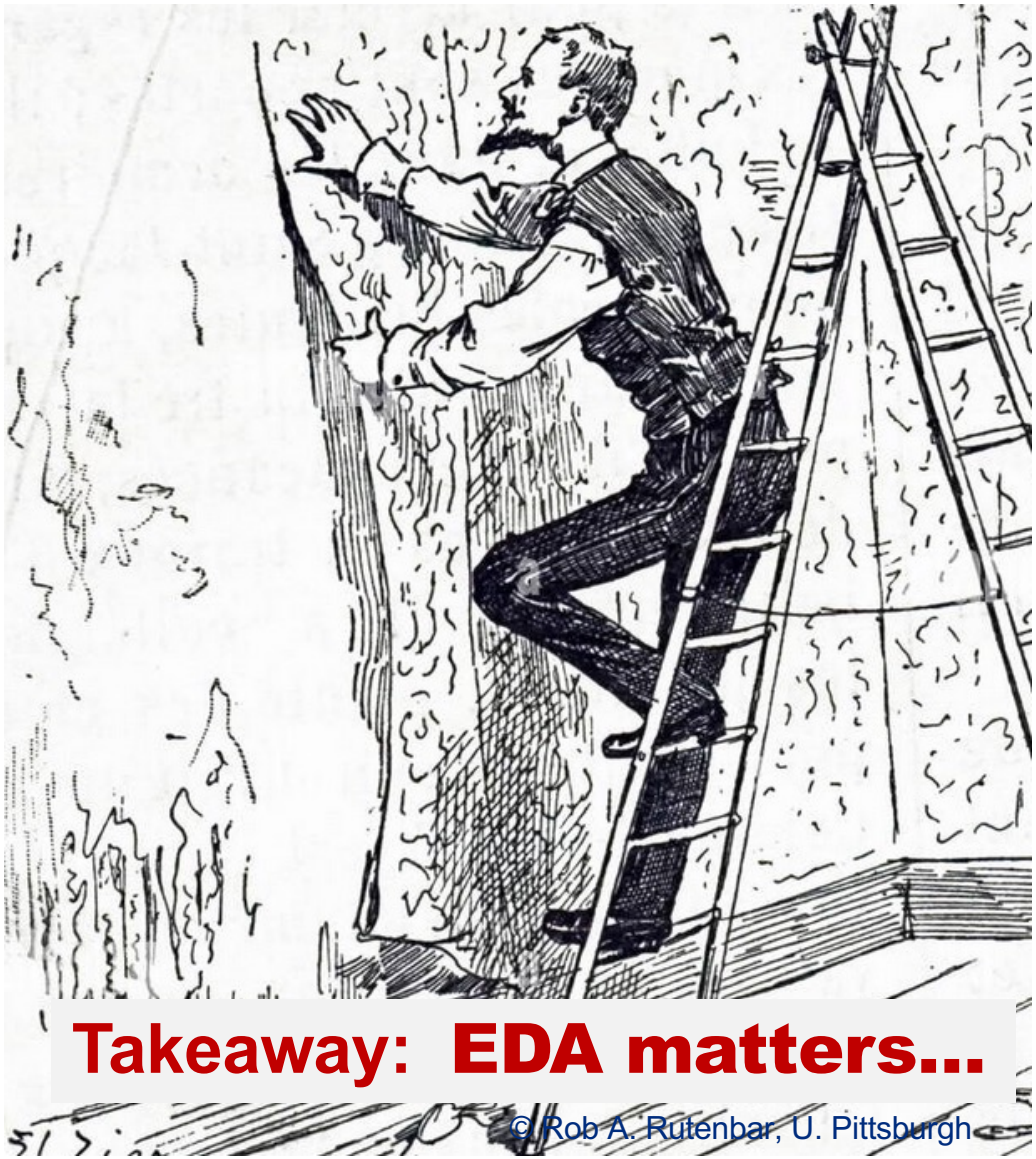

Takeaway: EDA matters…

Disruption #2

2

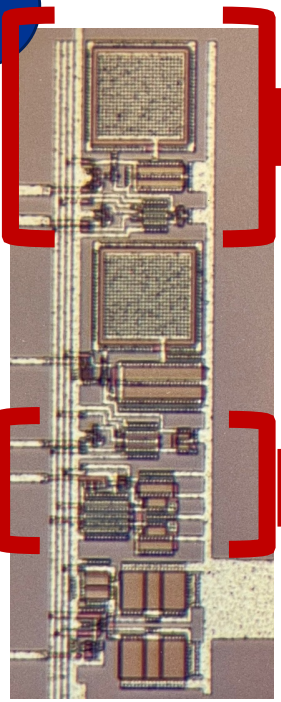

OTA opamp

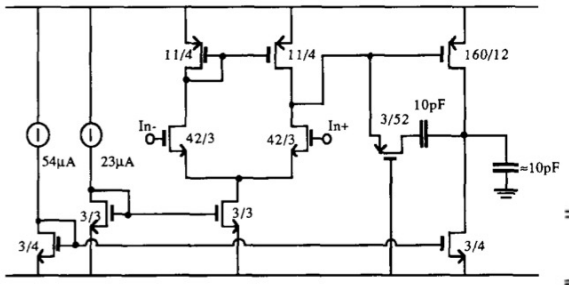

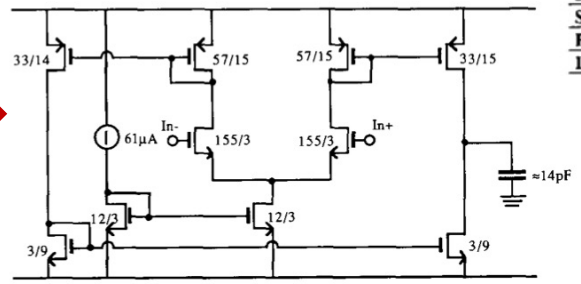

2-Stage opamp

INPUT SPECIFICATIONS AND RESULTS FOR FABRICATED OP AMPS

Parameter	Units	Circuit OTA-1 Spec	OTA-1 SPICE Results	OTA-1 Chip Results	Circuit 2Stage-1 Spec	2Stage-1 SPICE Results	2Stage-1 Chip Results
Gain	dB	≥ 60	60	69.5	≥ 45	45.2	48.4
Unity Gain Freq.	MHz	≥ 1.0	1.25	0.993	≥ 2.8	2.52	2.6
Phase Margin	Deg.	≥ 45	46	43	≥ 60	82	85
Load Capacitance	pF	10	10	~ 14	14	14	~ 14
Supply	V	± 2.5	± 2.5	± 2.5	± 2.5	± 2.5	± 2.5
Slew Rate	V/µS	≥ 2.0	2.1	-1.6/+2.0	≥ 2.5	2.42	-3.0/+4.5
Power	mW	≤ 3.0	0.63	0.62	-	0.56	0.56
Input Offset	mV	-	6.6	11.4	-	16	18.4

1988

3µm CMOS

Pitt Research

Disruption #2: Algorithmic Analog EDA

2

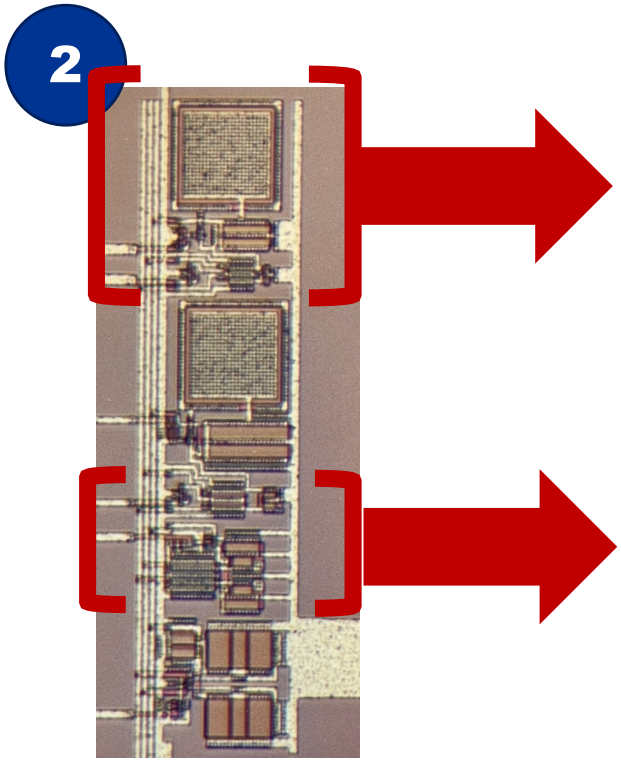

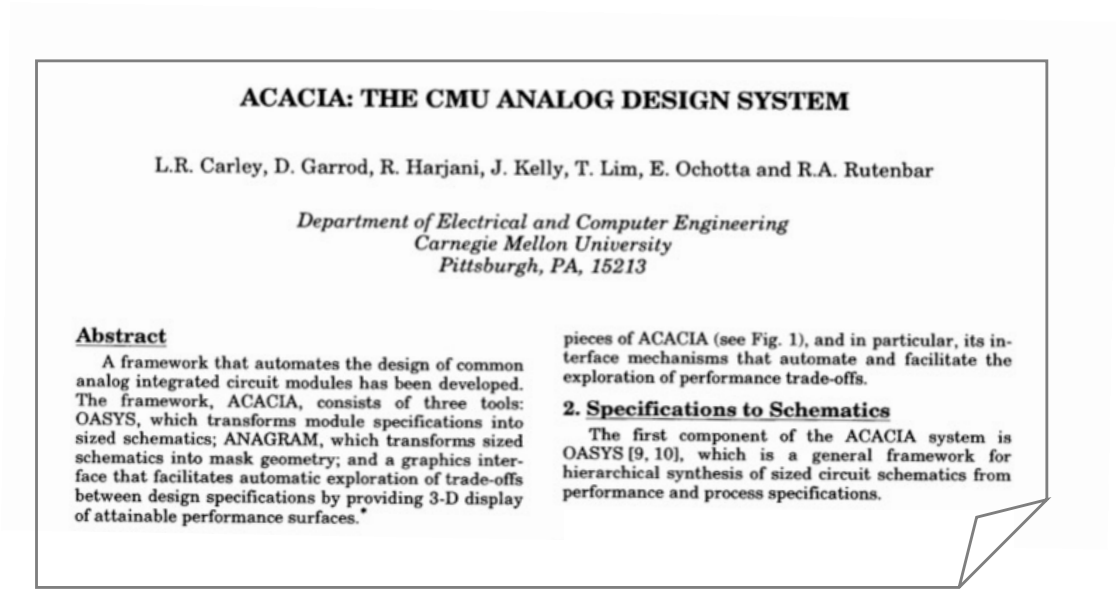

ACACIA: THE CMU ANALOG DESIGN SYSTEM

L.R. Carley, D. Garrod, R. Harjani, J. Kelly, T. Lim, E. Ochotta and R.A. Rutenbar

Department of Electrical and Computer Engineering
Carnegie Mellon University
Pittsburgh, PA, 15213

Abstract

A framework that automates the design of common analog integrated circuit modules has been developed. The framework, ACACIA, consists of three tools: OASYS, which transforms module specifications into sized schematics; ANAGRAM, which transforms sized schematics into mask geometry; and a graphics interface that facilitates automatic exploration of trade-offs between design specifications by providing 3-D display of attainable performance surfaces.*

pieces of ACACIA (see Fig. 1), and in particular, its interface mechanisms that automate and facilitate the exploration of performance trade-offs.

2. Specifications to Schematics

The first component of the ACACIA system is OASYS [9, 10], which is a general framework for hierarchical synthesis of sized circuit schematics from performance and process specifications.

First synthesized circuits from the CMU analog tools

Pitt Research

Aside: Timeline...

Carnegie Mellon University → neolinear® → cadence

1986 - 1998 **1998 - 2004** **2004 - today**

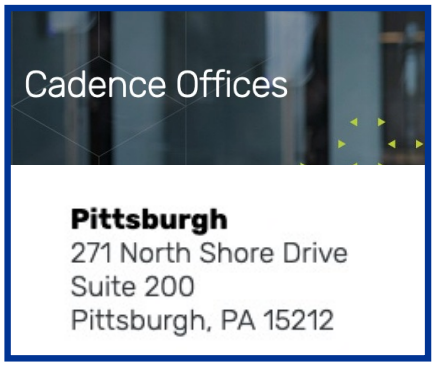

Aside: Timeline...

cadence®

neolinear®

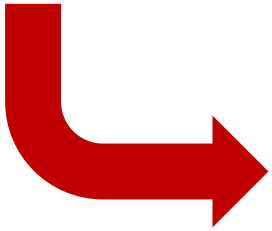

Virtuoso Analog Design Environment GXL

Variation analysis simplified

Cadence® Virtuoso® Analog Design Environment GXL provides all the capabilities of Virtuoso Analog Design Environment L and XL for thorough exploration and validation of a design. Additionally, Virtuoso Analog Design Environment GXL enables you to explore parasitic effects and sensitivities to improve yield; create design-specific worst-case corners; and find the optimum design over nominal, corner, or target yield. A flexible licensing methodology further ensures a cost-effective application of the tool's capabilities.

Neolinear tools became Virtuoso ADE-GXL

Pitt Research

Broader Disruption: *Algorithmic* EDA

Twin disruptions of 1980s, 90s

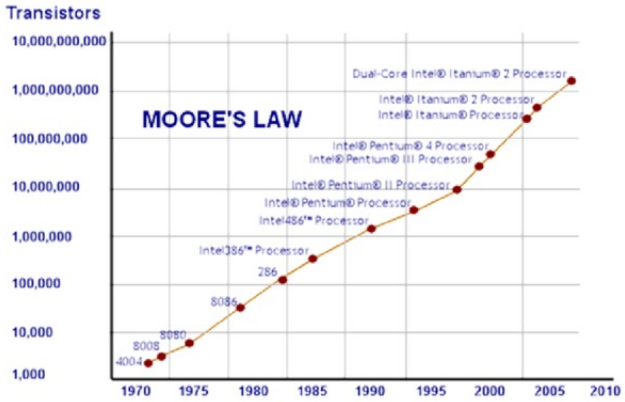

- **Moore's law**
  - Exponentially better, cheaper **hardware**
- **Rise of algorithms**
  - Model problems **formally, with principled goals**
  - Measure, approximate **circuit results, inside algorithms**
  - Focus on **accuracy, complexity & efficiency**
  - **We apply ideas from optimization**

Pitt Research

Synthesis = Algorithms + Optimization

- Smart models & engines to support **principled search** for best design to **meet designer constraints**

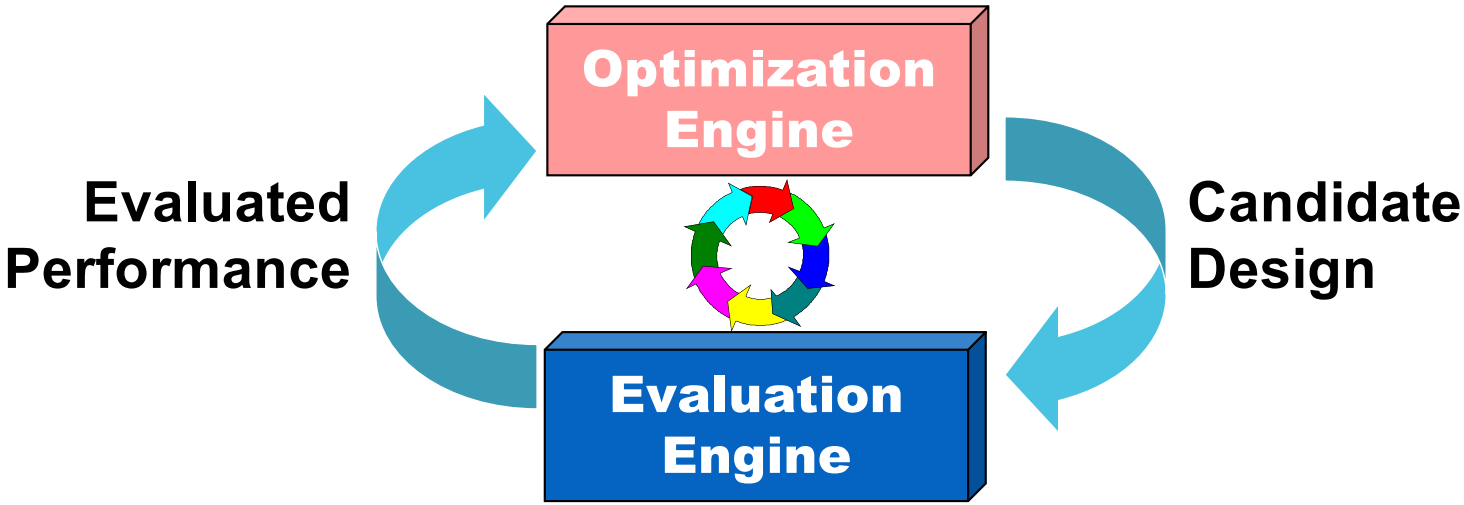

Pitt Research

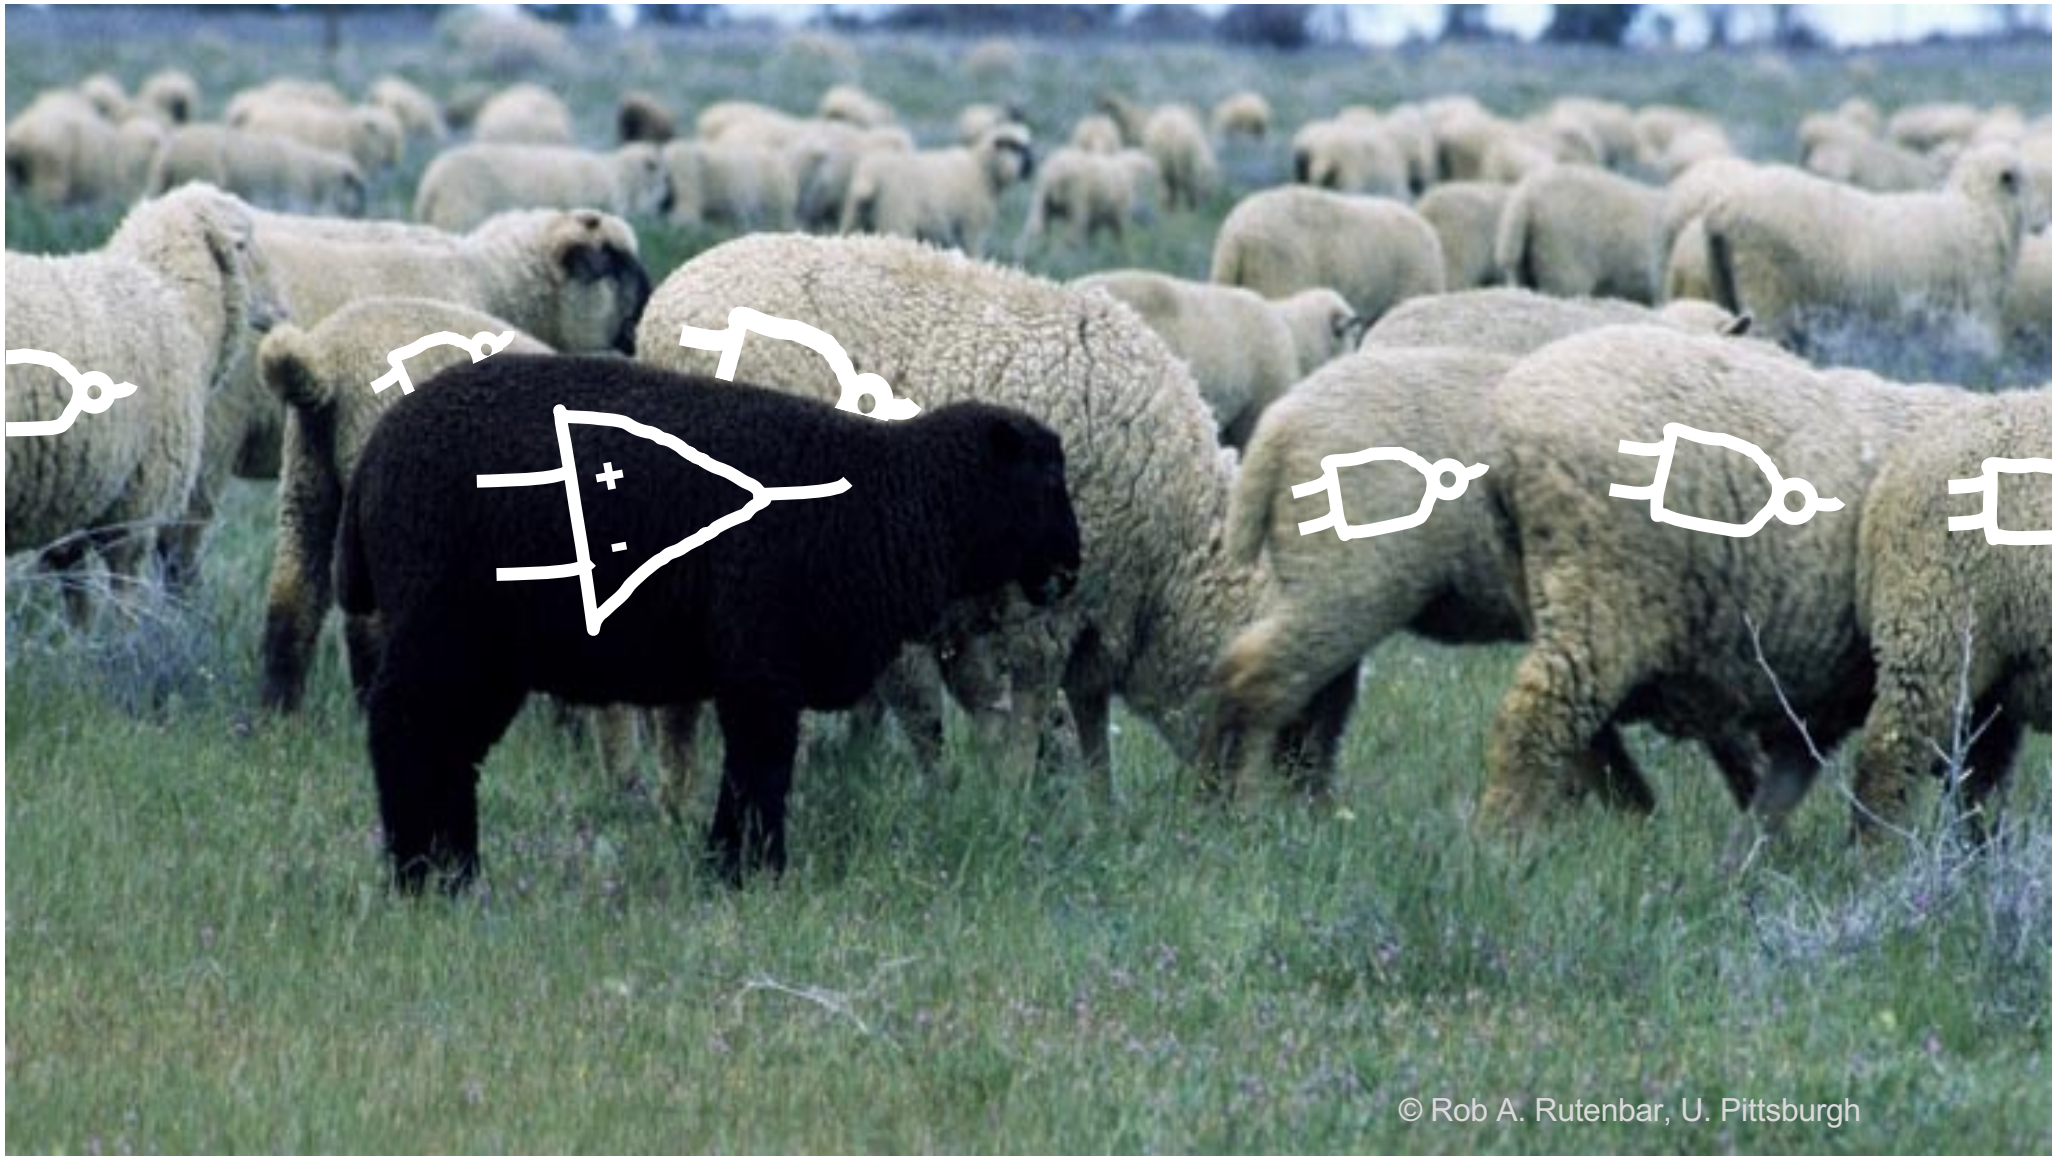
+
-
© Rob A. Rutenbar, U. Pittsburgh

Analog EDA: Fundamental Debate

- **Analog –** it's art, mysterious, all **patterns in my head**
- **Nope** – it's **optimization**: big hard optimization

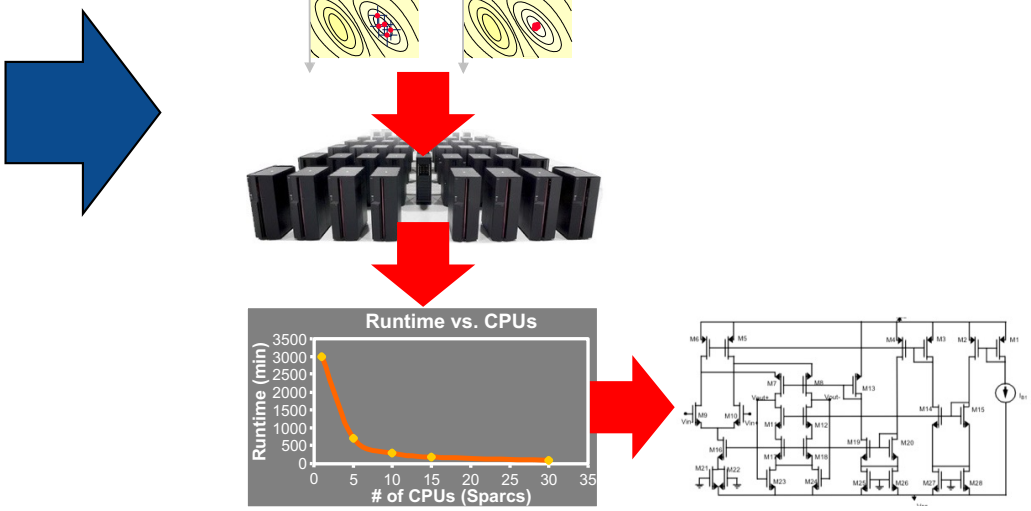

Pitt Research

Landscape of "Core" Analog Tools

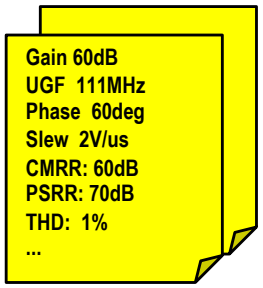

Constraint Capture

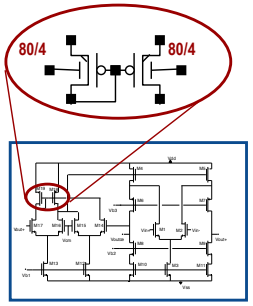

Sizing & Biasing

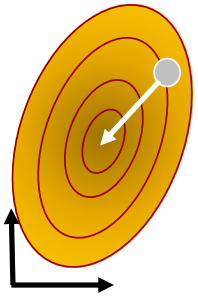

Yield Optimization

CIRCUIT

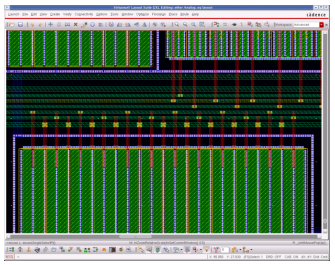

Device Gen

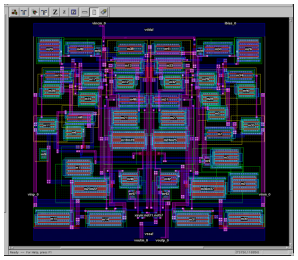

Placement

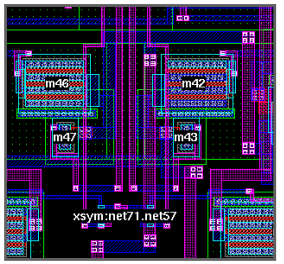

Routing

LAYOUT

Pitt Research

Two (Historical) Analog Examples

- **Practical Spice-in-the-Loop optimizers for sizing/biasing**

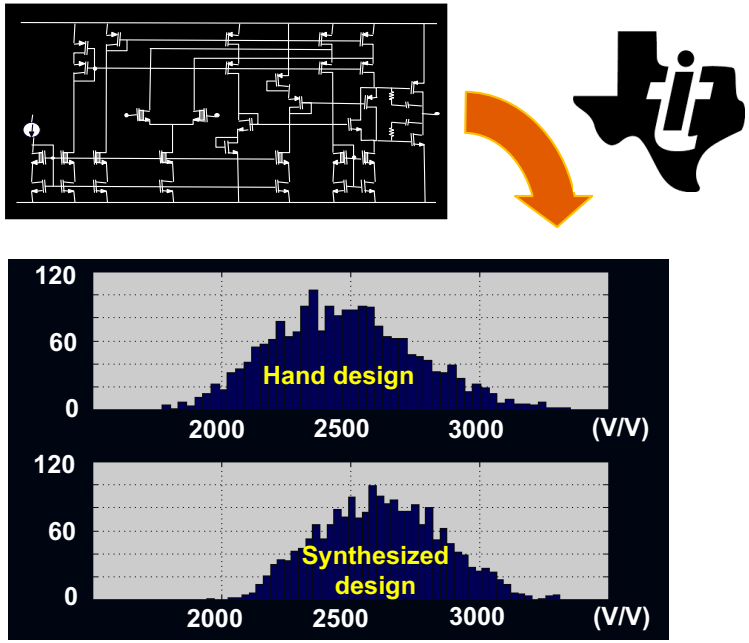

- **Automatic layouts with both aesthetics & performance**

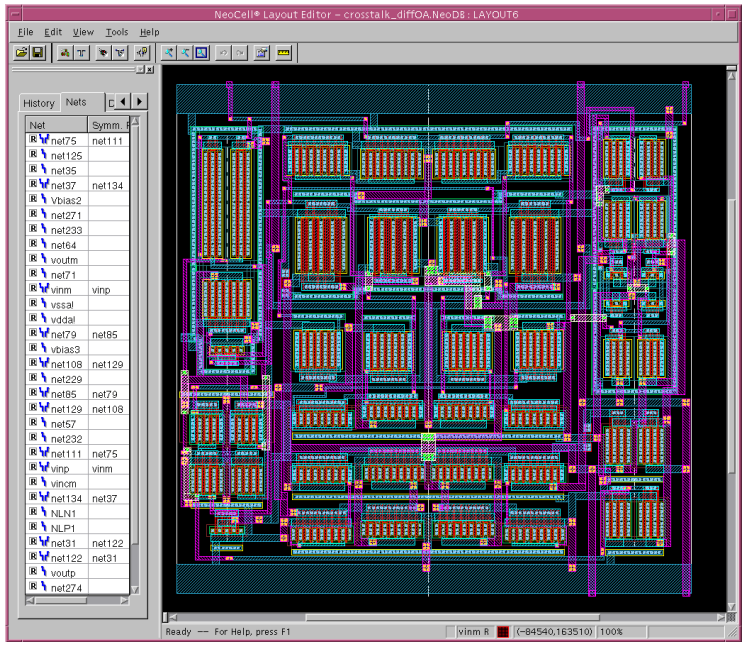

Tools from CMU

Tools from Neolinear

Pitt Research

Looking Back....

What did we get **right?**

What did we get **wrong?**

Simplify All EDA into 2-Dim Space

- **Claim***: There are only 2 kinds of EDA tools

INSURANCE

PAIN RELIEF

Pitt Research

Analog EDA on Pain/Fear Plot

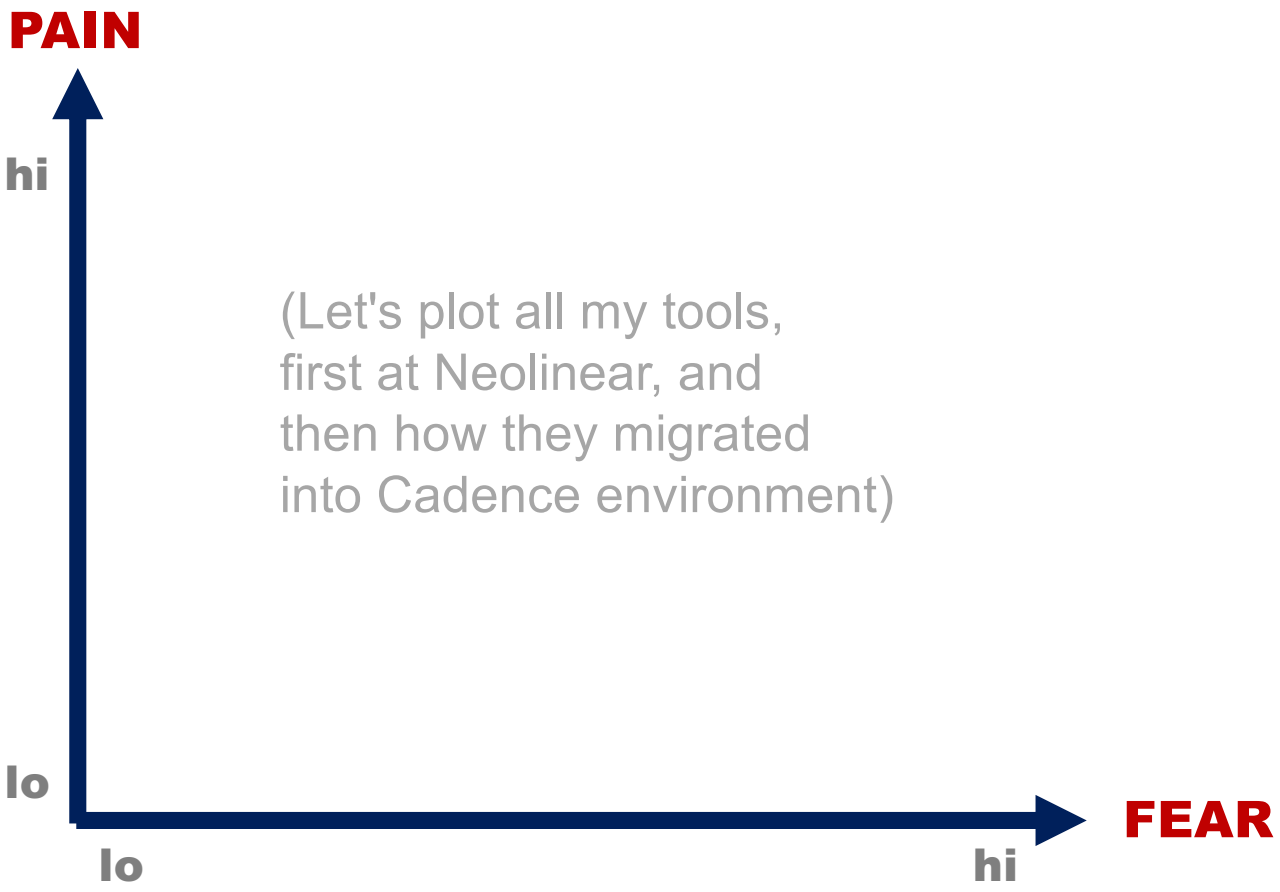

(Let's plot all my tools,
first at Neolinear, and
then how they migrated
into Cadence environment)

Pitt Research

Analog EDA on Pain/Fear Plot

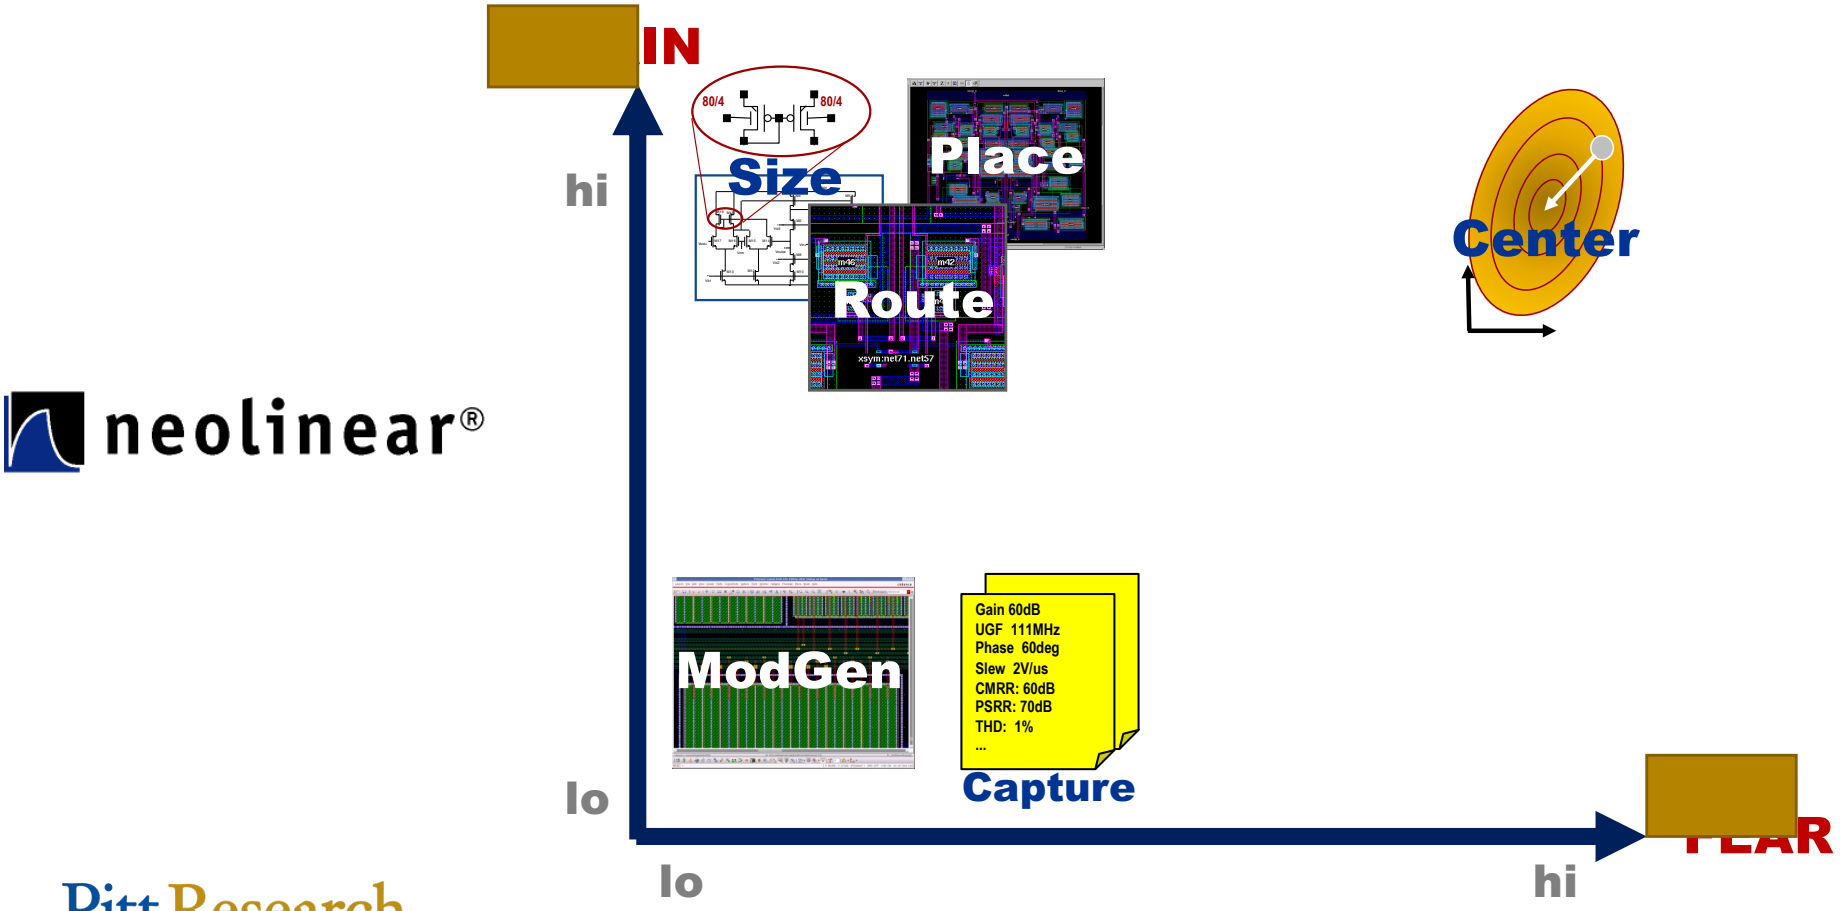

neolinear®

Pitt Research

Analog EDA on Pain/Fear Plot

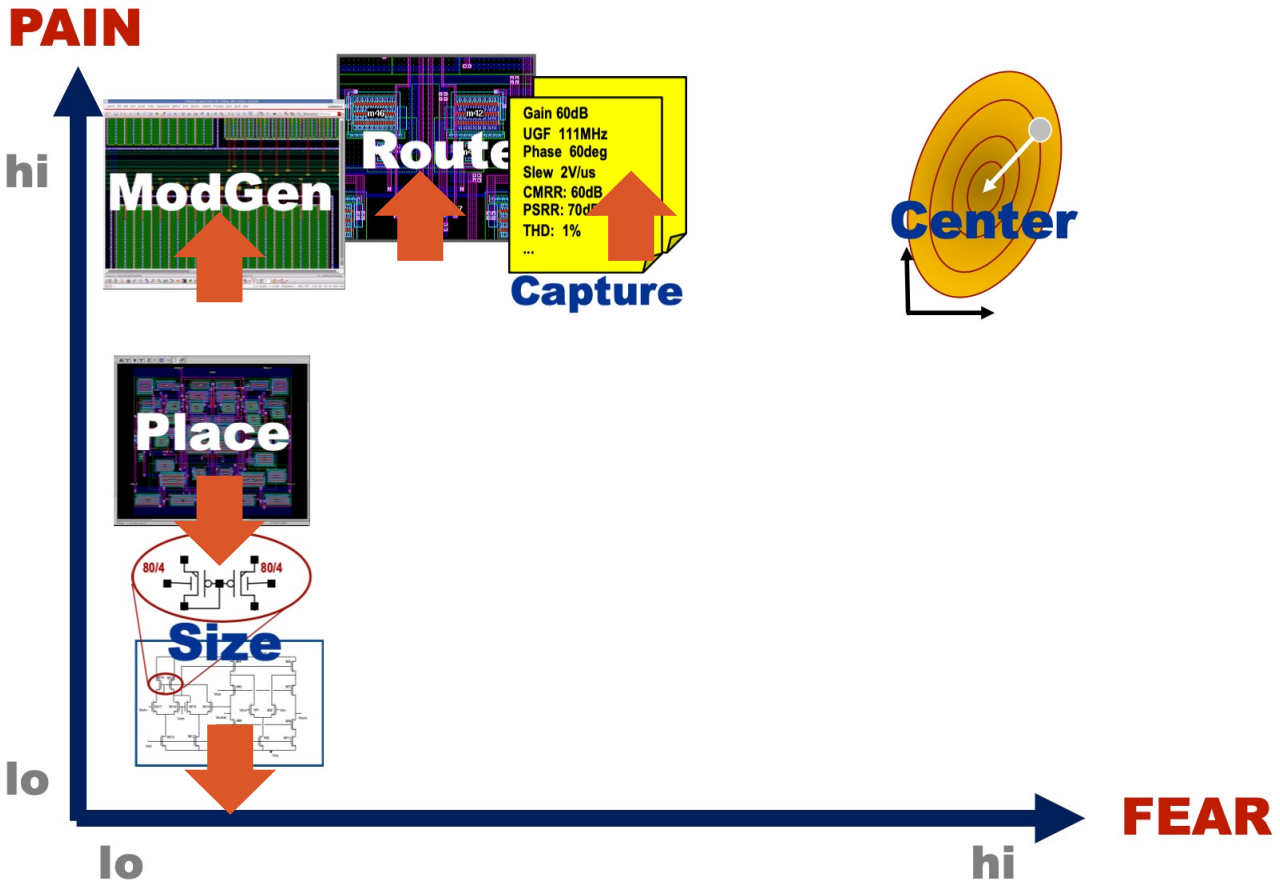

cadence™

(After connection with more customers at large scale…)

Pitt Research

Aside: 20+ Yrs of Moore's Law…

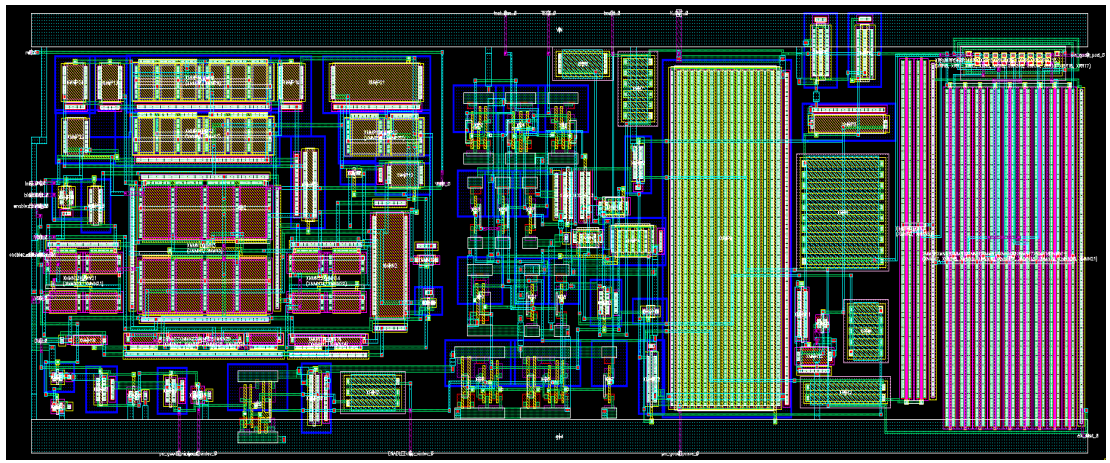

250nm
Place ~ floorplan

In either case, *nobody* wants to draw device shapes all by hand

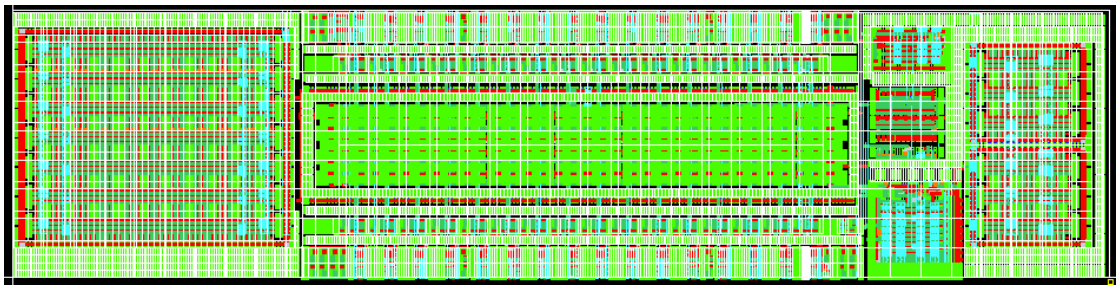

14nm
Place ~ fixed grid

Pitt Research

Teaser for Final Part of Talk

- Analog – **expert**

- Analog – **optimizer**

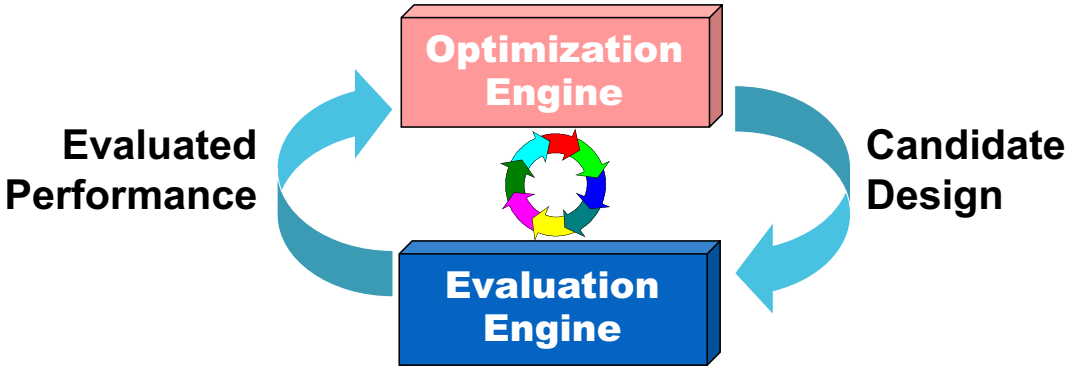

Pitt Research

Disruption #3

3

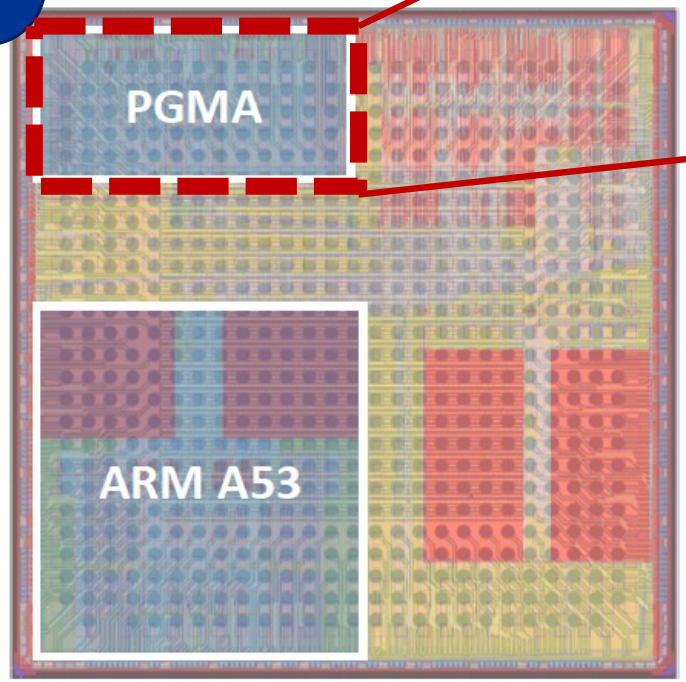

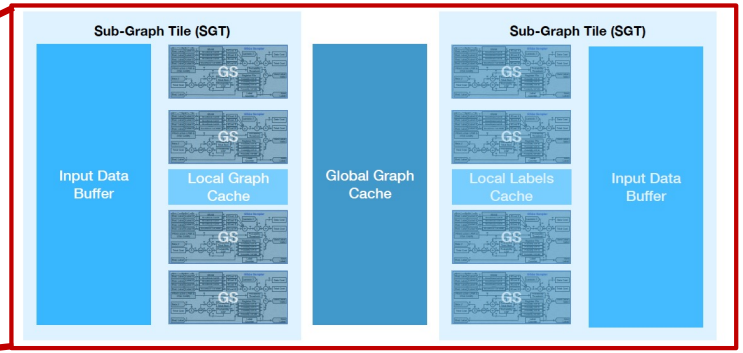

2020

A 3mm² Programmable Bayesian Inference Accelerator for Unsupervised Machine Perception using Parallel Gibbs Sampling in 16nm

Glenn G. Ko[1], Yuji Chai[1], Marco Donato[1], Paul N. Whatmough[1,2], Thierry Tambe[1], Rob A. Rutenbar[3], David Brooks[1], Gu-Yeon Wei[1]
[1]Harvard University, MA, [2]Arm Research, MA, [3]University of Pittsburgh, PA

- TSMC 16nm FFC
- ~2M gates, 450MHz
- Part of a larger SOC experiment at Harvard called **SM5: ML for IOT**

Pitt Research

PGMA vs Arm53 vs eFPGA (on SOC)

- **PGMA: 1000X+ throughput vs CPU; 6X+ ops/W vs eFPGA**

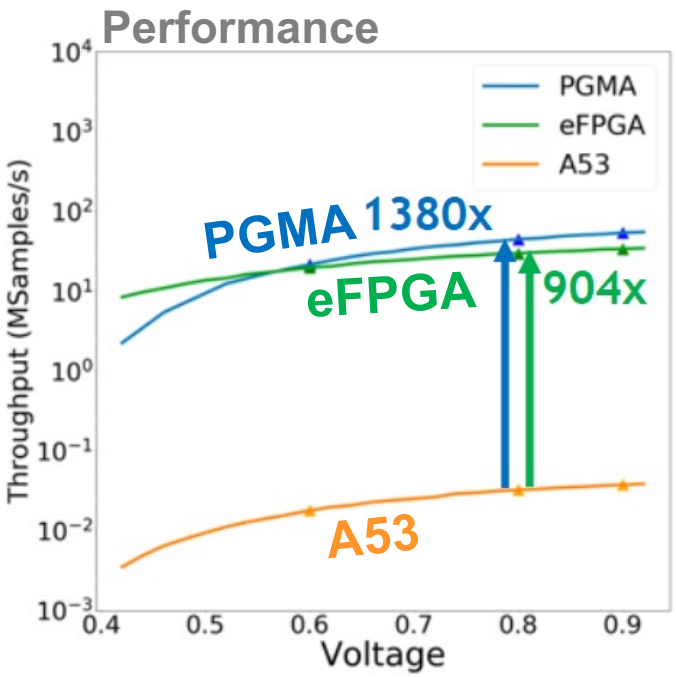

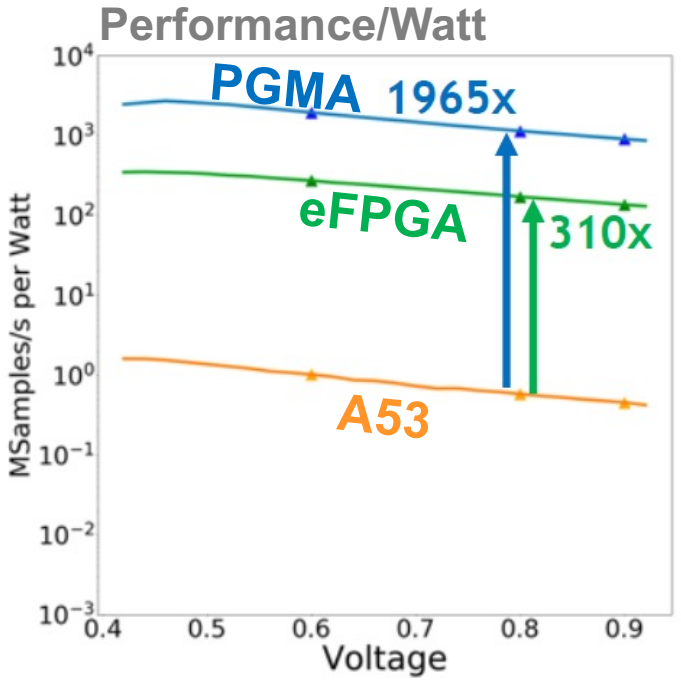

Pitt Research

Key Takeaways

- The world **changed** fundamentally in 2012
- **AI** is invading everything, everywhere (even **Rob**…)
- Even **EDA...**

ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky
University of Toronto
kriz@cs.utoronto.ca

Ilya Sutskever
University of Toronto
ilya@cs.utoronto.ca

Geoffrey E. Hinton
University of Toronto
hinton@cs.utoronto.ca

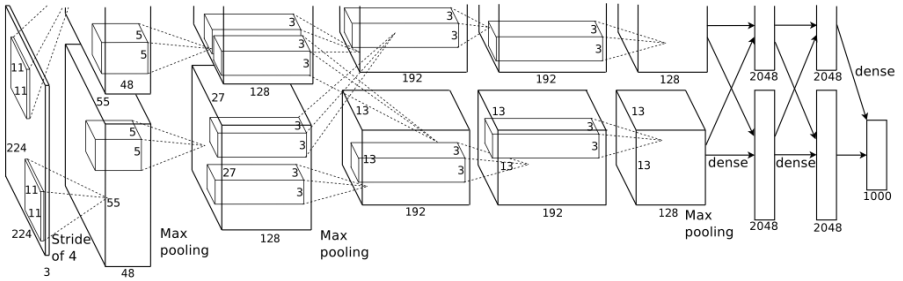

NIPS 2012

Pitt Research

Quantifying AI in EDA: ICCAD22

ICCAD22 program: Count every talk, tutorial, workshop, slot, that **mentions AI, ML, deep nets**…

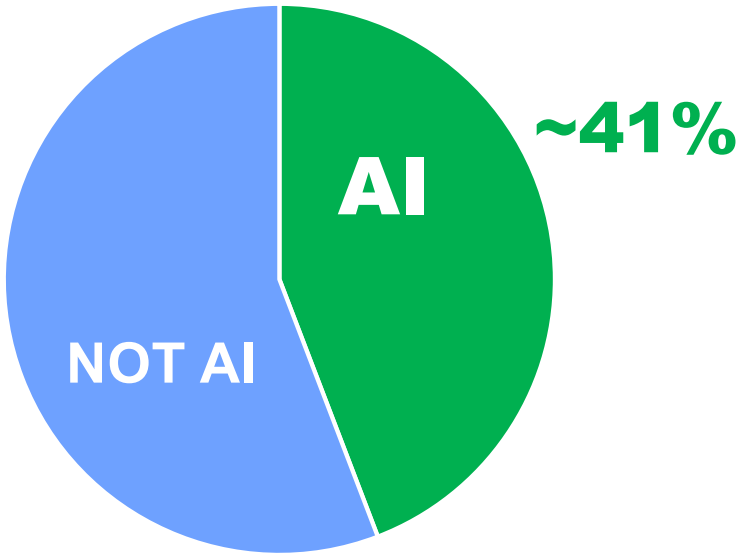

172 ICCAD event slots

Pitt Research

BTW – Same for DAC22

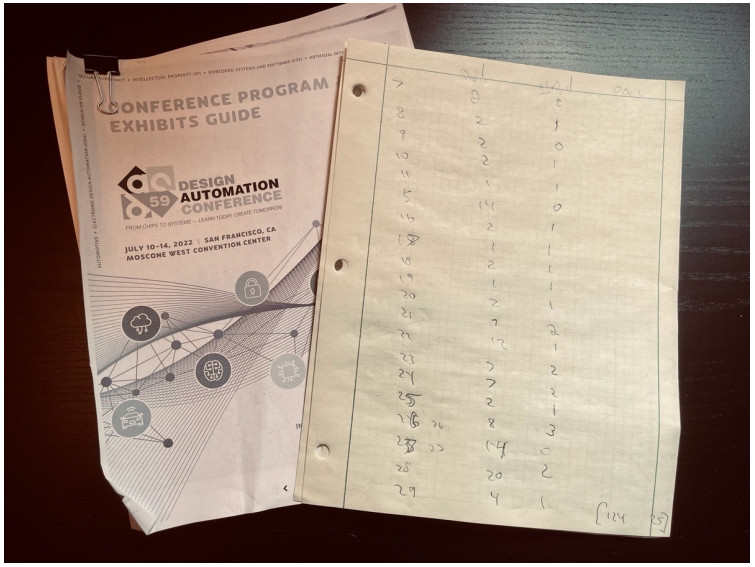

DAC22 program: Count every talk, tutorial, poster, slot, that **mentions AI, ML, deep nets**…

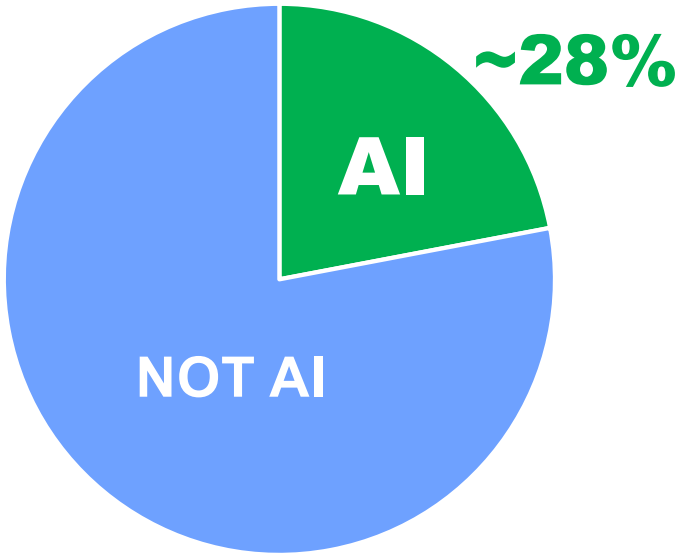

836 DAC event slots

Pitt Research

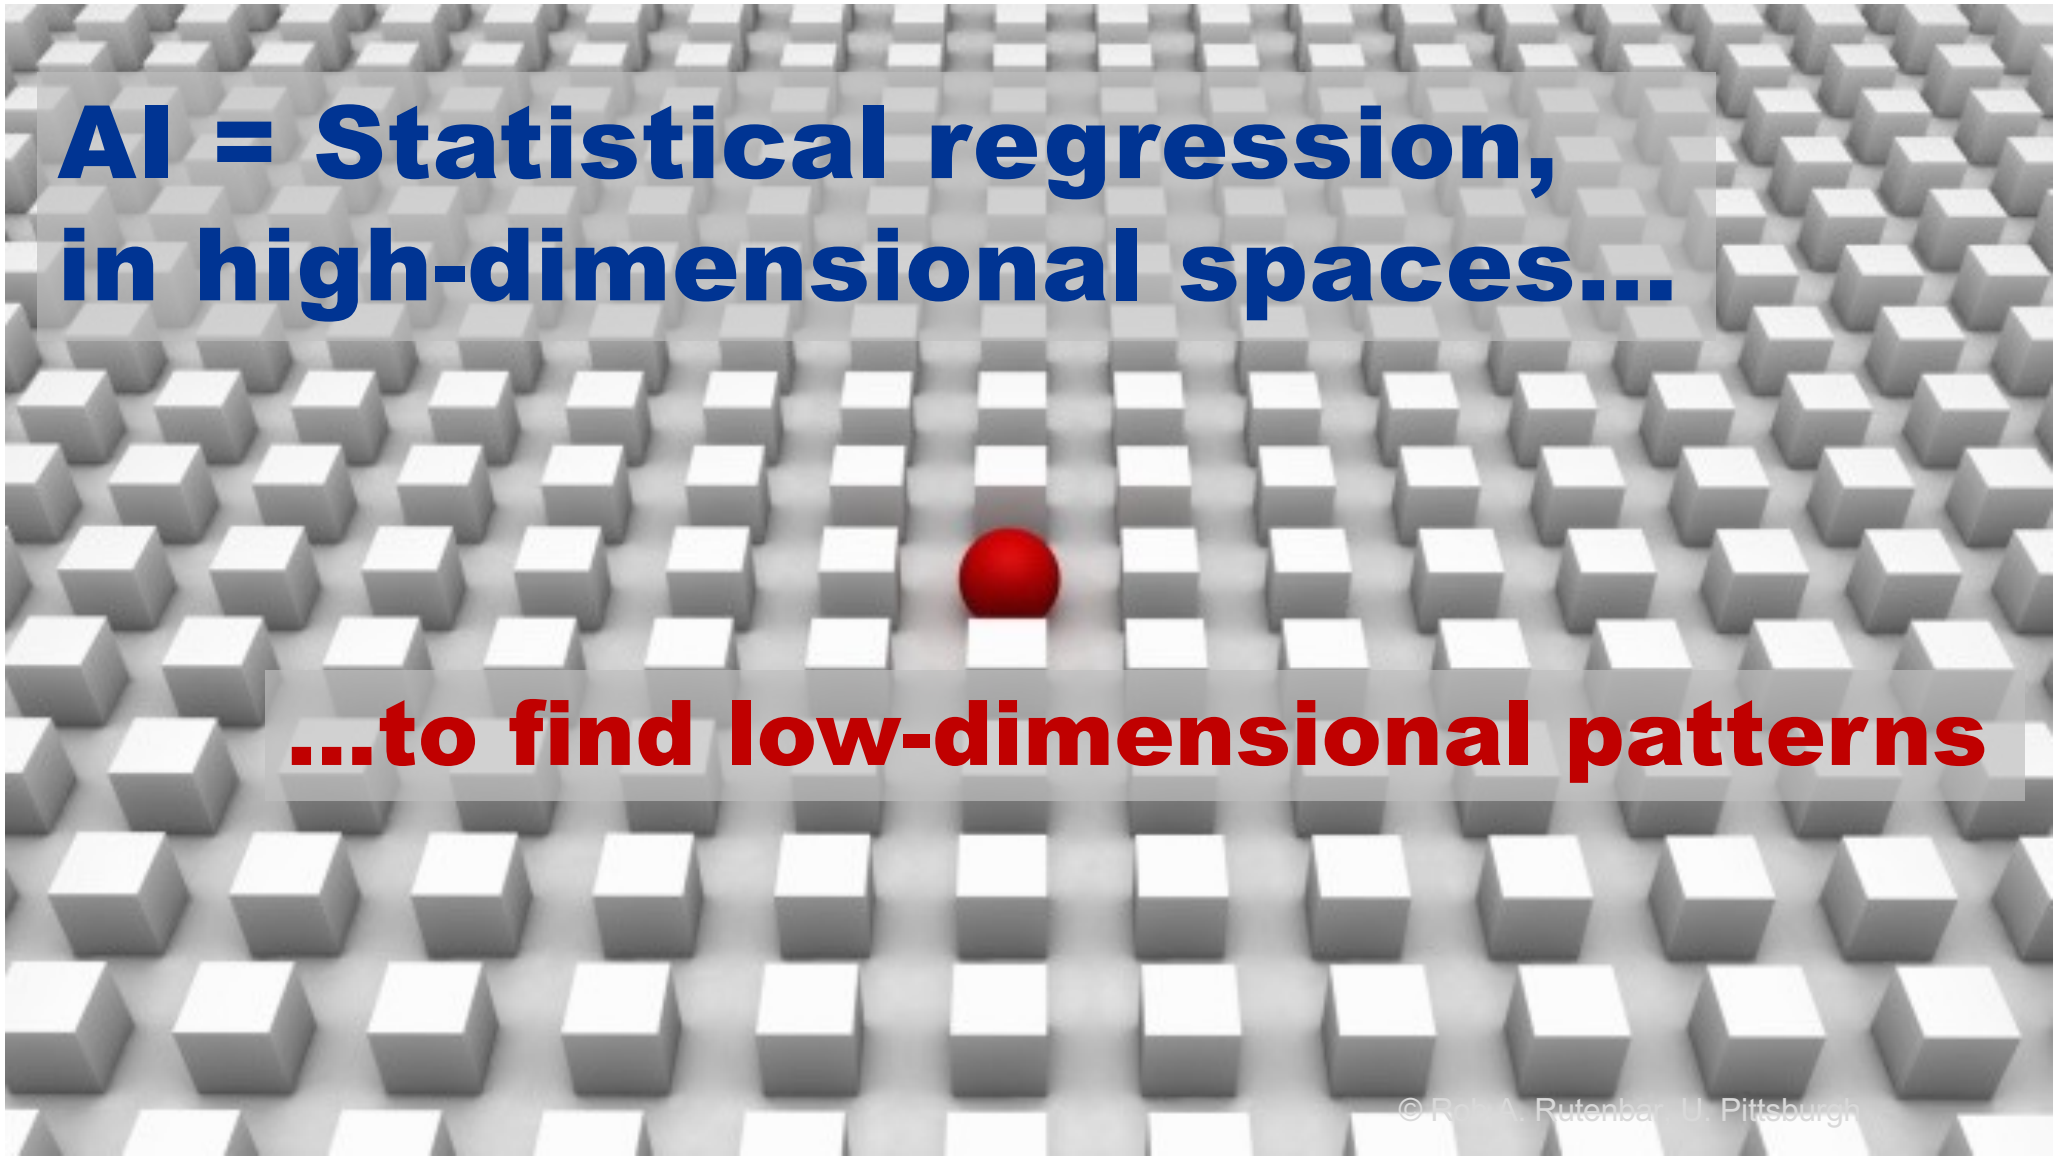

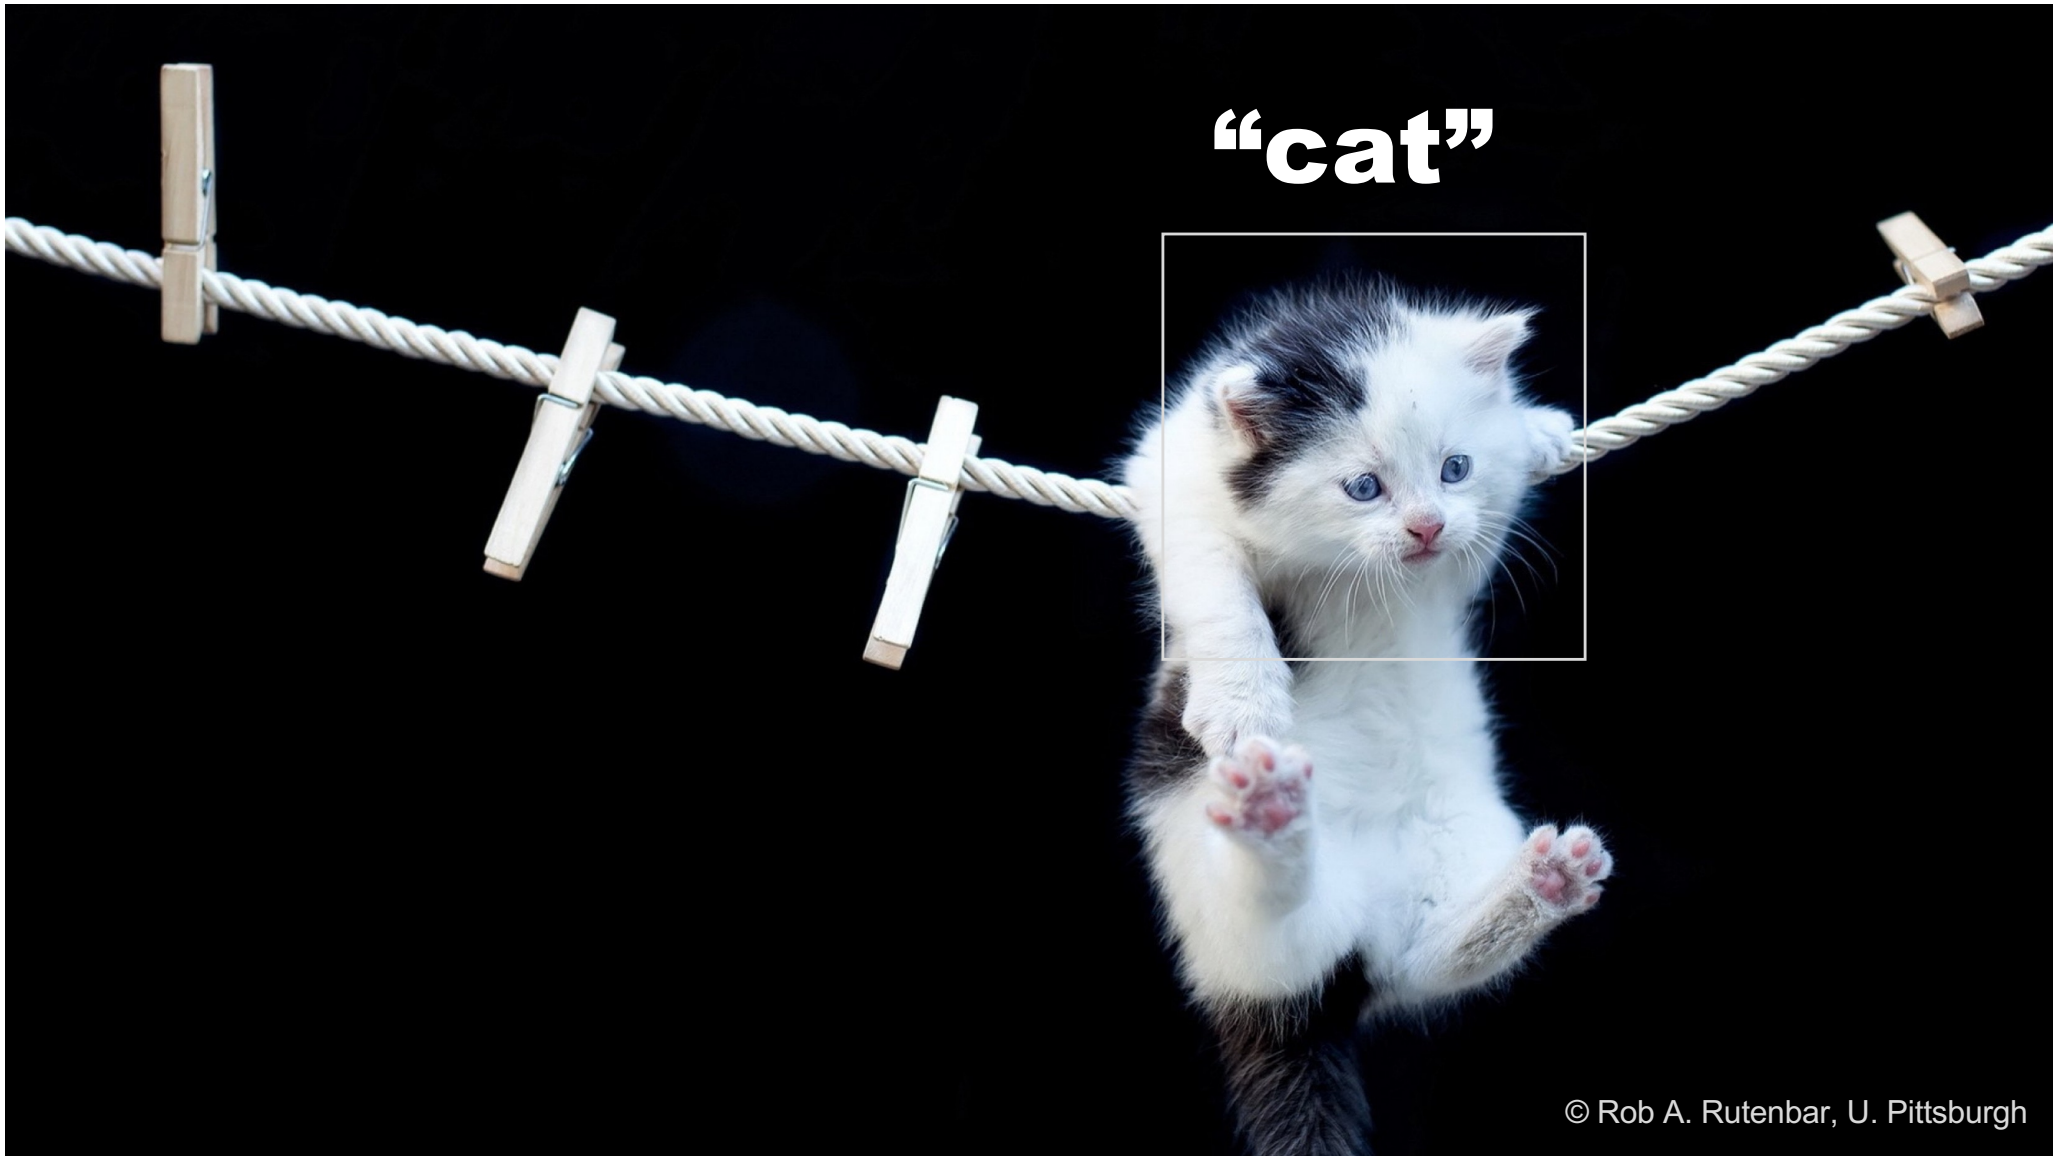
“cat”
© Rob A. Rutenbar, U. Pittsburgh

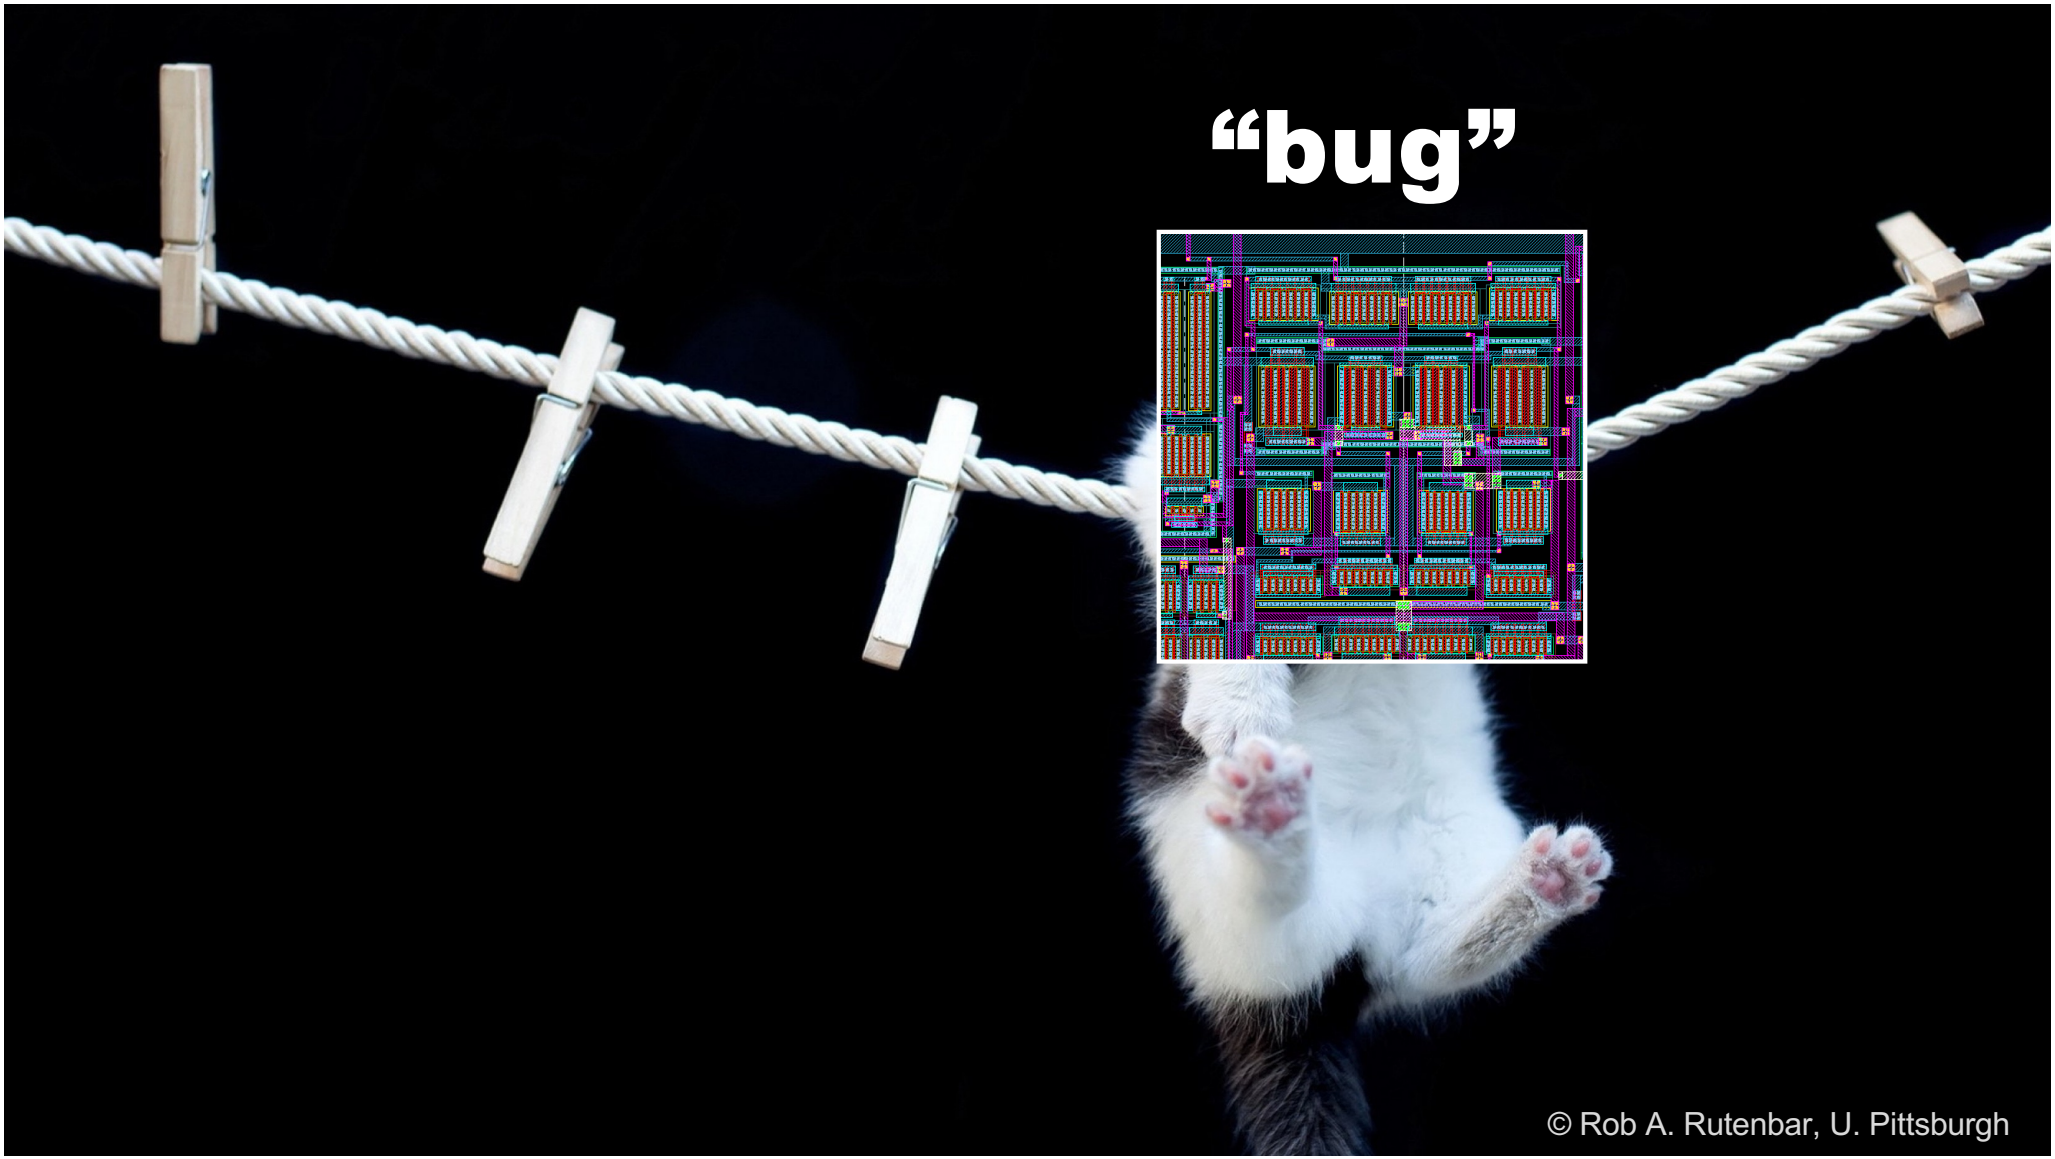
“bug”
© Rob A. Rutenbar, U. Pittsburgh

Analog and AI, Revisited

- Analog – **expert**

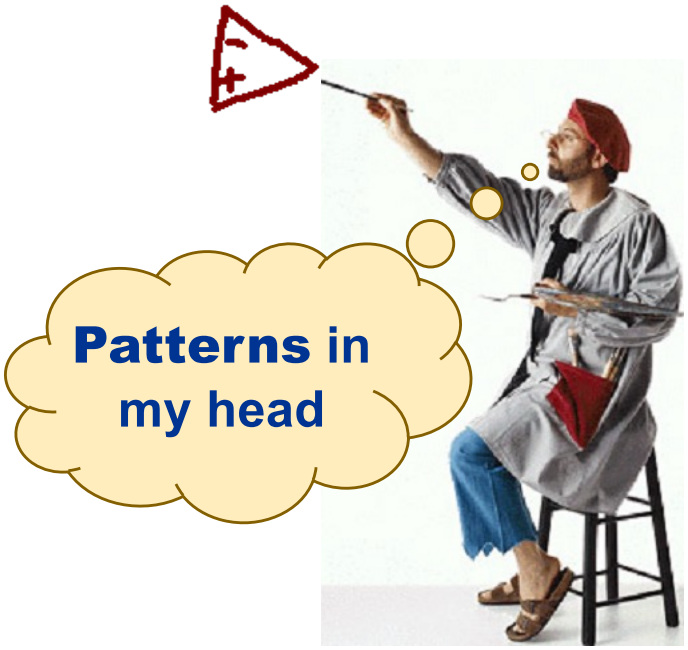

- Analog – **optimizer & deep patterns**

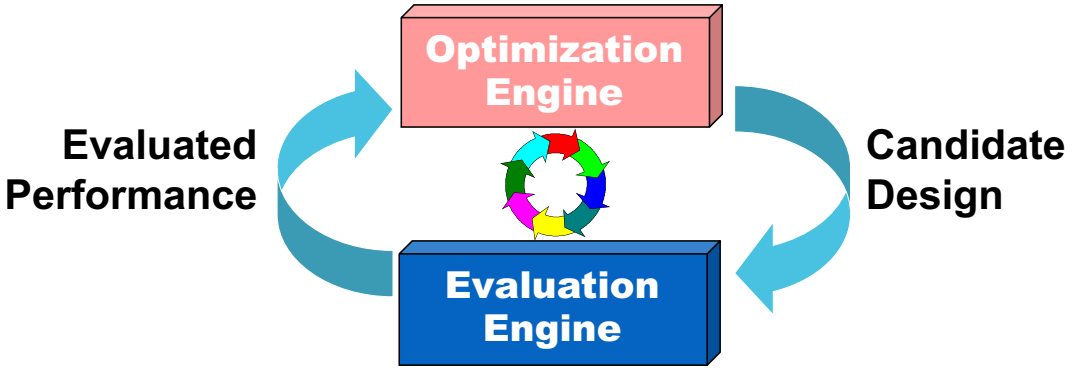

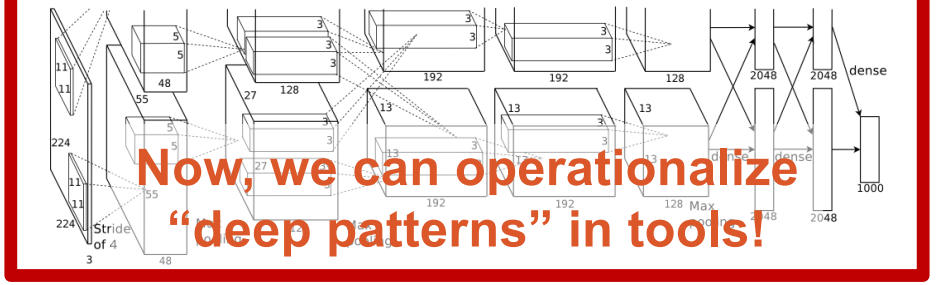

Pitt Research

Analog as “Operationalized Patterns”

“This is ugly”

“This is not what I had in mind”

Fundamental truth about analog:
Aesthetics often a surrogate for correctness

Pitt Research

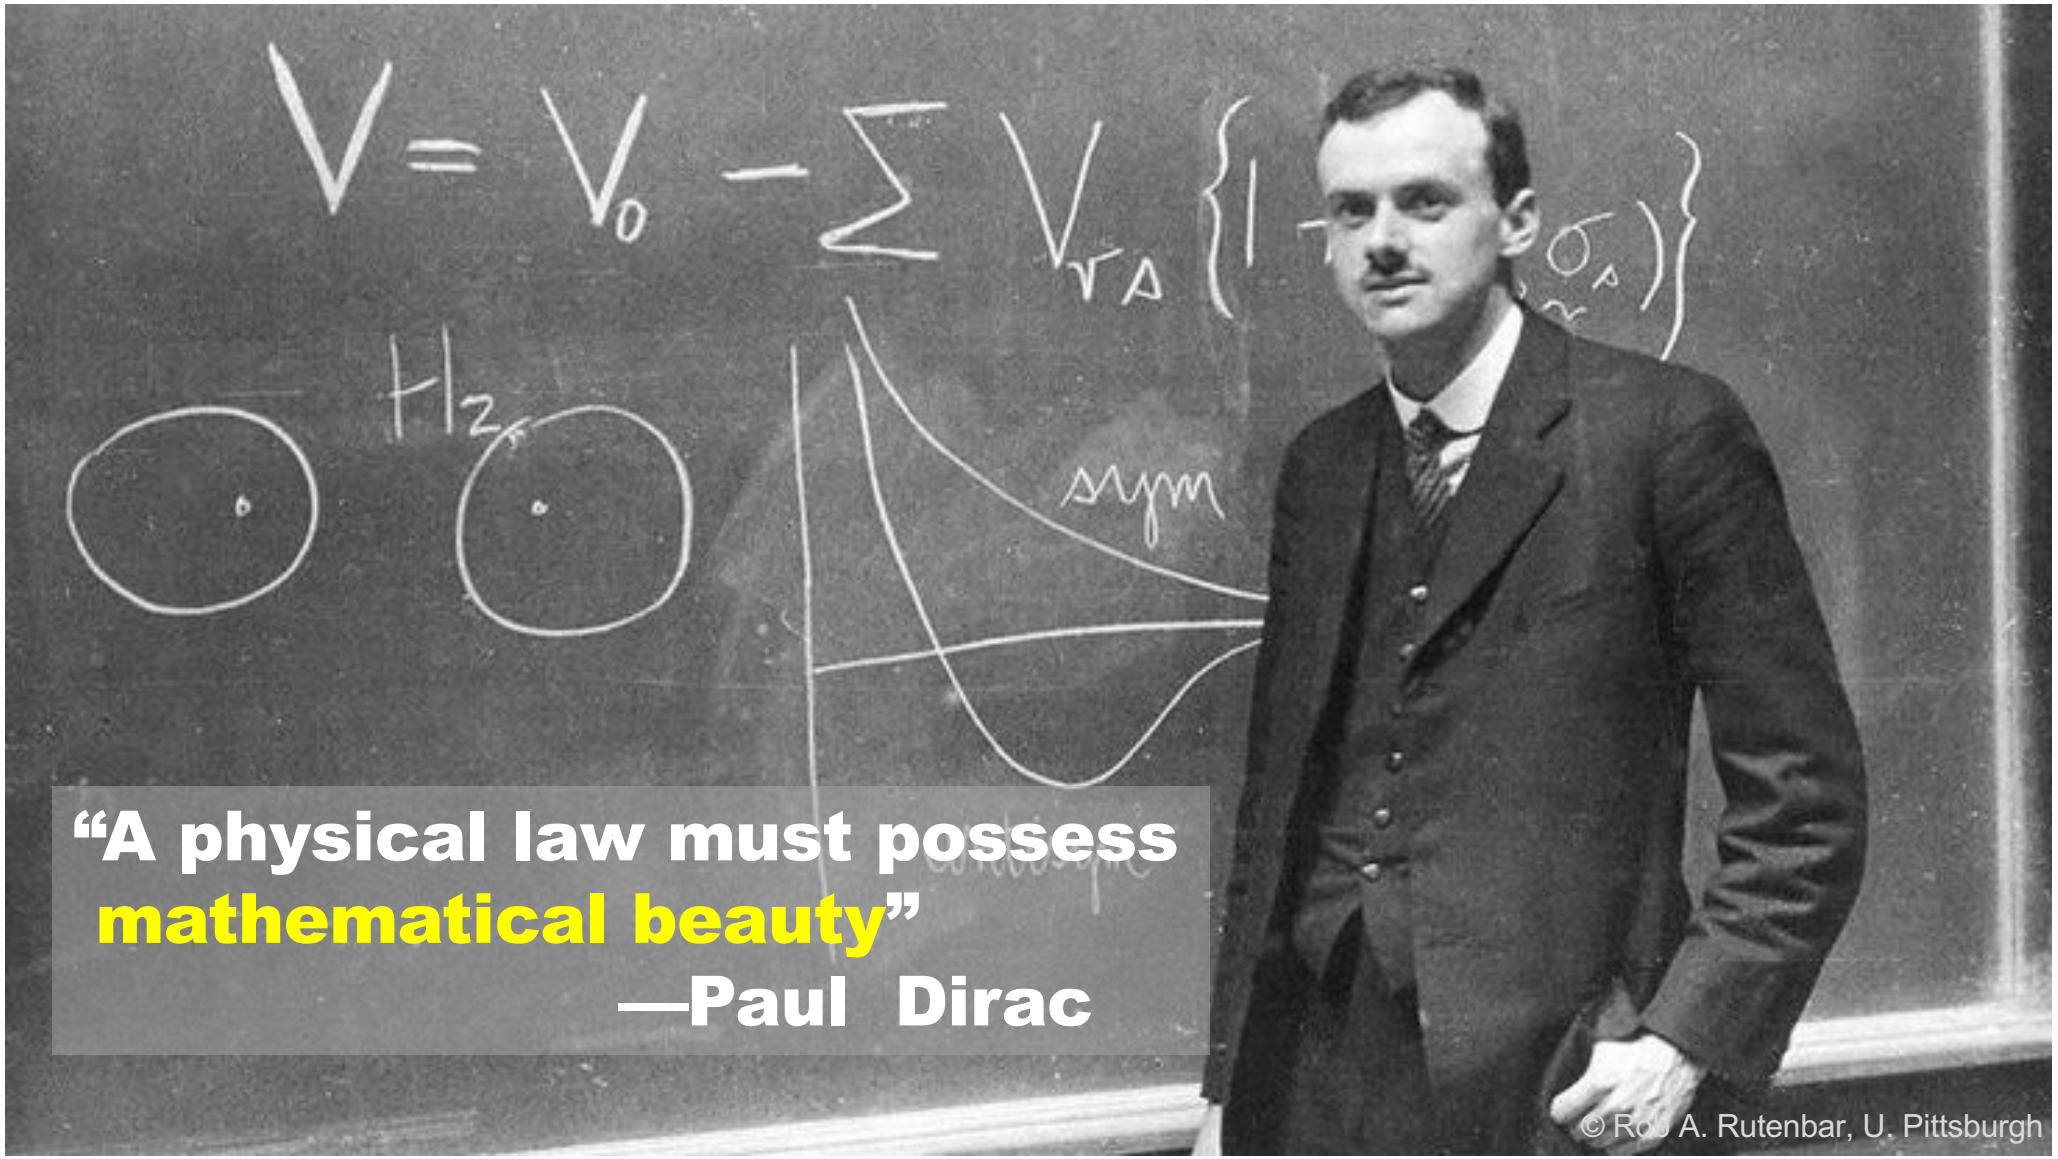

"A physical law must possess mathematical beauty"

—Paul Dirac

AI Meets "Core" Analog Tools...?

AI

Gain 60dB
UGF 111MHz
Phase 60deg
Slew 2V/us
CMRR: 60dB
PSRR: 70dB
THD: 1%
...

Capture

AI

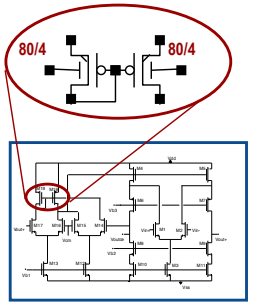

Sizing

AI

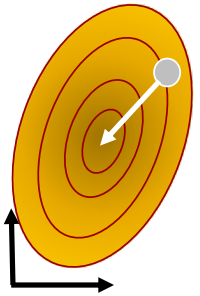

Center

AI

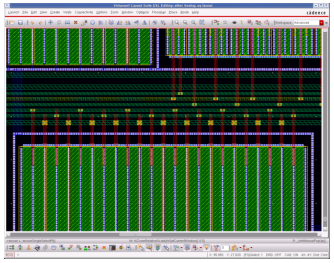

ModGen

AI

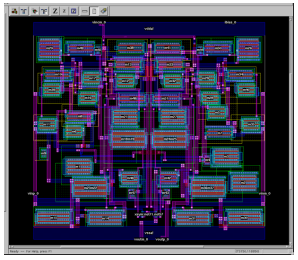

Place

AI

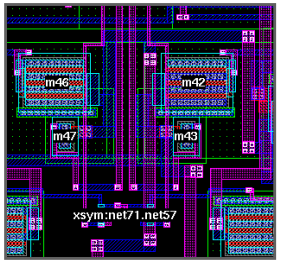

Route

Pitt Research

Digital Ex: AI Meets Floorplanning

Article

A graph placement methodology for fast chip design

https://doi.org/10.1038/s41586-021-03544-w
Received: 3 November 2020
Accepted: 13 April 2021
Published online: 9 June 2021

Azalia Mirhoseini[1,4✉], Anna Goldie[1,3,4✉], Mustafa Yazgan[2], Joe Wenjie Jiang[1], Ebrahim Songhori[1], Shen Wang[1], Young-Joon Lee[2], Eric Johnson[1], Omkar Pathak[2], Azade Nazi[1], Jiwoo Pak[2], Andy Tong[2], Kavya Srinivasa[2], William Hang[2], Emre Tuncer[2], Quoc V. Le[1], James Laudon[1], Richard Ho[2], Roger Carpenter[2] & Jeff Dean[1]

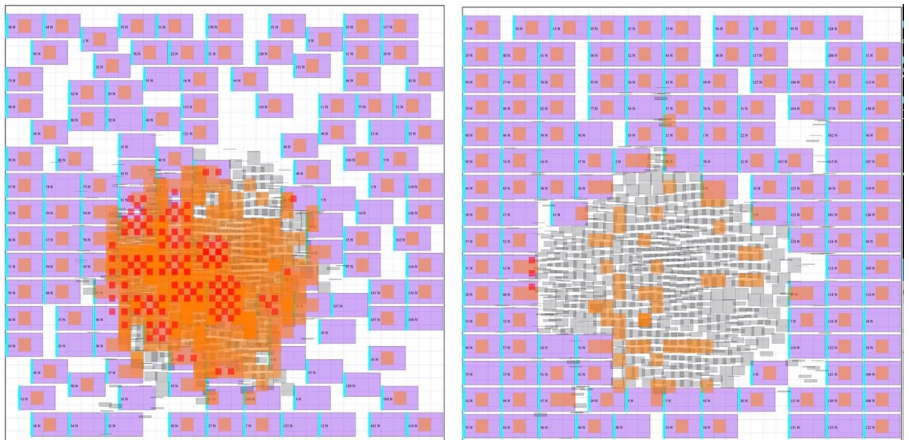

NEWS AND VIEWS | 09 June 2021

AI system outperforms humans in designing floorplans for microchips

A machine-learning system has been trained to place memory blocks in microchip designs. The system beats human experts at the task, and offers the promise of better, more-rapidly produced chip designs than are currently possible.

EDITORIAL | 09 June 2021

Google's AI approach to microchips is welcome – but needs care

Artificial intelligence can help the electronics industry to speed up chip design. But the gains must be shared equitably.

Pitt Research

Bonus Disruption…

1

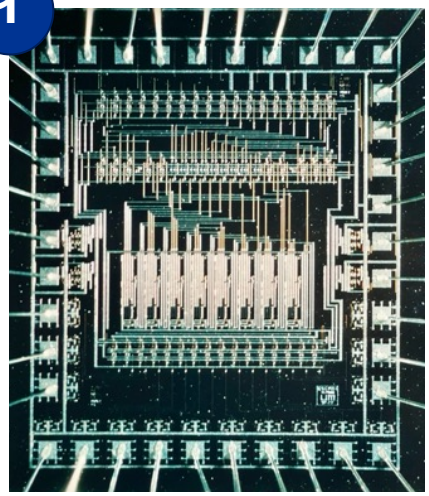

2

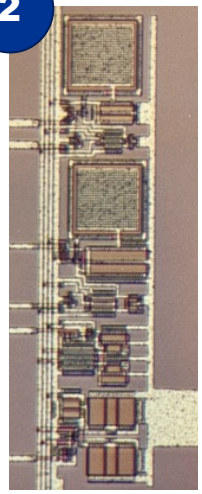

3

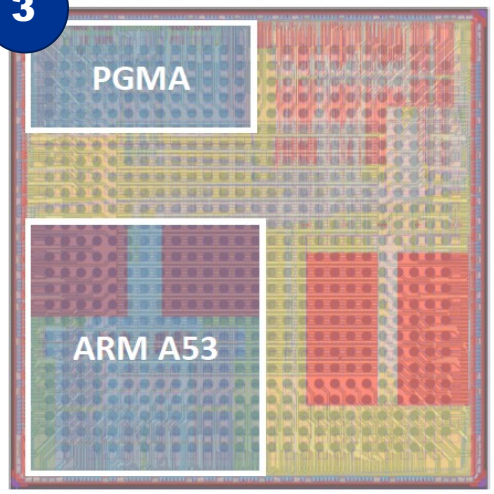

4

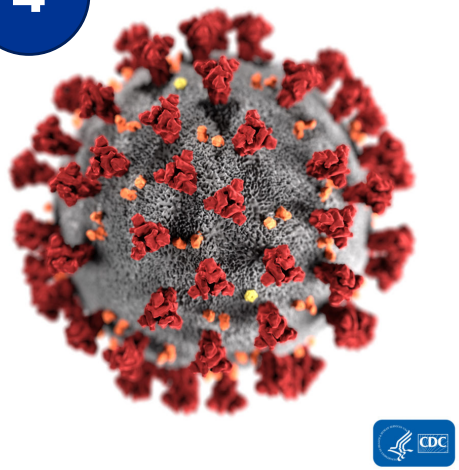

Pitt Research

VLSI CAD MOOCs

- Since 2013, I'm running MOOCs on **VLSI CAD**

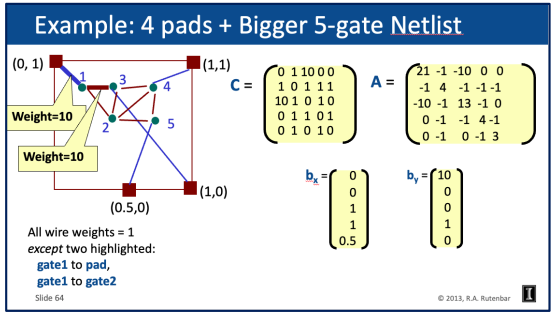

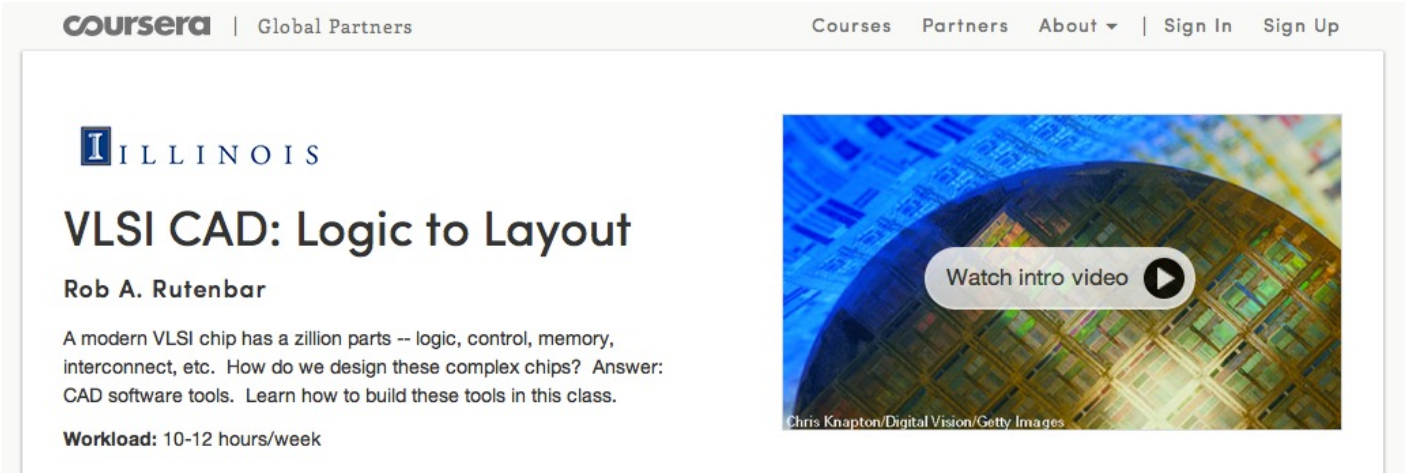

Pitt Research

Enrollment: Coursera CAD MOOCS

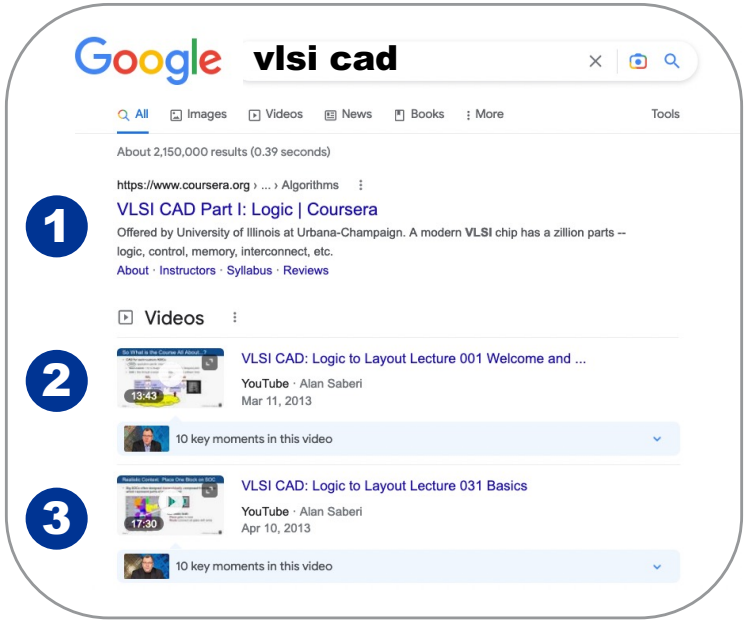

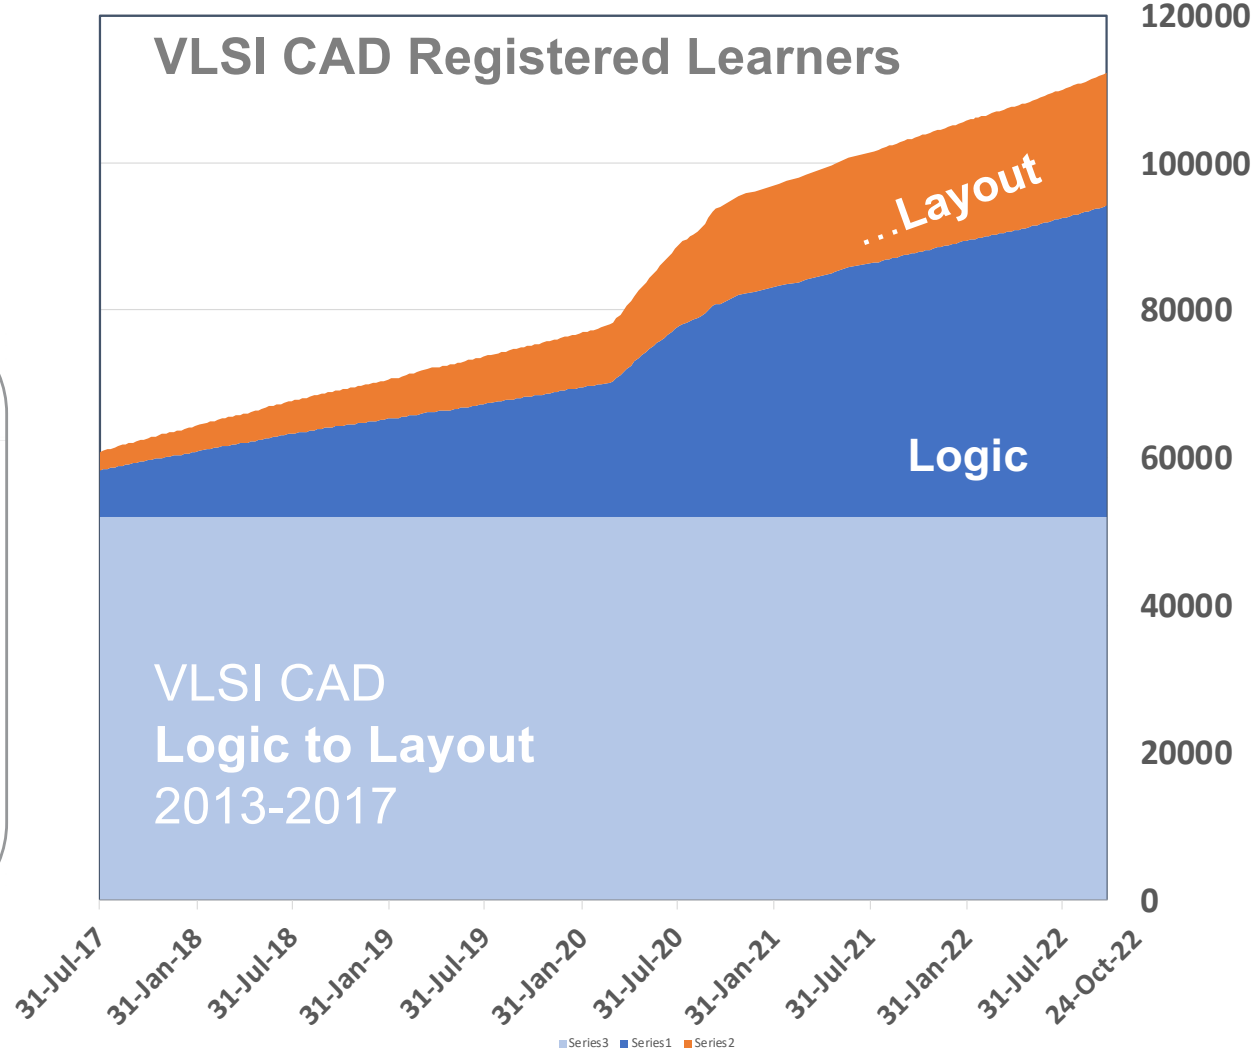

Pitt Research

Enrollment: Coursera CAD MOOCS

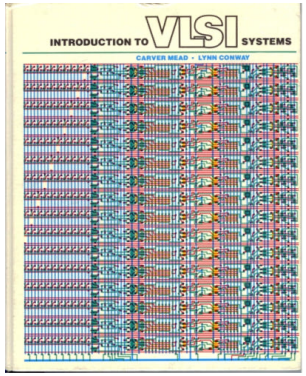

Hypothesis: COVID is responsible for more global Interest in EDA than anything we've seen – *since this book*

Pitt Research

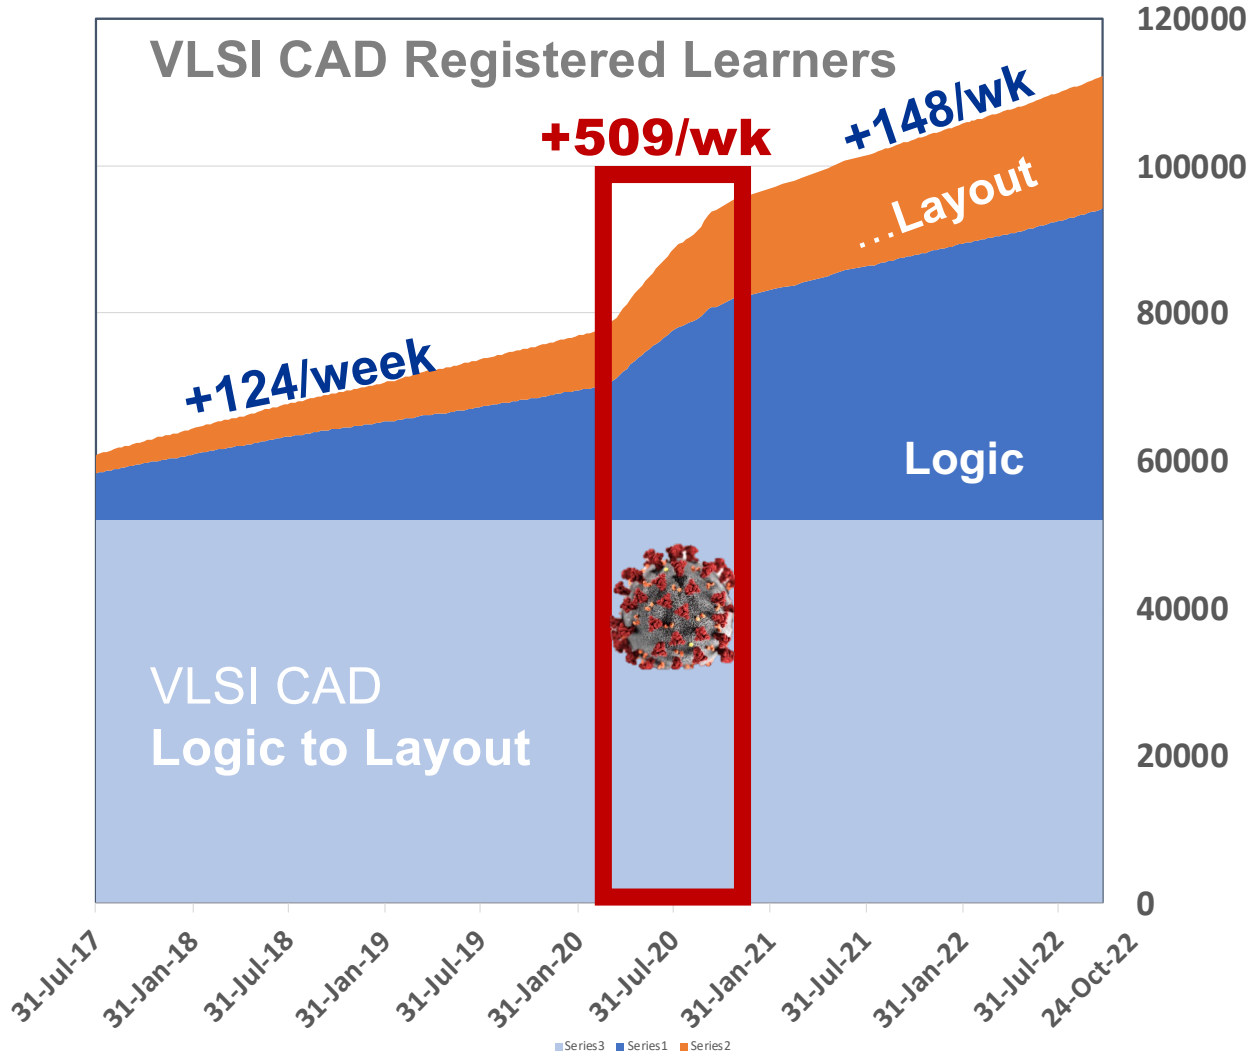

ICCAD & SIGDA:

Thank You!

Pitt Research