Walking Pads: Fast Power-Supply Pad-Placement Optimization

K. Wang*, **B.H. Meyer†**, R. Zhang*, K. Skadron*, M. Stan*
*University of Virginia, †McGill University

Research made possible with the support of the NSF, SRC, MARCO, and DARPA

Background: Power Delivery

- *Objective*: provide a spatially uniform and temporally stable voltage supply
- *Challenge*: non-uniform current demand
- On-chip power delivery interface: *controlled collapse chip connection* (C4) pads
  - – Shared by power and I/O subsystems

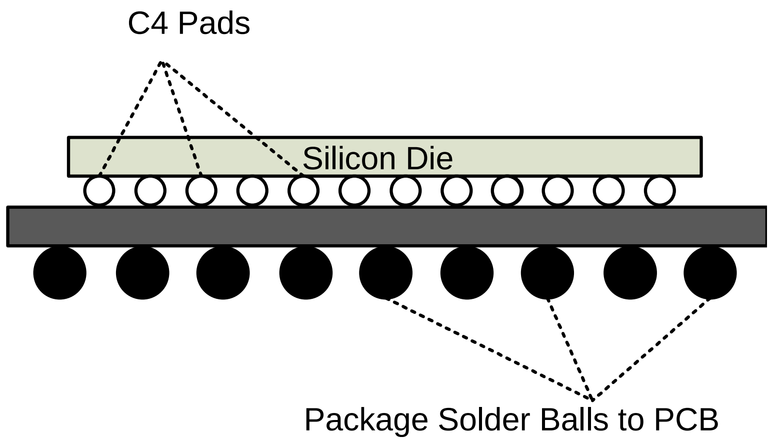

Motivation

- Process scaling complicates power delivery
  1. Imperfect supply voltage scaling ⇒ Increasing current density
  2. C4 pads are a scarce resource
     - C4 density has not scaled
     - Increasing I/O demands
- Power pads must be placed close to current sources

⇒ Pad allocation must be optimized!

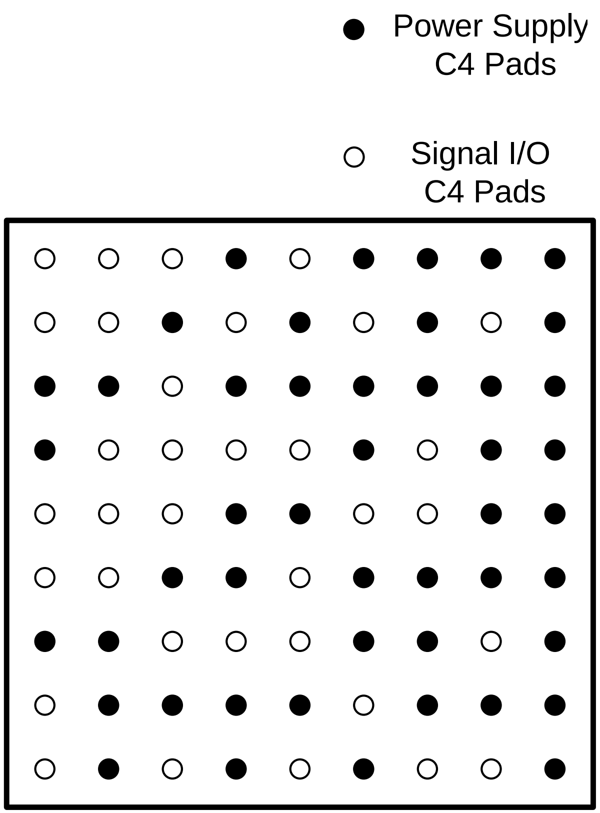

C4 Pad Allocation Matters

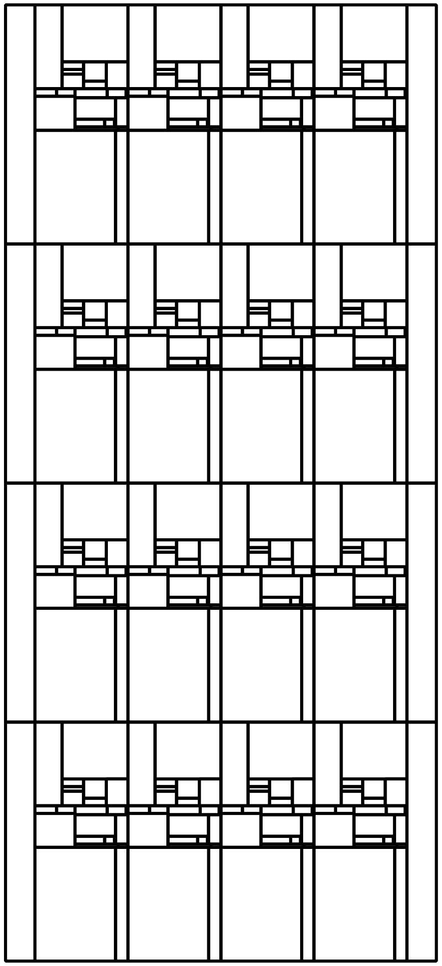

16-Core System
ArchFP, VLSI-SOC'12

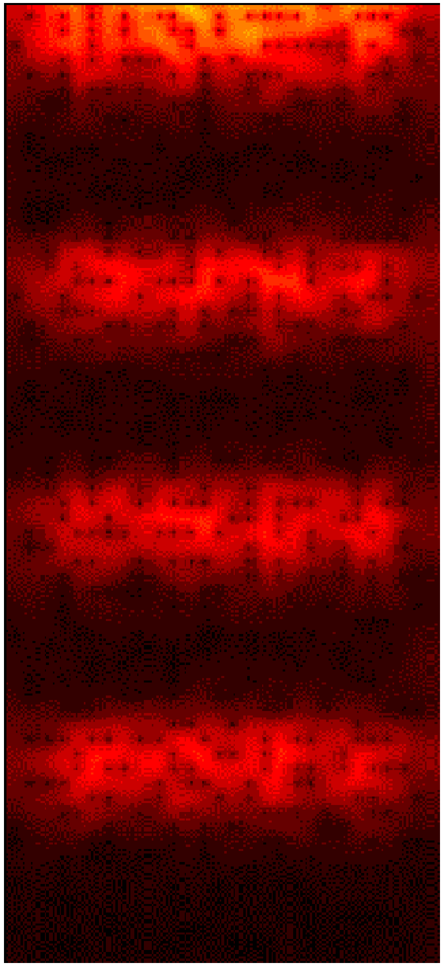

Too Few Pads
(540)

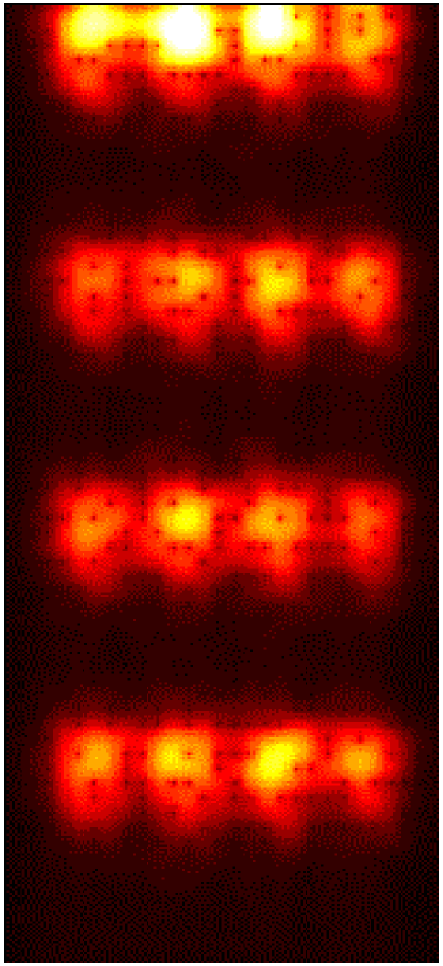

Bad Pad Placement
(960)

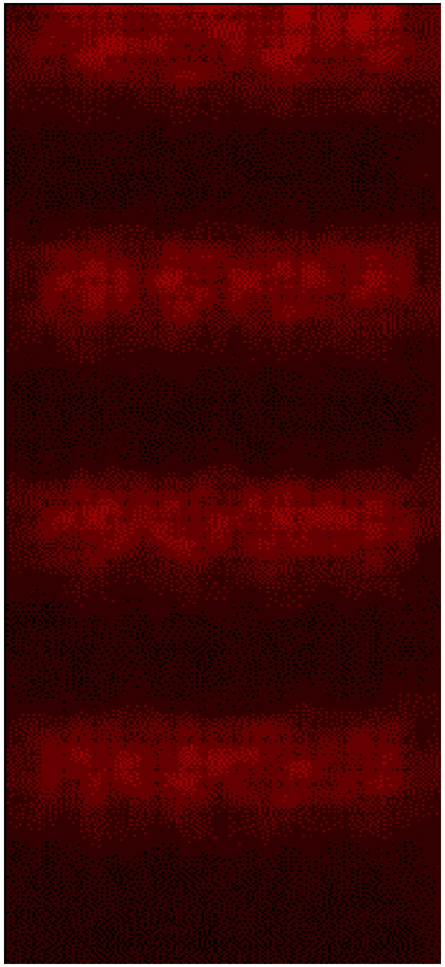

Optimal Placement
(960)

Walking Pads

- Challenges
  - *Massive* design space ($10^{200} - 10^{1400}$)
  - Computationally *expensive* design evaluation
- Recent approaches
  - Placement for pad rings [Sato ASPDAC'05]
  - Analytical models [Rius TVLSI 21(3)]
  - Simulated annealing (SA) [Zhong ASPDAC'07]
- *Idea*: Allow all pads to simultaneously *walk* toward current demands in the system
  - Reduces the number of designs to evaluate
  - Results in *substantial* speedup over SA, up to 634X!

Overview

- Modeling and Optimization Framework
- Experimental Setup and Results
- Analytical Model for Fast Pad Count Optimization
- Conclusions and Future Work

Power Delivery Network Model

- Architectural blocks as ideal current sources
- C4 pads on a coarse grid
- Fine-grained PDN grid modeling
- Given a configuration, solve for $V(x, y)$

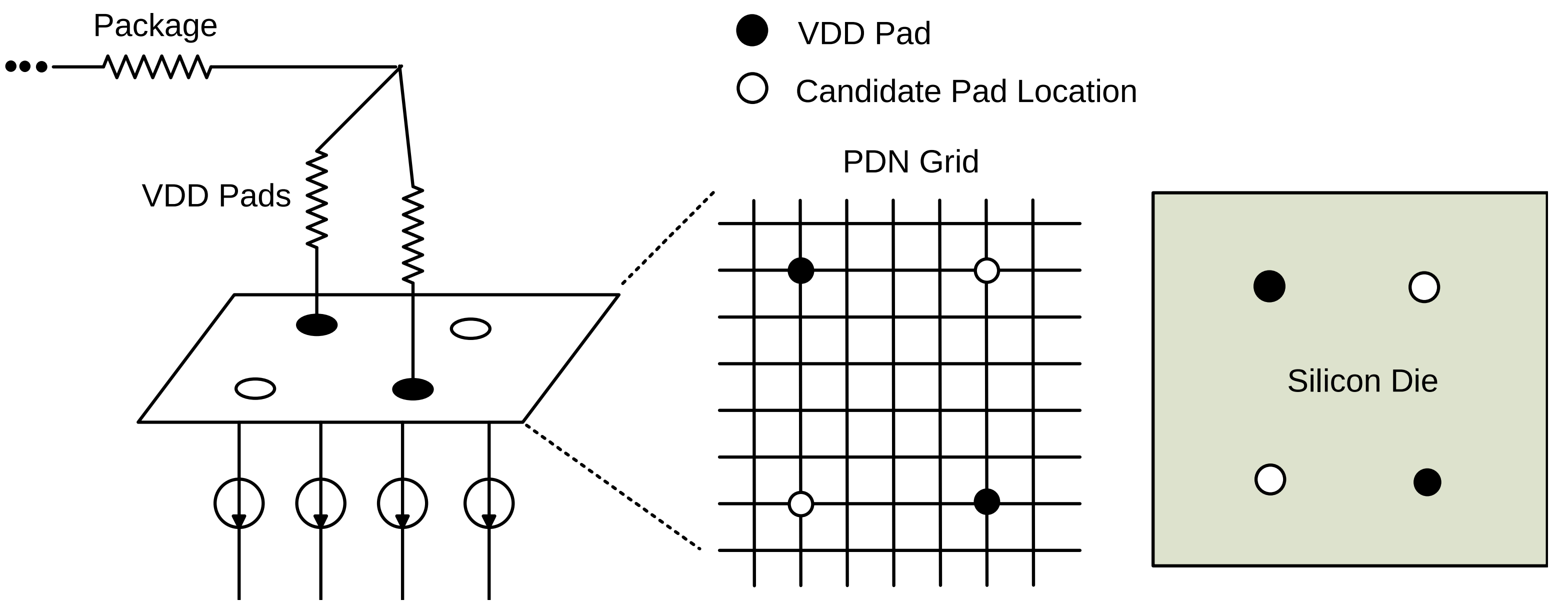

Walking Pads Framework

- PDN voltage field obeys Poisson's equation

$$\frac{\partial^2 V}{\partial x^2} + \frac{\partial^2 V}{\partial y^2} = I_{xy}R$$

- *Key observation*: electrostatic field looks similar

$$\frac{\partial^2 V}{\partial x^2} + \frac{\partial^2 V}{\partial y^2} = \frac{\rho_{xy}}{\varepsilon_{xy}}$$

- Treat pad allocation as a force balance problem
  – Allow virtual electrostatic forces on pads to push and pull them into place

Virtual Forces

- Pads: *mobile positive charges*
- Current sources: *negative surface charges*
- Virtual forces move pads
  - Sources pull pads to optimal locations
  - Other pads push each them away
- Virtual forces calculated using Gauss's Law
  - Direct superposition is too expensive (100s of pads!)
  - Instead, use voltage gradient in N, S, E, and W to determine total *virtual force*

Walking Pads Algorithms

- Basic structure:
  1. Calculate steady-state equations
  2. Determine direction and displacement of movement for each pad using *virtual forces*
  3. Move pads
- WP-N: move all pads in one step, constant displacement, one step in the grid
- WP-F: move all pads in one step, initially large displacement that decreases gradually until frozen
- WP-R: greedily move one pad at a time (pads sorted by distance to max IR drop location)

WP Ex: WP-F Optimization

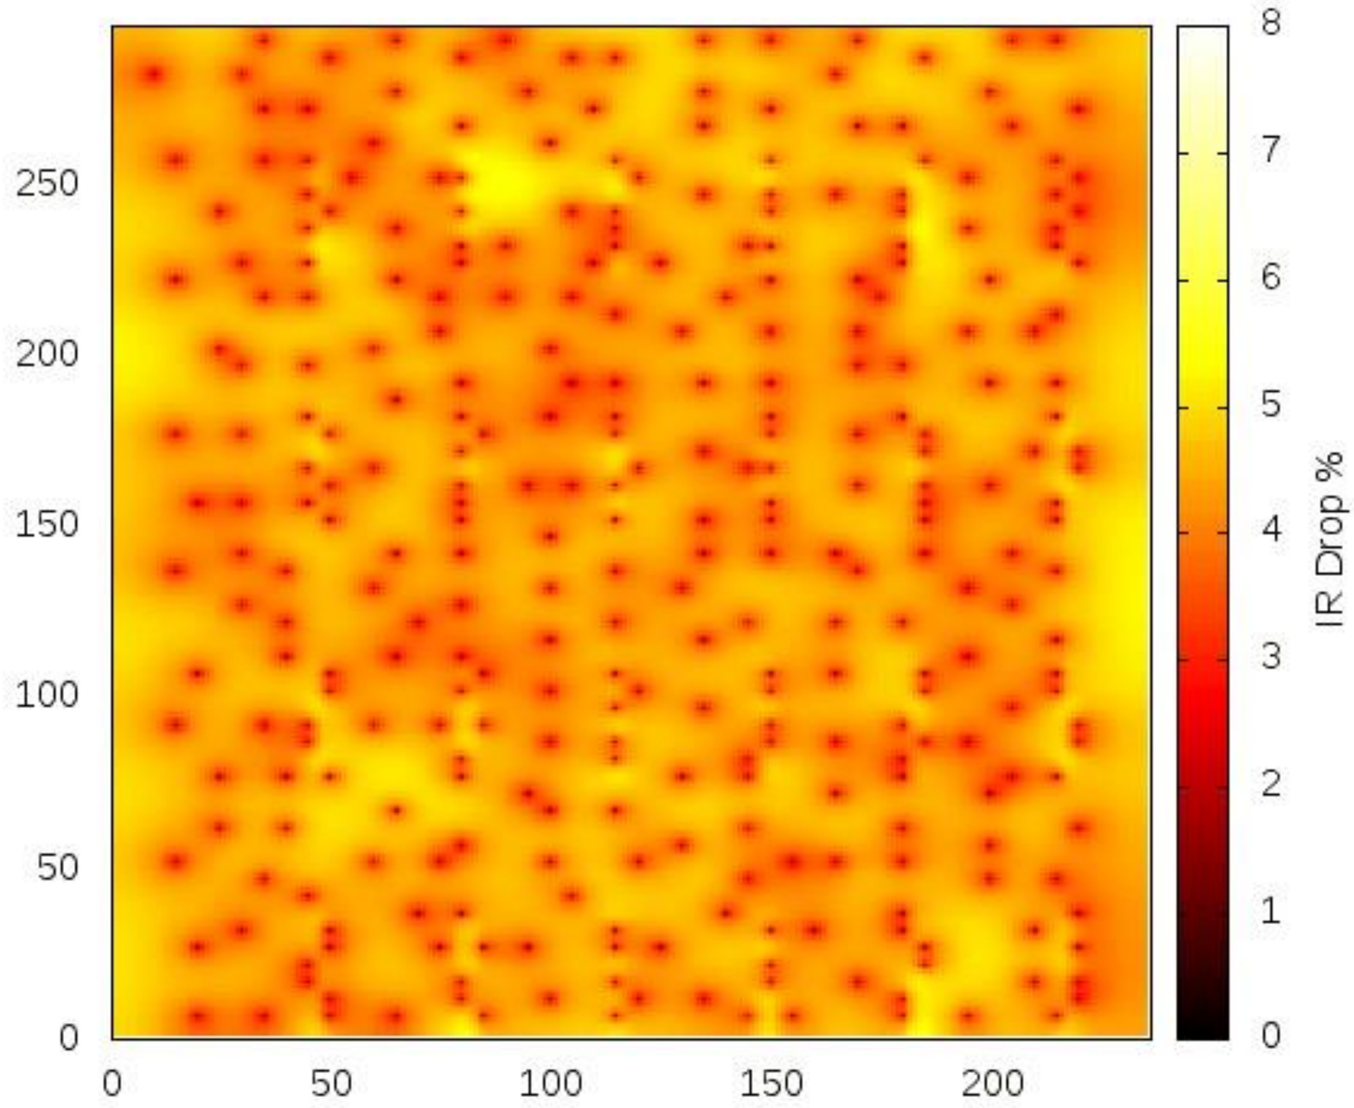

WP Ex: WP-R Optimization

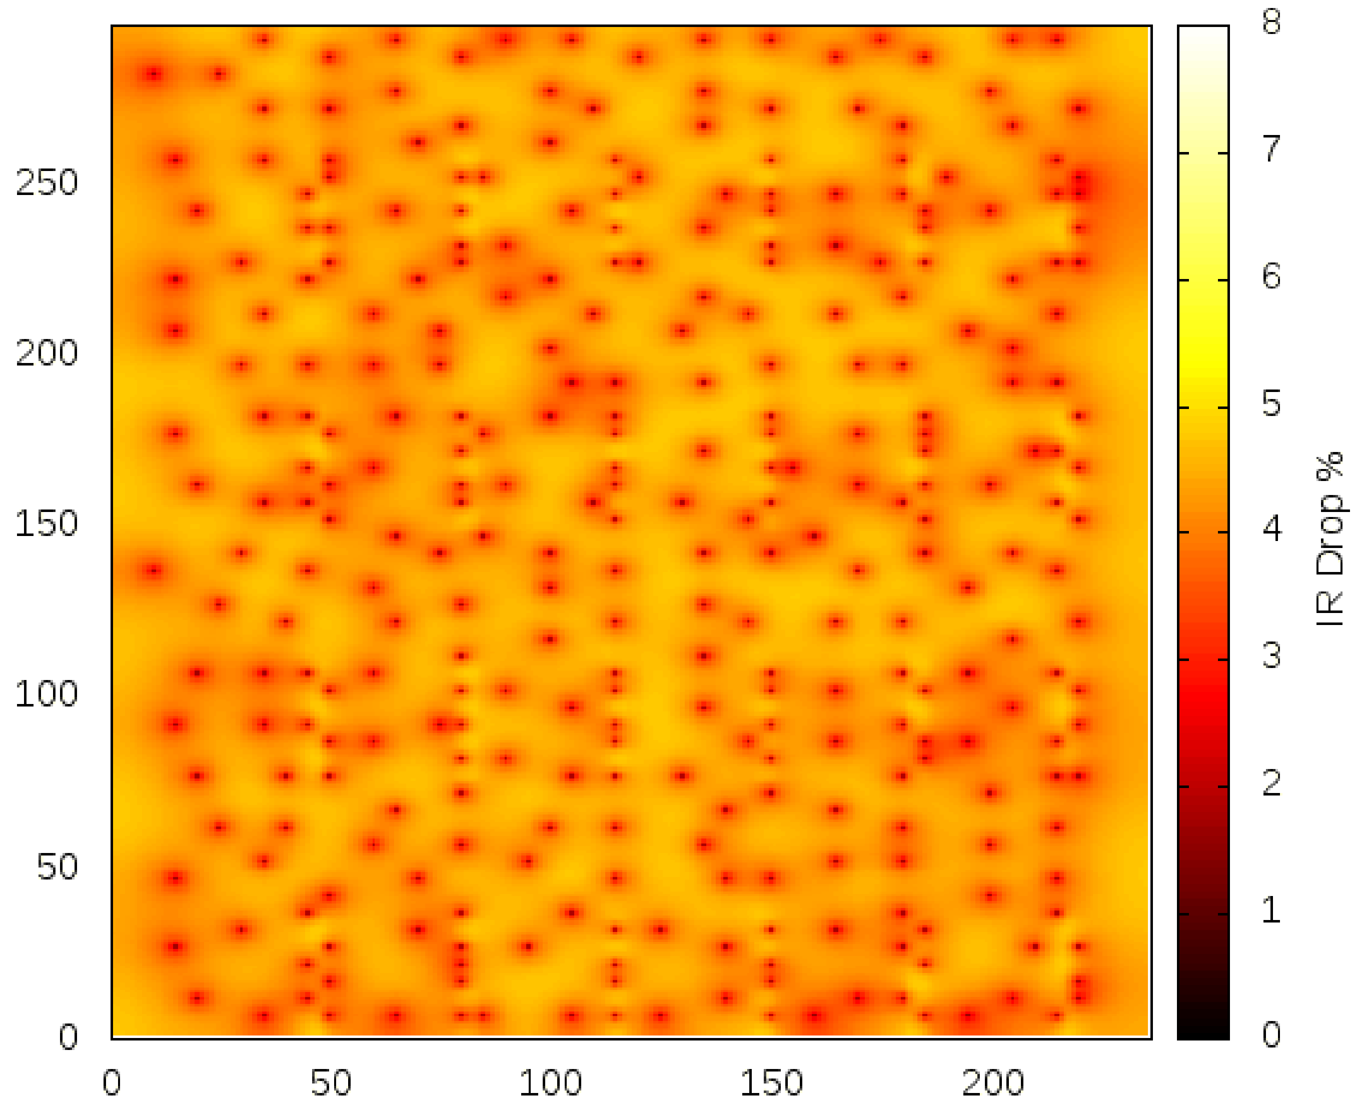

Experimental Setup

- Architectures
  - 16 to 48-core multicore systems
  - Synthetic benchmarks
- Comparison with Zhong and Wong (ASP-DAC'07)
  - Simulated annealing technique
  - Cooling rates of 0.999 (SA-S– ground truth) and 0.98 (SA-P– practical optimization)

Results: Pad Placement Quality

- 24 cores, 180 pads
- McPAT for power (85% of peak)
- Uni: 12.5% IR
- SA-P: 6.9% IR
- WP-N: 10.2%
- WP-F: 7.5%
  - 157X faster
  - 0.6% higher

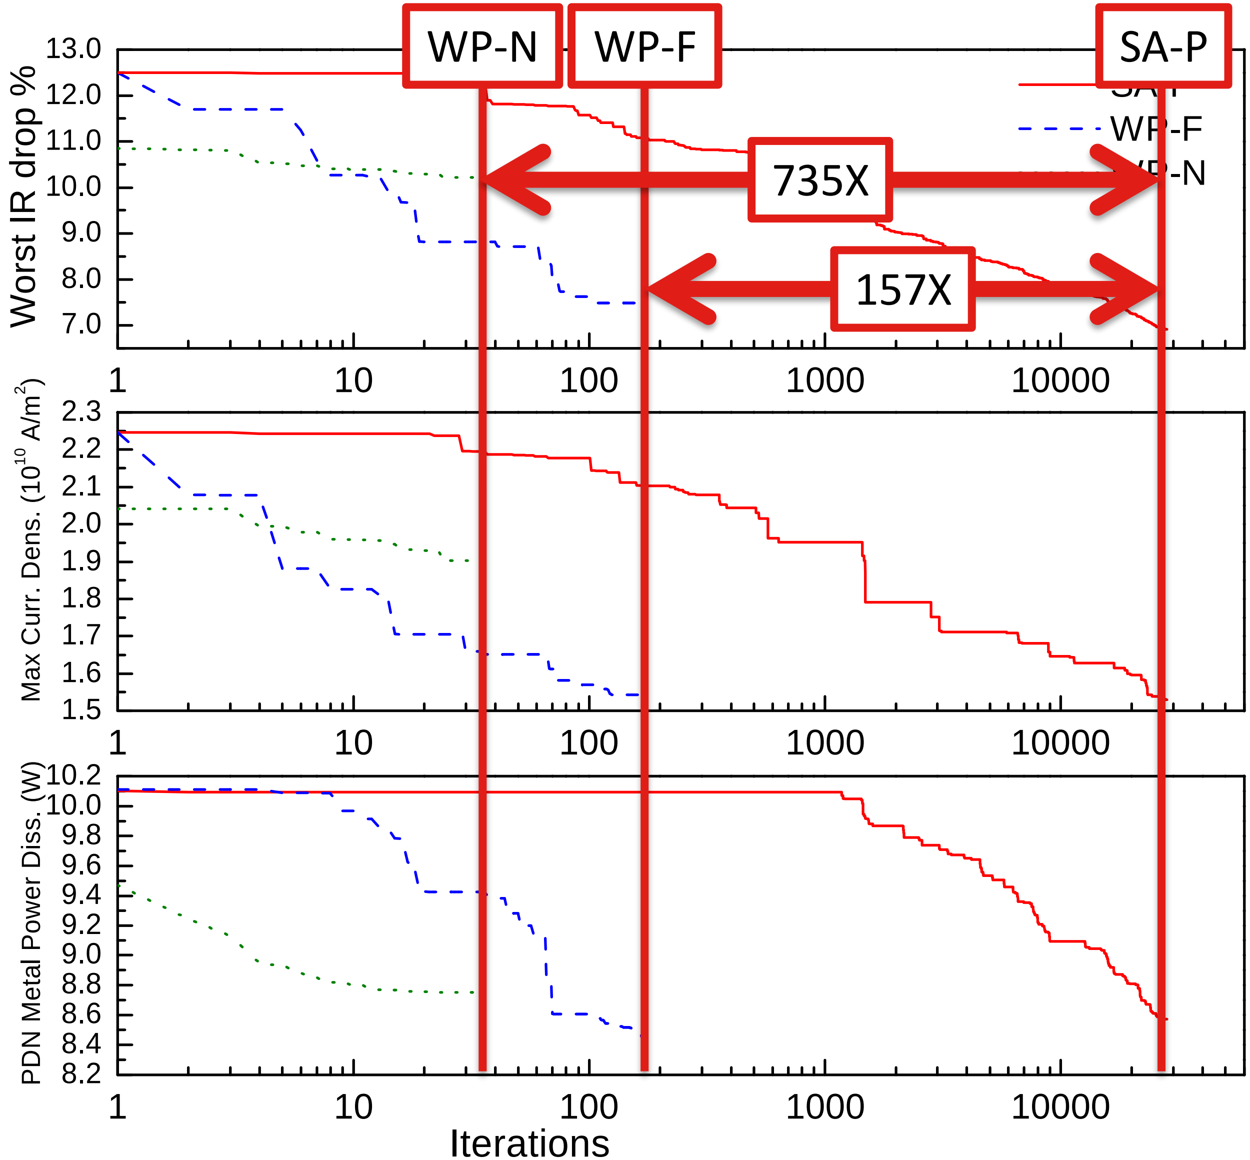

Results: WP-F and WP-R, 24 Cores

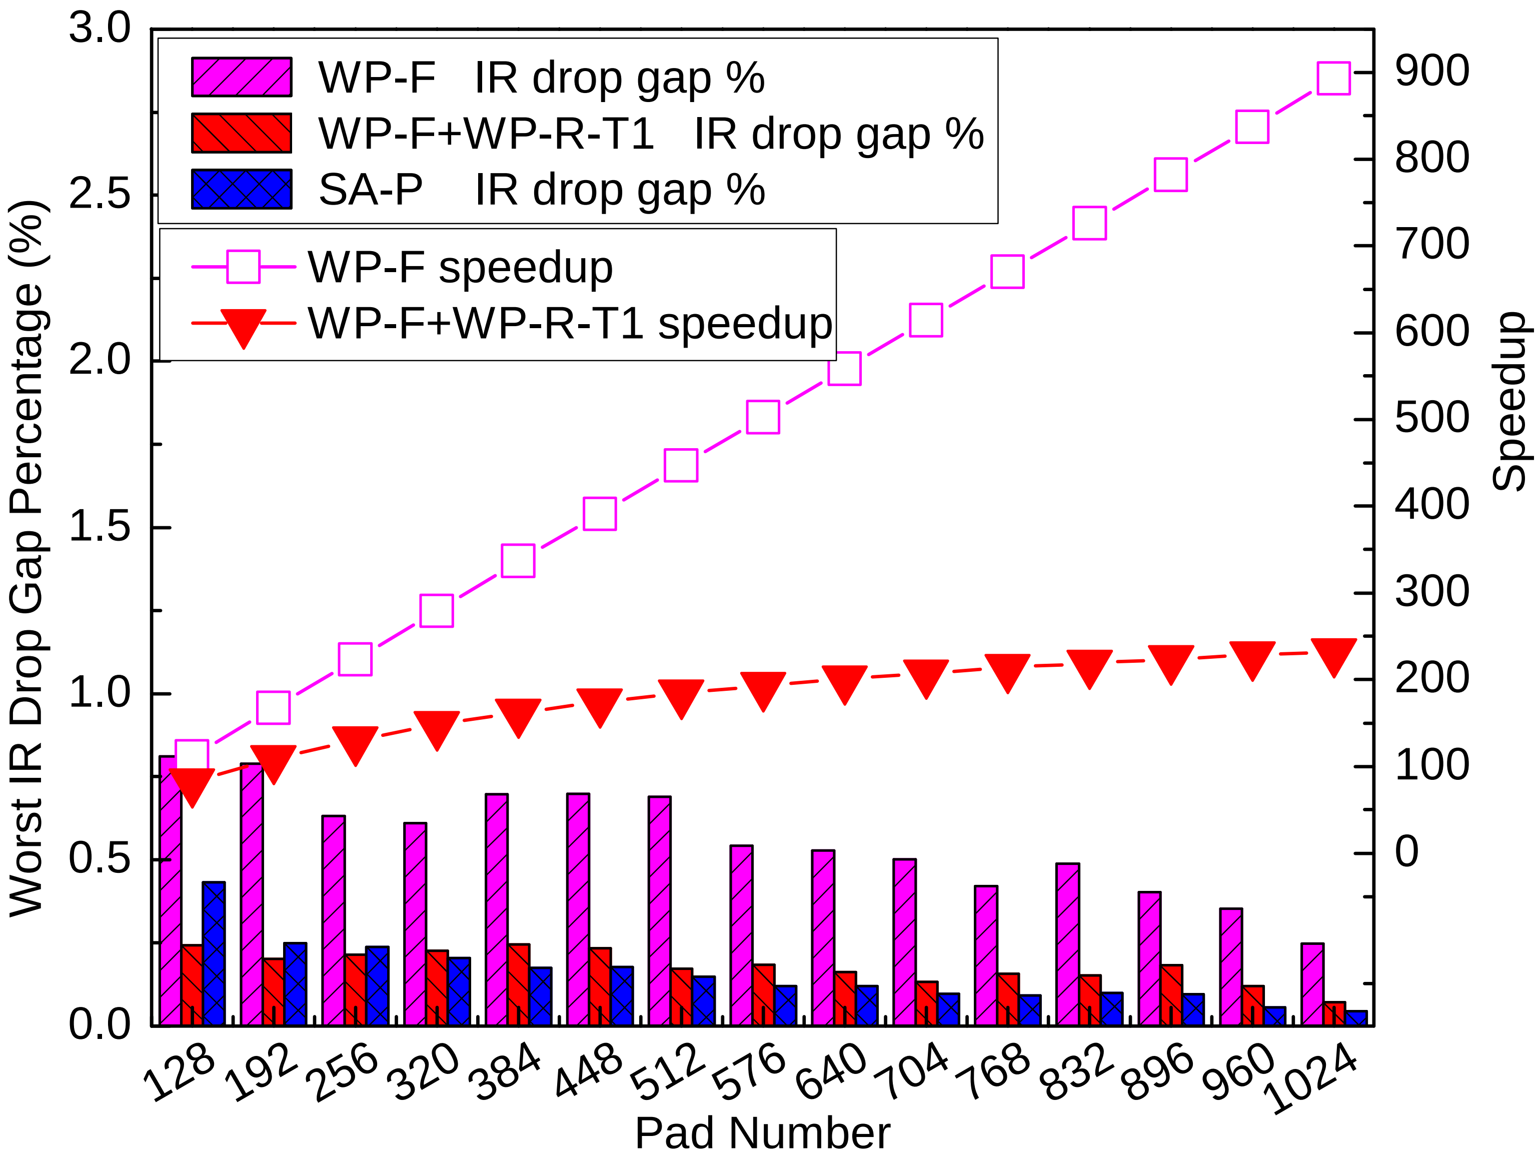

Synthetic Benchmarks

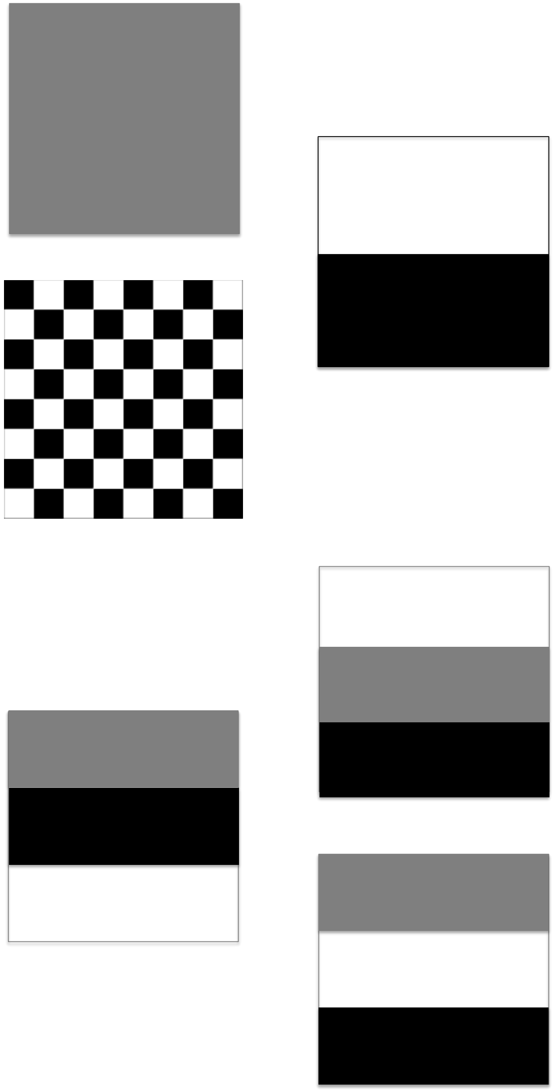

- All 20x20 mm^2, 150 W total
- Uniform (Uni)
- Half-half (HH, 4:1, black:white)
- Checkerboard (CB, 4:1, black:white)

- Three-level variations (TL*x*)
- 3:2:1, black:gray:white

Results: Scaled and Synthetic Systems

Bench.	# pads	# loc	Speedup		% Gap		
			F	R-T1	F	R-T1	R
S-Uni	512	4900	498	206	0.18	-0.03	-0.11
S-HH	512	4900	498	206	0.23	-0.03	-0.11
S-CB	512	4900	498	206	0.20	-0.06	-0.12
S-TL1	512	4900	498	206	0.15	-0.03	-0.10
S-TL2	512	4900	498	206	0.24	0.01	-0.12
S-TL3	512	4900	498	206	0.19	-0.03	-0.11
16-Core	512	1914	375	155	0.42	0.16	-0.07
24-Core	768	2880	670	277	0.33	0.071	-0.04
32-Core	1024	3844	961	397	0.41	0.070	-0.07
48-Core	1536	5776	1536	634	0.39	0.055	-0.09

How Many Pads?

- WP: Given *n* pads, determines placement
- How do we determine *n*? An analytical model.
- Simplifying assumptions:
    1. On-chip current density is uniform
    2. All pad currents are equal (I_0)
    3. Each pad serves a circular area with radius r_0
- $$V_{drop} = a\frac{1}{N_p}\log\left(\frac{1}{N_p}\right) + b\frac{1}{N_p} + c,$$

  where a, b and c are constant functions of I_0, r_0, r_ε, ρ, l, or the voltage drop across the package

PDN Design Space

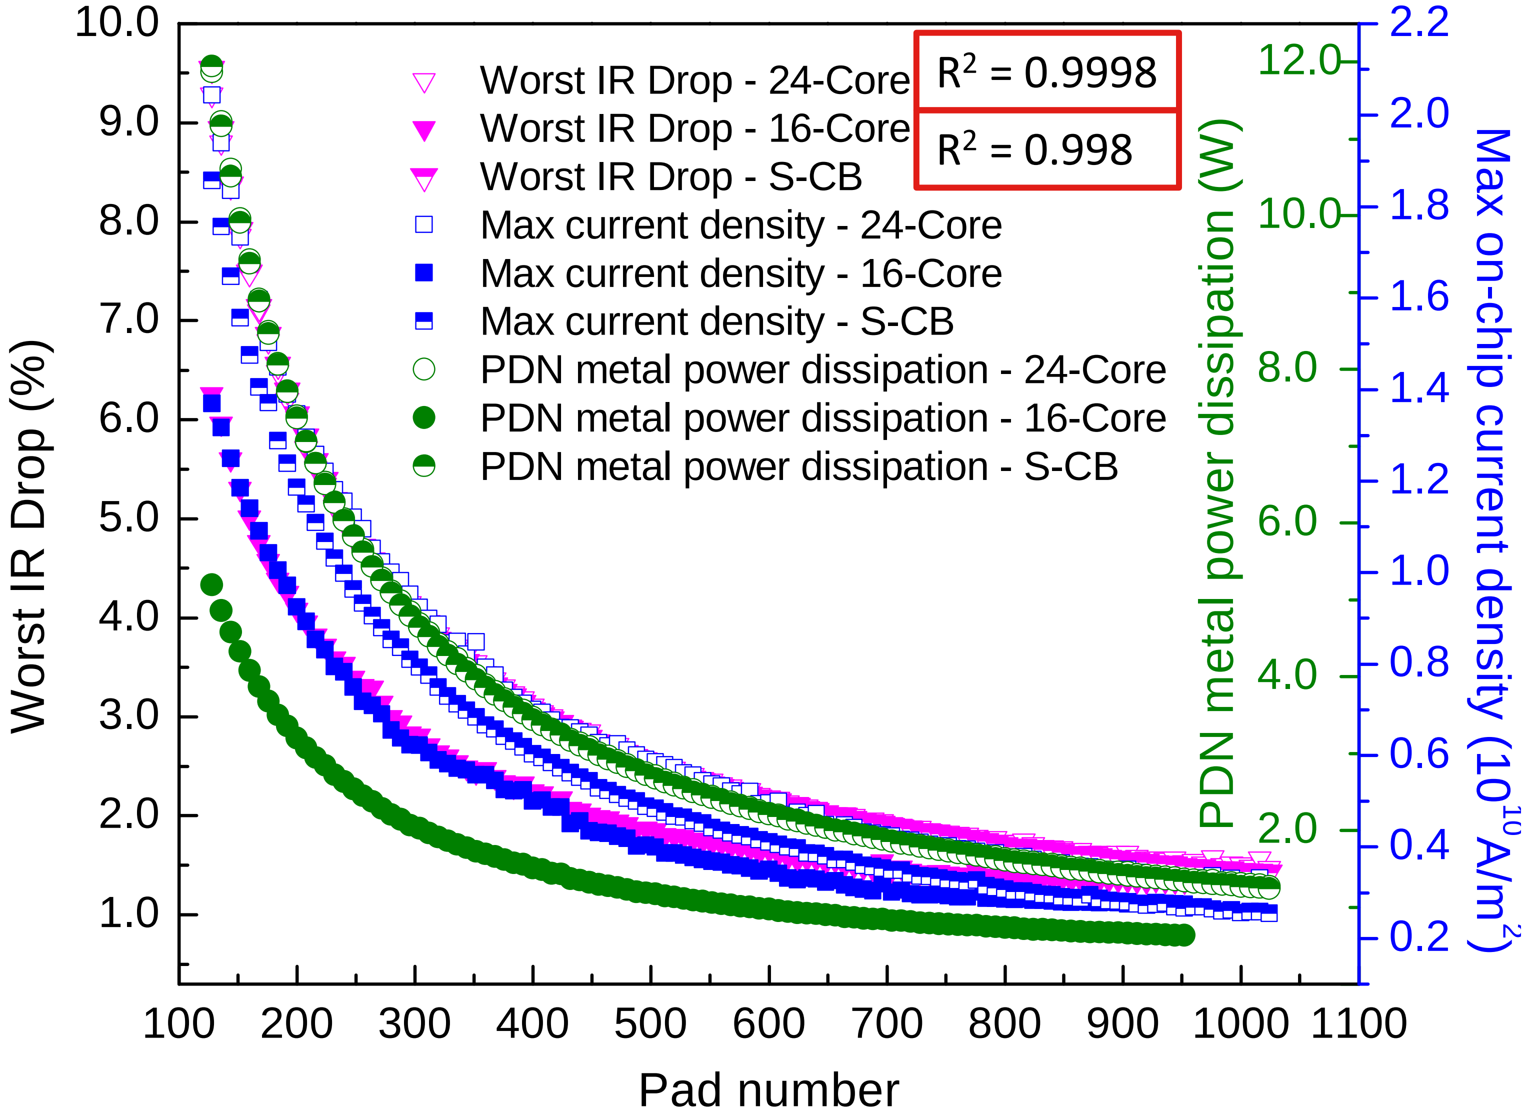

Fast Pad Count Optimization

1. Choose three pad counts; run WP; fit *a*, *b*, *c*.
2. Predict the minimum pad number *n* given an IR drop budget.
3. Perform a local search around *n*.

IR Drop Budget	Pred.	Optimal	Actual IR Drop
5%, 35mV	240	238	34.63mV
4%, 28mV	304	306	27.97mV
3%, 21mV	416	418	20.77mV
2%, 14mV	673	672	13.99mV

24-core System

Conclusions

- Walking Pads
  - Improves pad location optimization by allowing all pads to *walk* toward their local balance positions
  - Results in speedup of up to 634X relative to simulated annealing (SA)
  - Produces allocations within 0.1% of the IR drop of SA
- Analytical Model for Fast Pad Count Optimization
  - Relates pad count and max IR drop; performs fast, effective prediction after minimal evaluation and fitting

Ongoing and Future Work

- VDD-GND joint optimization
  - So far, only VDD placement
  - *Challenge*: requires transient voltage noise simulation
- Pad optimization for transient voltage noise
  - So far, only steady-state
  - *Challenge*: transient voltage noise simulation is *very, very* computationally expensive
- Pad optimization with spatial constraints
  - So far, unrestricted placement
  - I/O pads must be placed relative to I/O blocks, constraining the available pad locations
- IR-drop aware floorplanning, thermal-aware design, and power delivery for 3D ICs, too!