DATE 13
Design, Automation & Test in Europe
18-22 March, 2013 - Grenoble, France
The European Event for Electronic System Design & Test

Cherry Picking: Exploiting Process Variations in the Dark Silicon Era

Siddharth Garg
University of Waterloo

Co-authors:
Bharathwaj Raghunathan, Yatish Turakhia and Diana Marculescu

Dark Silicon Challenge

- **Transistor dimensions scale by a factor of *S***
  - S^2 more transistors in same die area
  - Each transistor dissipates $\approx \frac{1}{S}$ times power
  - Fixed power budget results in **dark silicon**

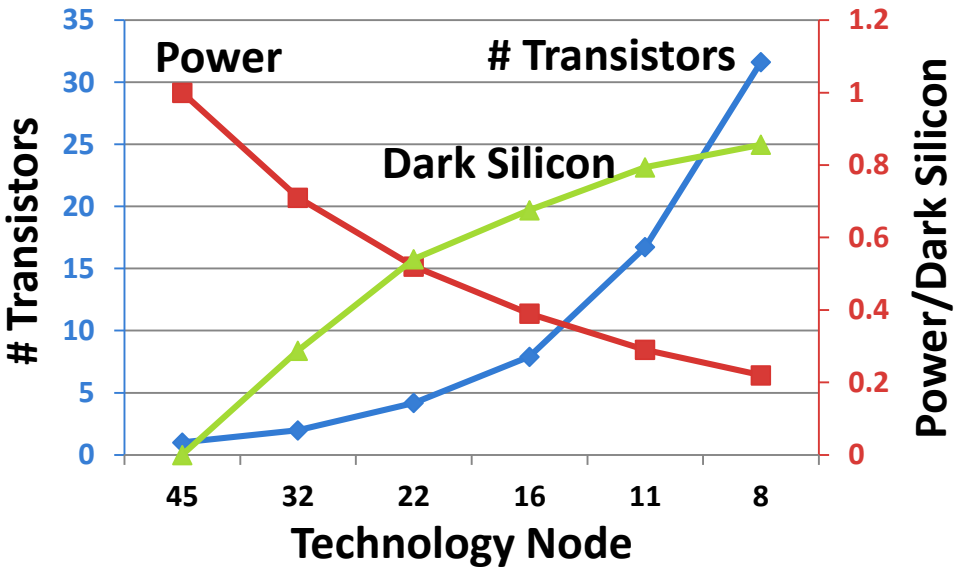

80% dark silicon at the 8 *nm* technology node

[Esmaeilzadeh *et al.*, ISCA'11]

Progress Hits Snag: Tiny Chips Use Outsize Power

By JOHN MARKOFF
Published: July 31, 2011

For decades, the power of computers has grown at a staggering rate as designers have managed to squeeze ever more and ever tinier transistors onto a silicon chip — doubling the number every two

RECOMMEND
TWITTER
LINKEDIN

Dark Silicon Architectures

- **How to best utilize dark silicon for performance enhancement?**
  - Heterogeneity

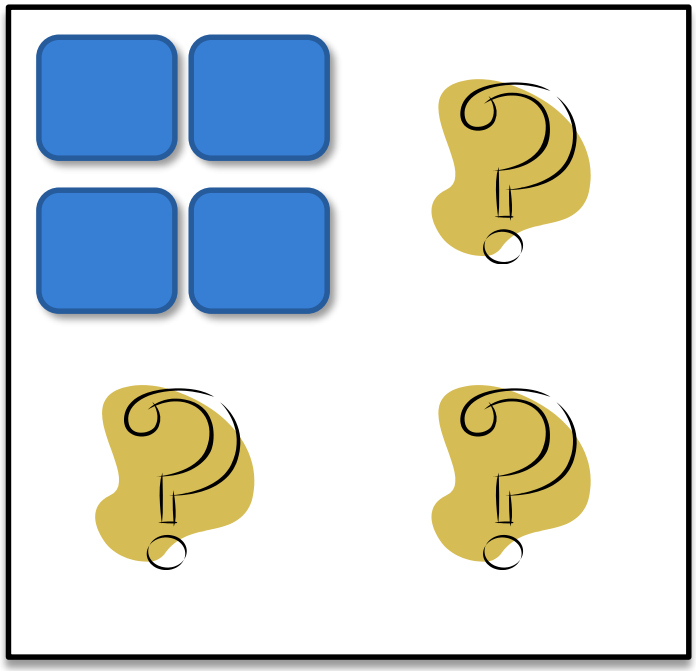

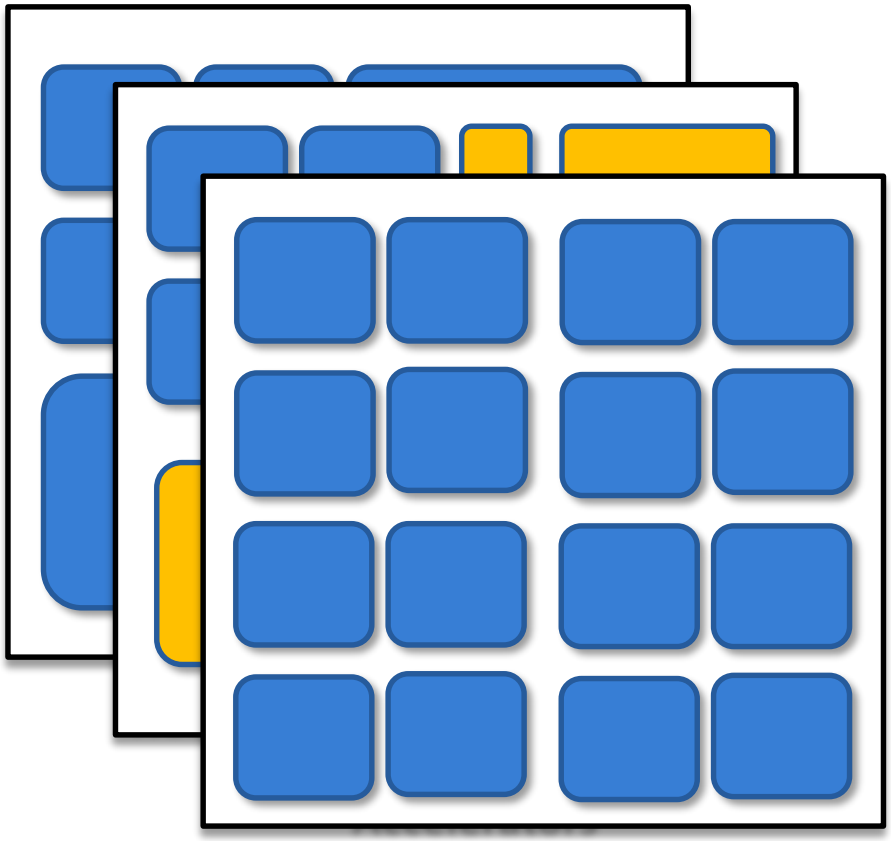

Homogeneous Cores?

Process Variations

- **Inability to precisely manufacture transistors**
  - Chip-to-chip variations
  - Within-chip variations

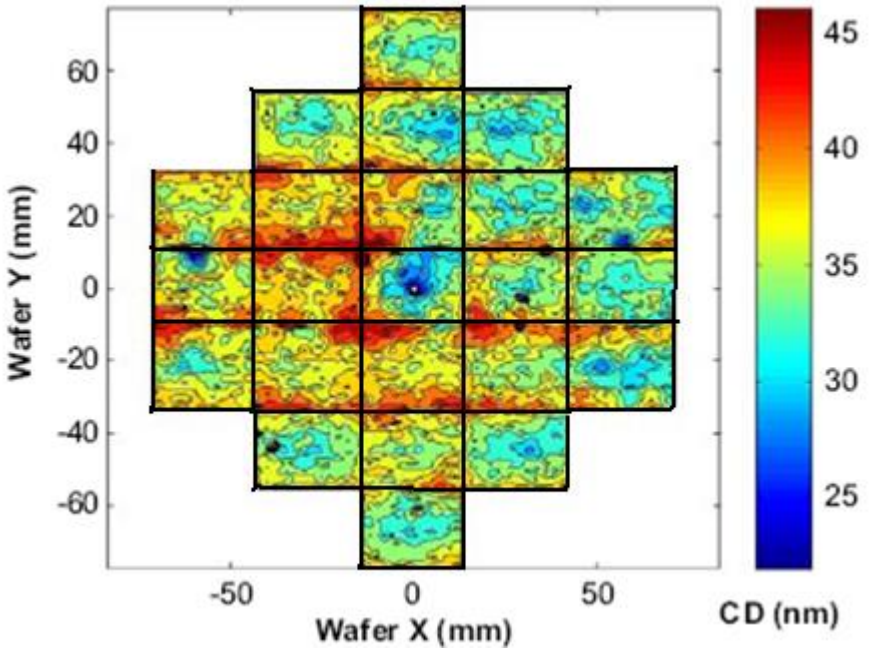

[Source: Friedberg *et al.*, ISQED'05]

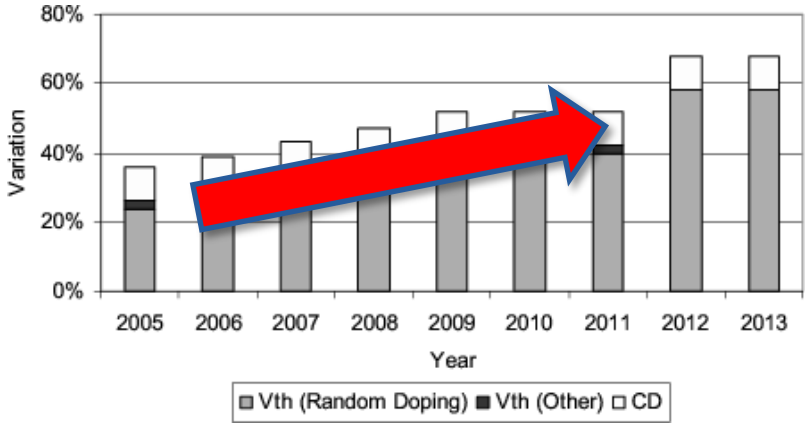

Increasing proportion of within-chip variations

Process Variation Impact

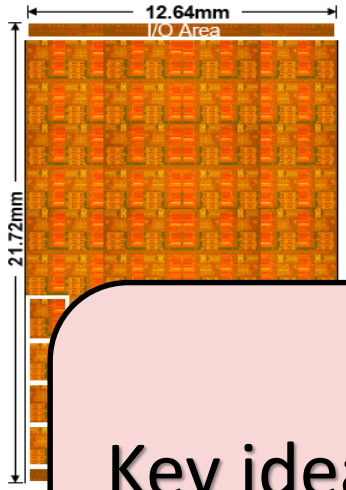

Intel 80-core Teraflop

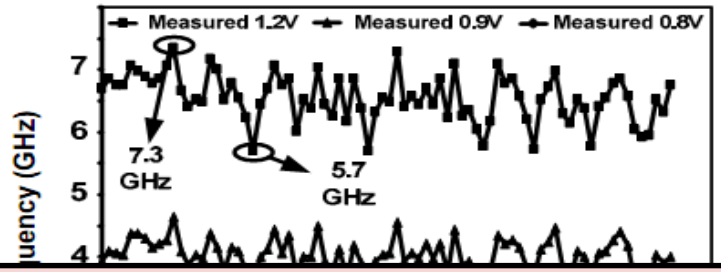

Key idea: exploit heterogeneity that arises from the impact of process variations

[Dighe *et al.*, JSSC'11]

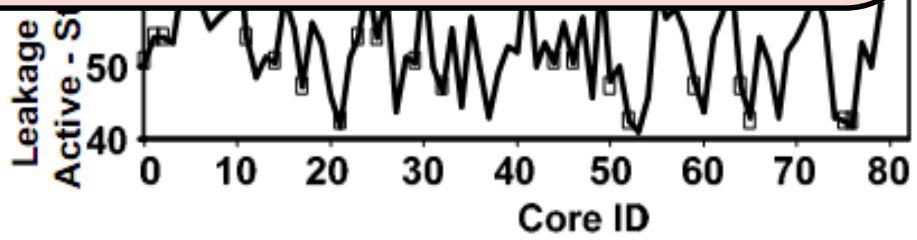

1.7X deviation in leakage power

Motivation: Best 1 of *N*

- **Best 1 of *N* statistics**
  - Provision chip with *N* identical cores and **cherry-pick** core with highest frequency

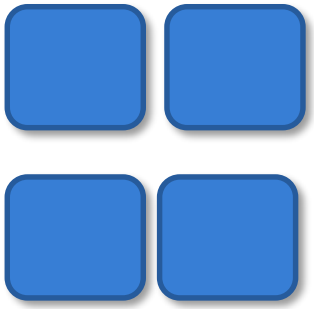

Best 1 of 4
Mean = 1.2

1-σ increase in frequency with core doubling

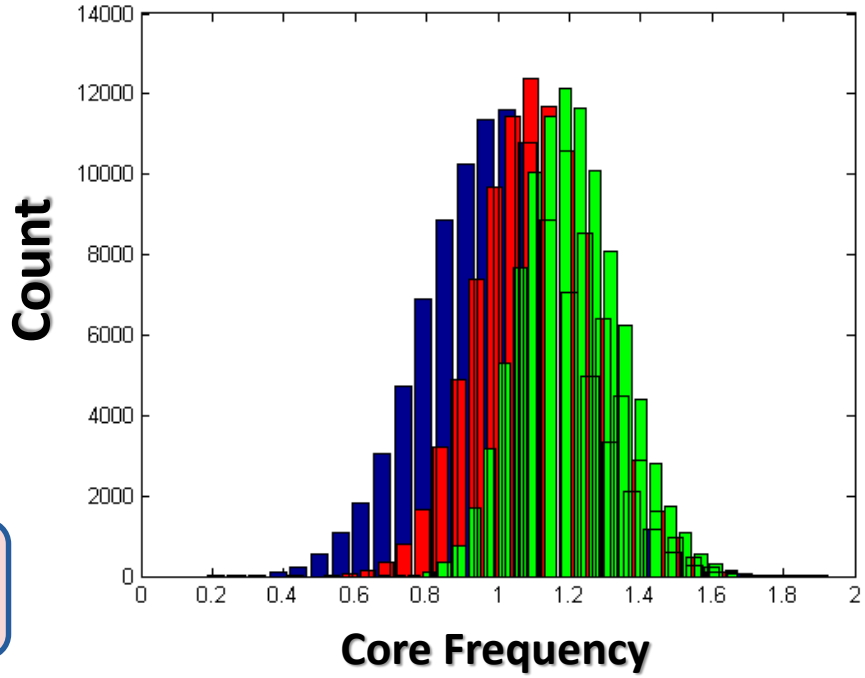

Best 1 of *N* for Leakage

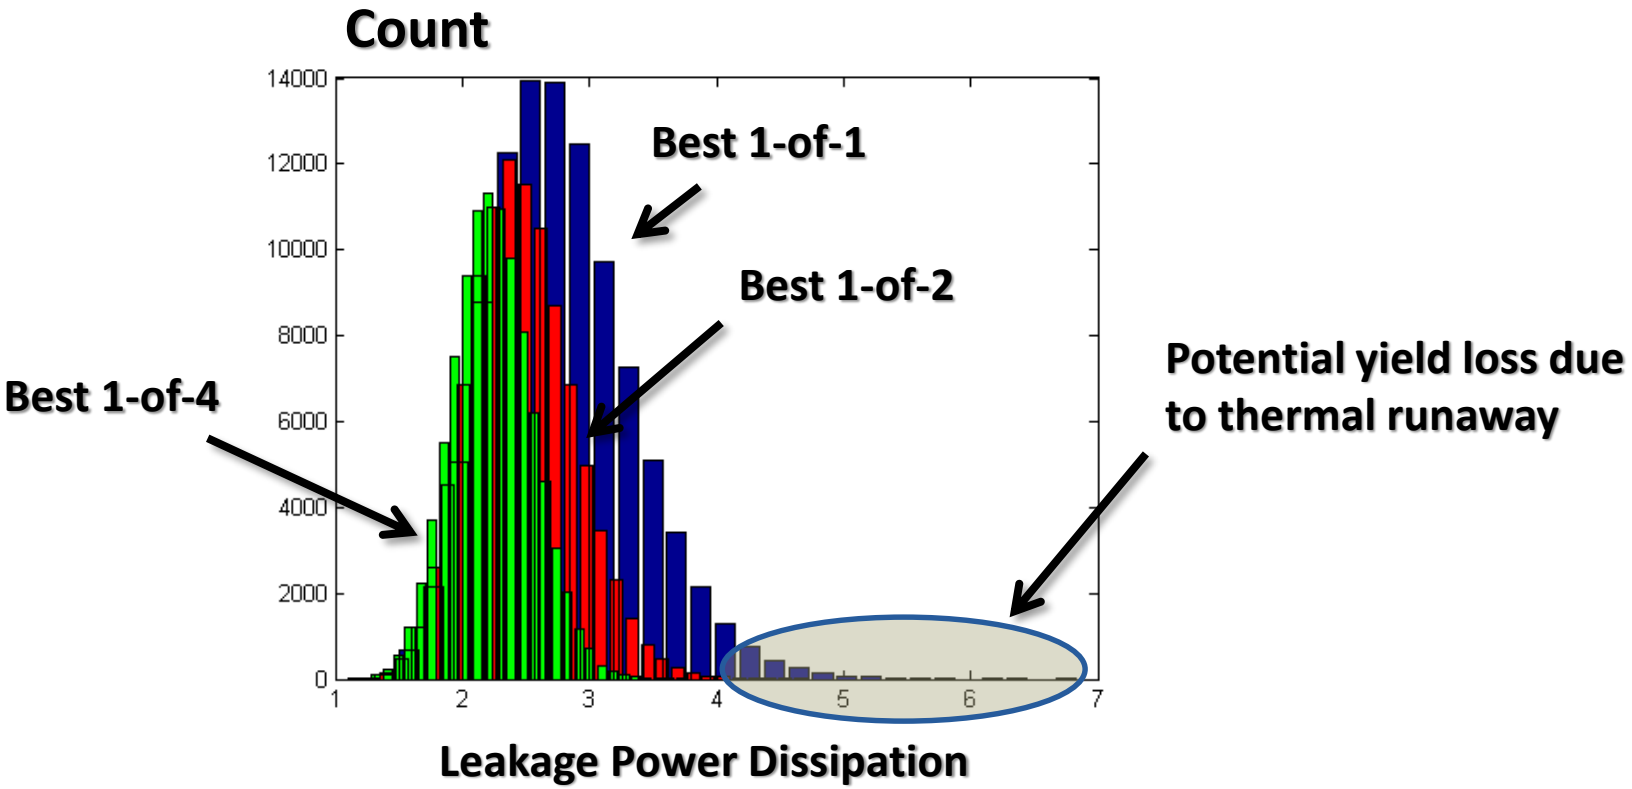

- 30% reduction in average leakage power
- 2X reduction in worst-case leakage power

Related Work

- **BubbleWrap [Karpuzcu et al.,MICRO'09]**
  - Use redundant cores to increase lifetime
  - Cores run in Turbo mode till they "pop"
- **Dark silicon architectures**
  - Heterogeneous cores [Esmaeilzadeh et al.,ISCA'11]
  - Accelerators [Venkatesh et al.,MICRO'11]
- **Statistical Element Selection**
  - Increasing immunity of analog circuits to process variations [Keskin et al.,CICC'10]
- **Process variation aware scheduling**
  - ILP based solution for multi-programmed apps [Teodorescu et al.,ISCA'08]

Variability Modeling

- **Generate die map of process variations**

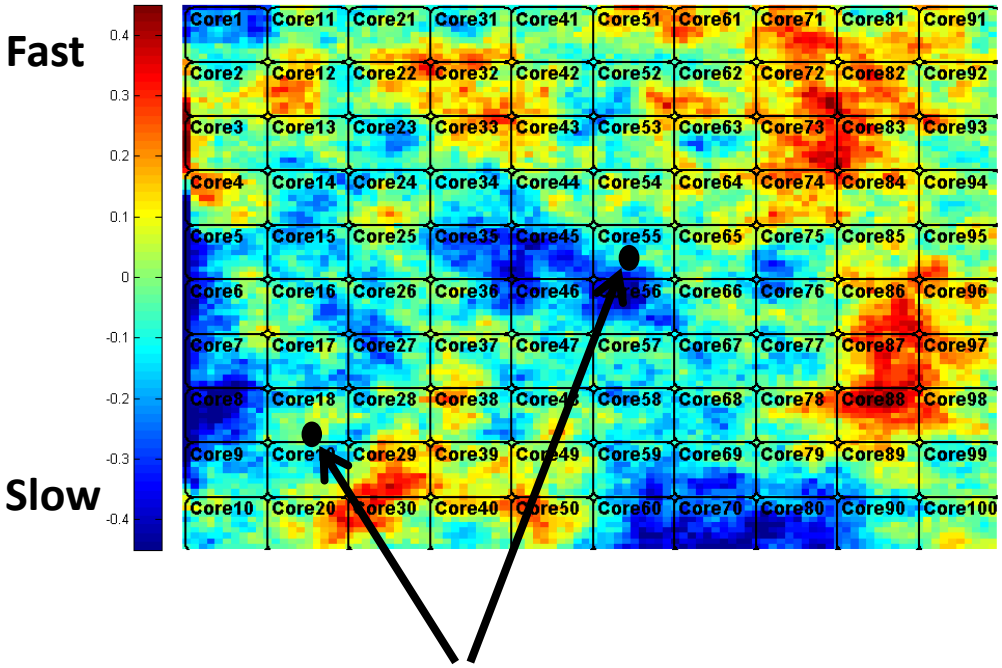

Single Gaussian random variable to model impact of process variations at each location

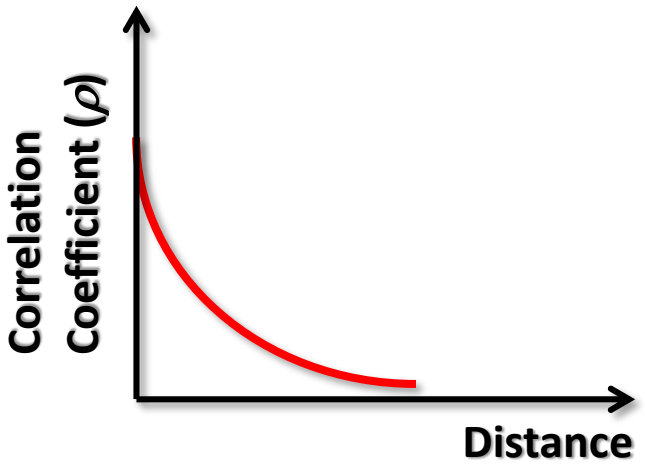

Spatial correlations modeled using an exponentially decaying function of distance

[Zhiong *et al.*, TCAD'07]

Frequency and Leakage

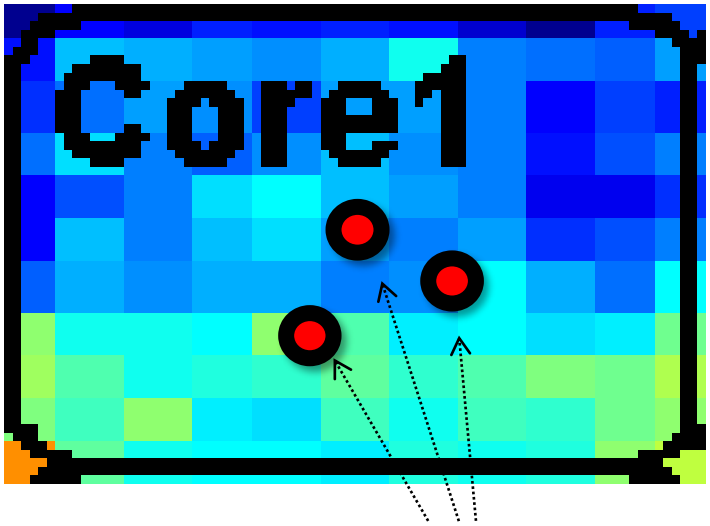

Critical Paths (CP)

- **Each core has N_{cp} identical critical paths**
  - Core frequency limited by slowest critical path
  - Critical path delay inversely proportional to process parameter

$$\textbf{Frequency} \propto \min_{i \in N_{cp}} \frac{1}{p_i}$$

- **Leakage is summed over all N_{core} grid points**
  - Exponential dependence on process parameters

$$\textbf{Leakage} \propto \sum_{i=1}^{N_{core}} \gamma e^{\alpha p_i}$$

Cherry Picking for Single Threads

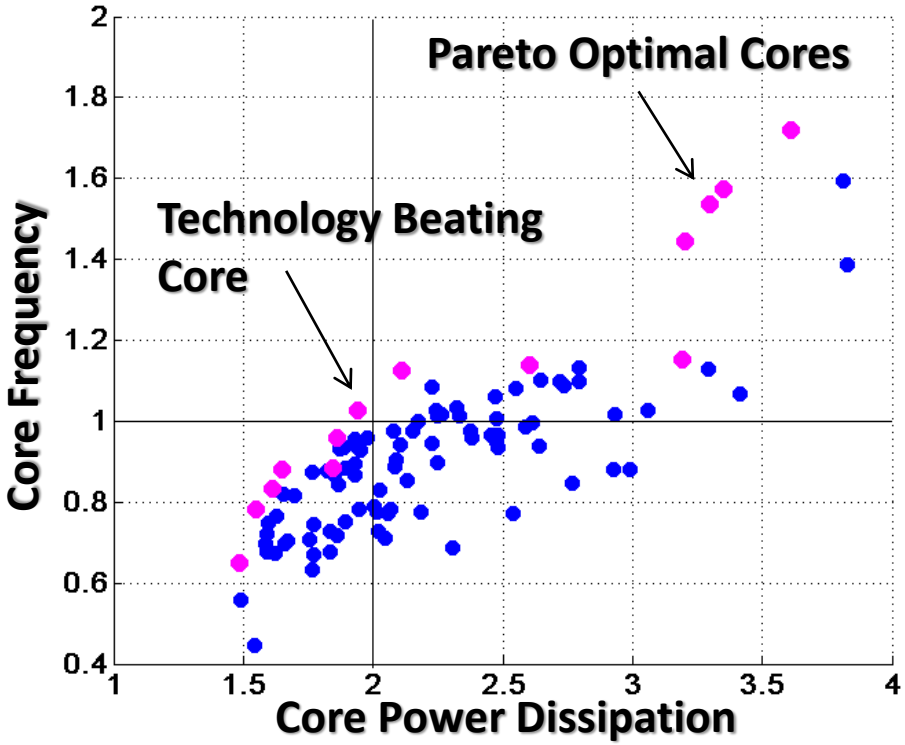

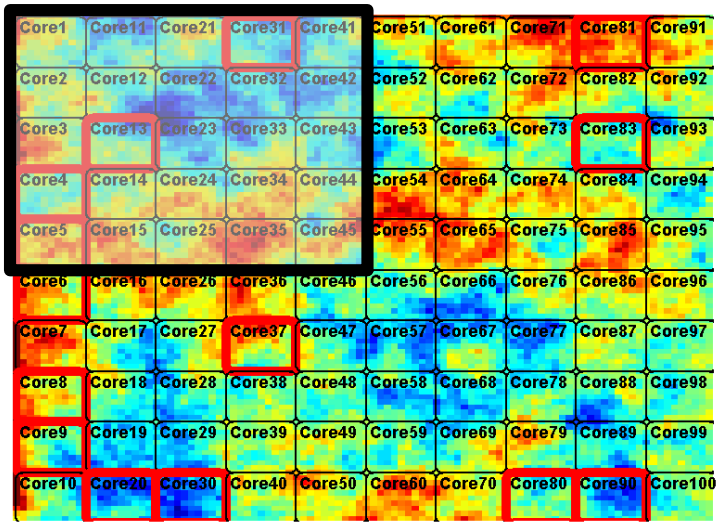

Only 4 Pareto optimal cores in the original design without spare cores

- **Wide range of power and frequency values**
- **One "technology beating" core**
  - Likelihood increases with more % dark silicon

Cherry Picking: Multi-program Workloads

- **Maximize performance within a P Watt budget**
  - Performance measured as the sum of frequencies of cores that are selected

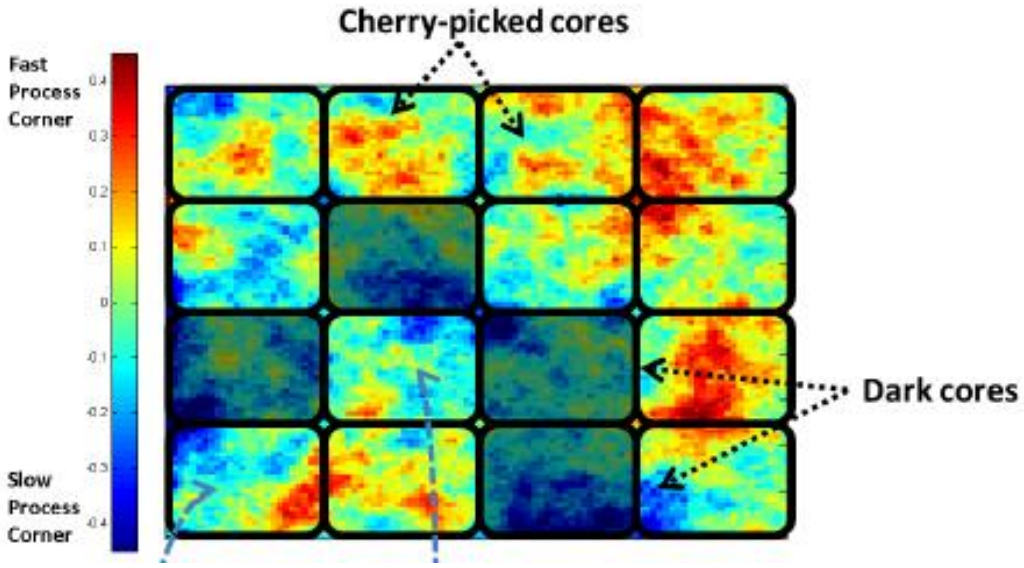

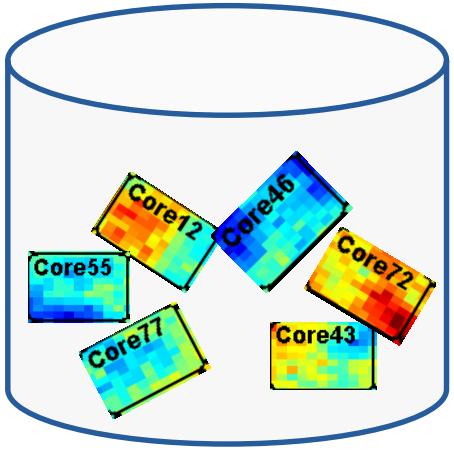

P Watt Bin

Instance of the knapsack problem
Pseudo-polynomial time solution

Cherry Picking: Multi-threaded Wkloads

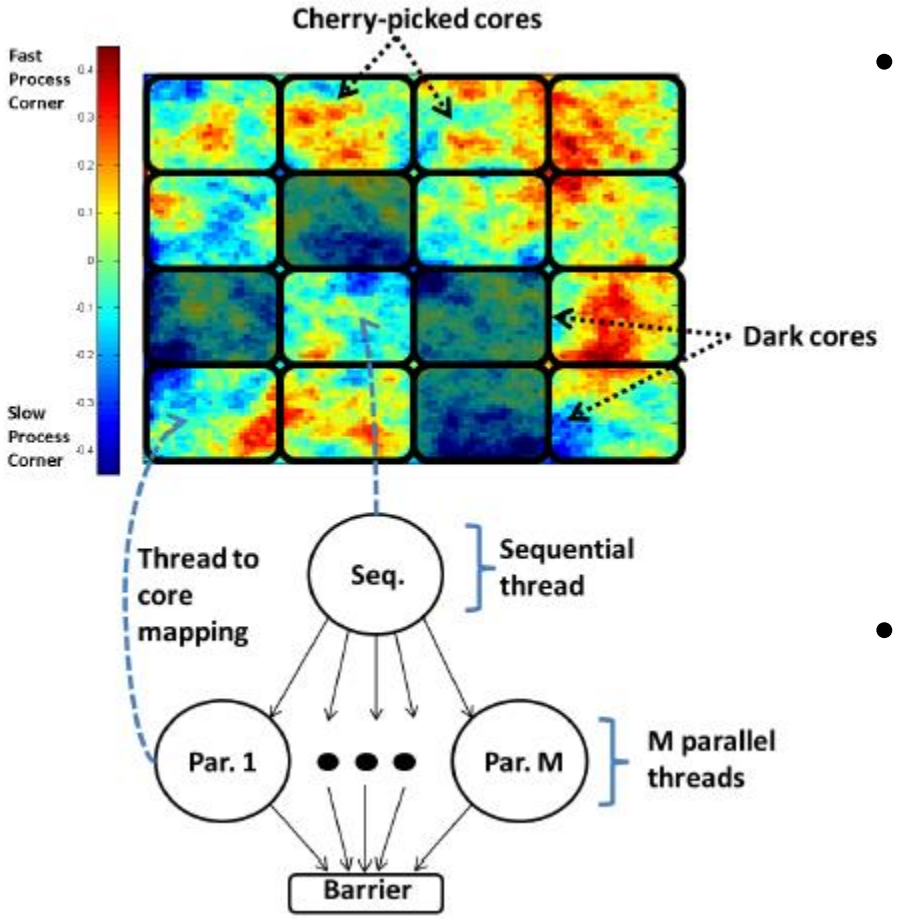

- **Common execution template for a number of parallel benchmarks**
  - Sequential phase followed by barrier based synchronization of parallel threads
- **Optimal mapping of threads to cores such that:**
  - Performance is maximized within a power budget

Performance Model

- **Goal: analytical + accurate performance model that is amenable to optimization**

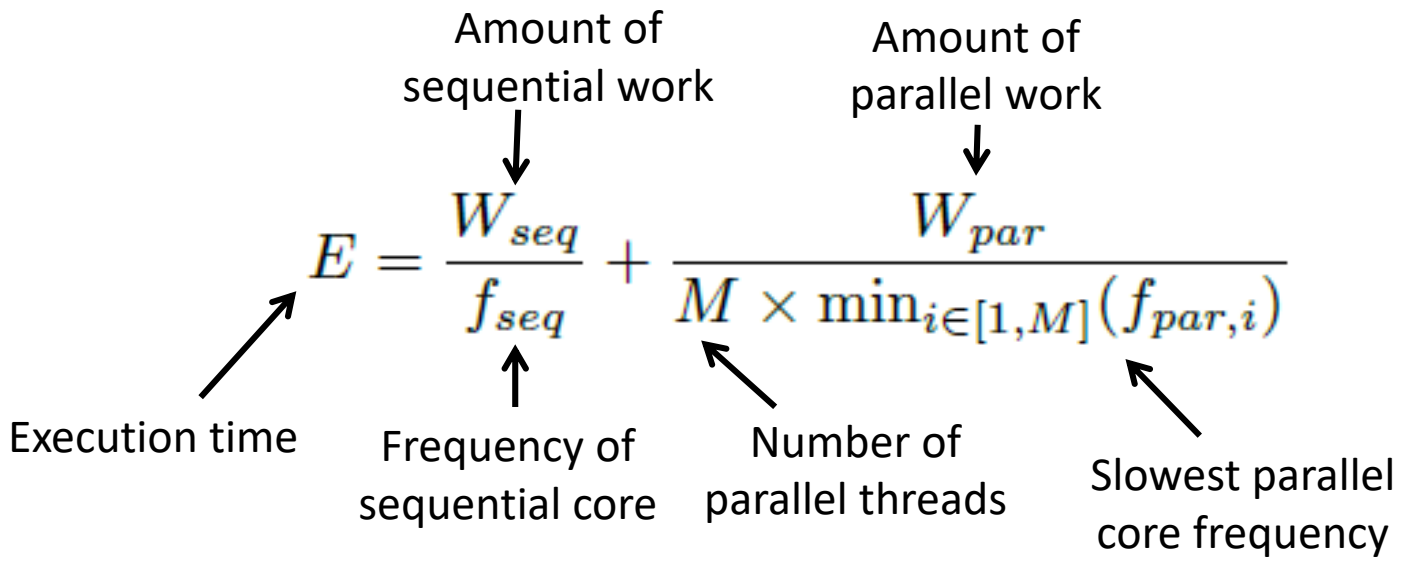

- **Execution time limited by sequential thread and slowest parallel thread**
  - Surprisingly accurate, although grossly simplified

Validation

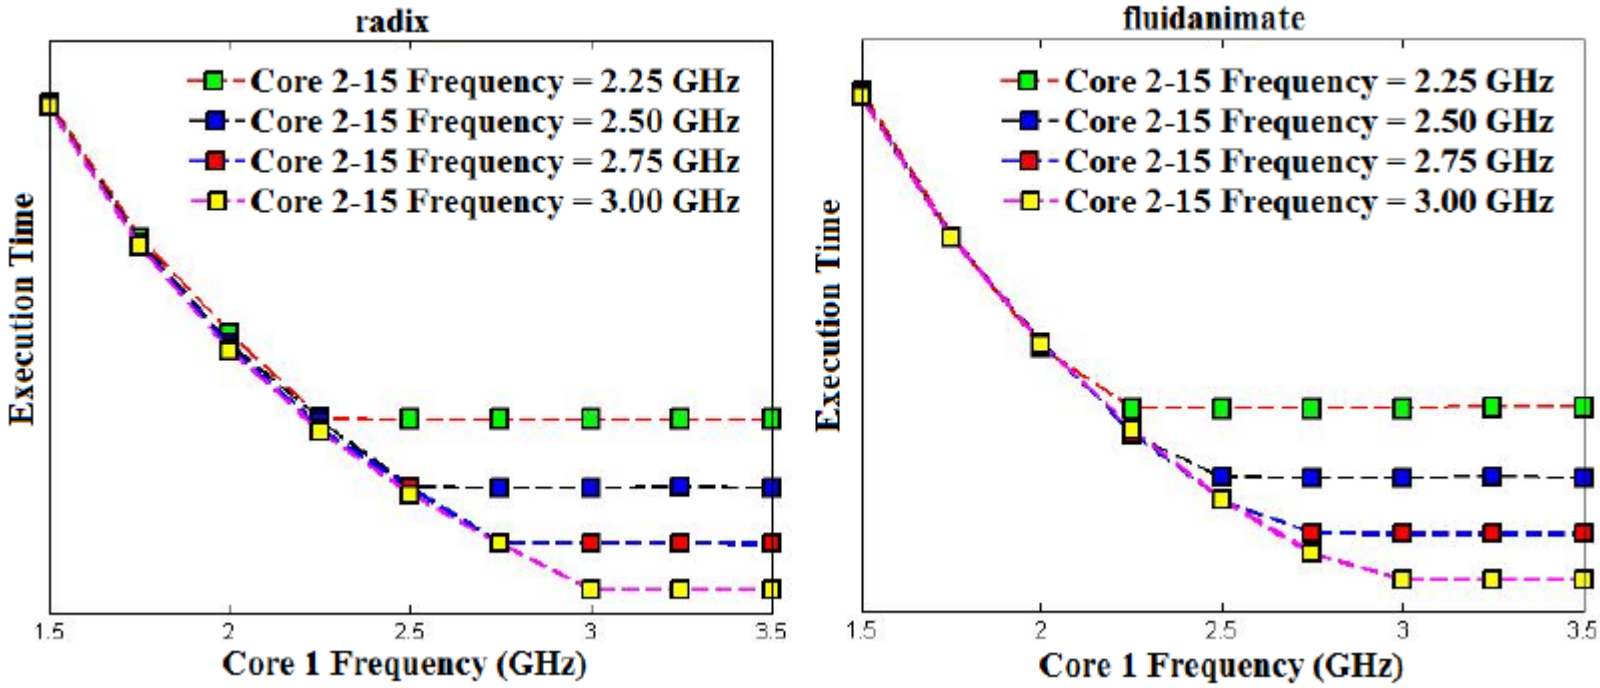

- When core 1 frequency is **lower than** frequency of other cores, **lower execution time** with increasing frequency

- When core 1 frequency is **higher than** frequency of other cores, **fixed execution time** with increasing frequency

Optimal Mapping

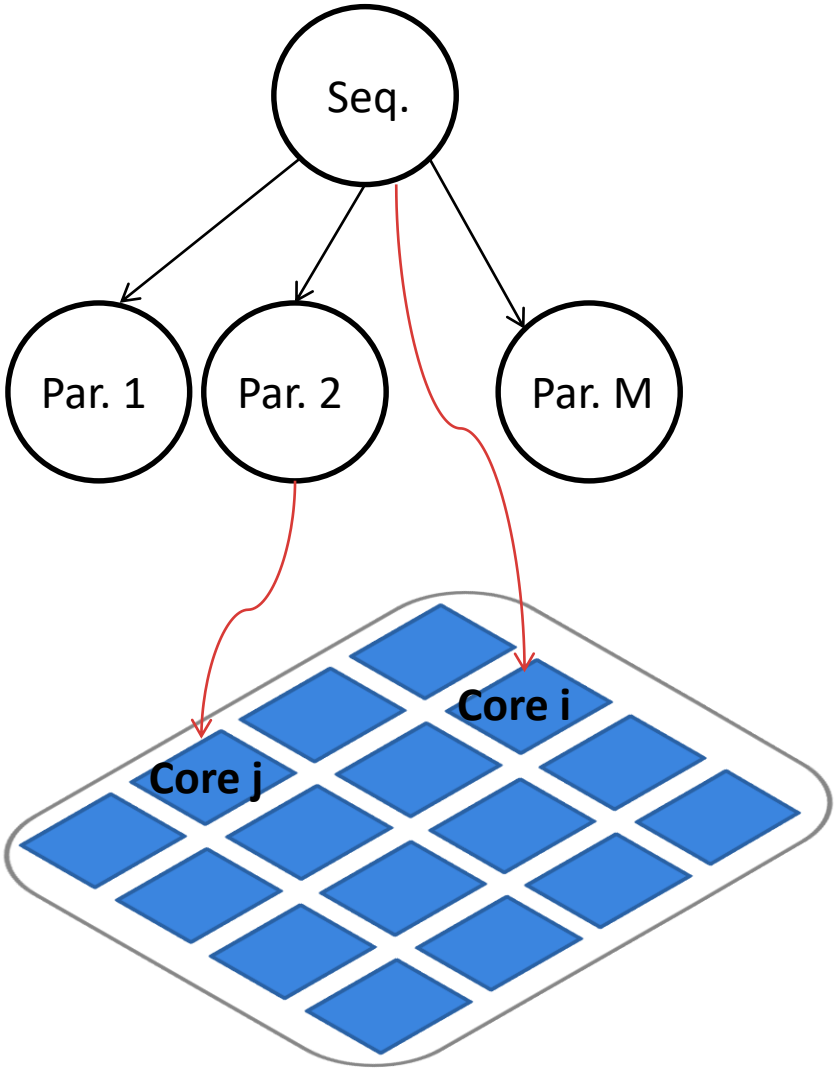

- **Assume that:**
  - Seq. thread executes on core i
  - Slowest parallel thread executes on core j
  - Q is a set of *M*-1 other cores:

$$f_k \geq f_j \quad \forall k \in Q$$

$$\sum_{k \in Q} P_k \leq P_{budget} - P_j$$

- **Execution time:**

$$E_{ij} = \frac{W_{seq}}{f_i} + \frac{W_{par}}{M f_j}$$

Frequency Scaling

- **For some <i,j> combinations, there might not exist *M*-1 faster cores that meet the power budget**
  - Frequency scaling can be used to meet power constraints at expense of performance
  - Frequency of all parallel cores scaled to the same frequency f_{par} such that:

$$MC_L V^2 f_{par} + \sum_{k \in Q \cup j} P_k(leak) \leq P_{budget}$$

  - Sufficient to only look at *M*-1 lowest leakage cores

Experimental Set-up

- **All experimental results based on the Sniper x86 multi-core simulator**
  - Interval core model, cycle-accurate cache, network and memory models
- **Parsec and SPLASH benchmarks with M=16**
  - Blackscholes
  - FFT
  - Radix
  - Fluidanimate
  - Swaptions

Core Parameters	Value
Nominal Frequency	3.0 GHz
Area	10.3 mm^2
Peak Dynamic Power	4.08W
Peak Leakage Power	2.1W

L2 Cache Parameters (per core)	Value
Size	2 MB
Area	11.2 mm^2
Peak Dynamic Power	0.76W
Peak Leakage Power	1.56W

Performance Model Validation

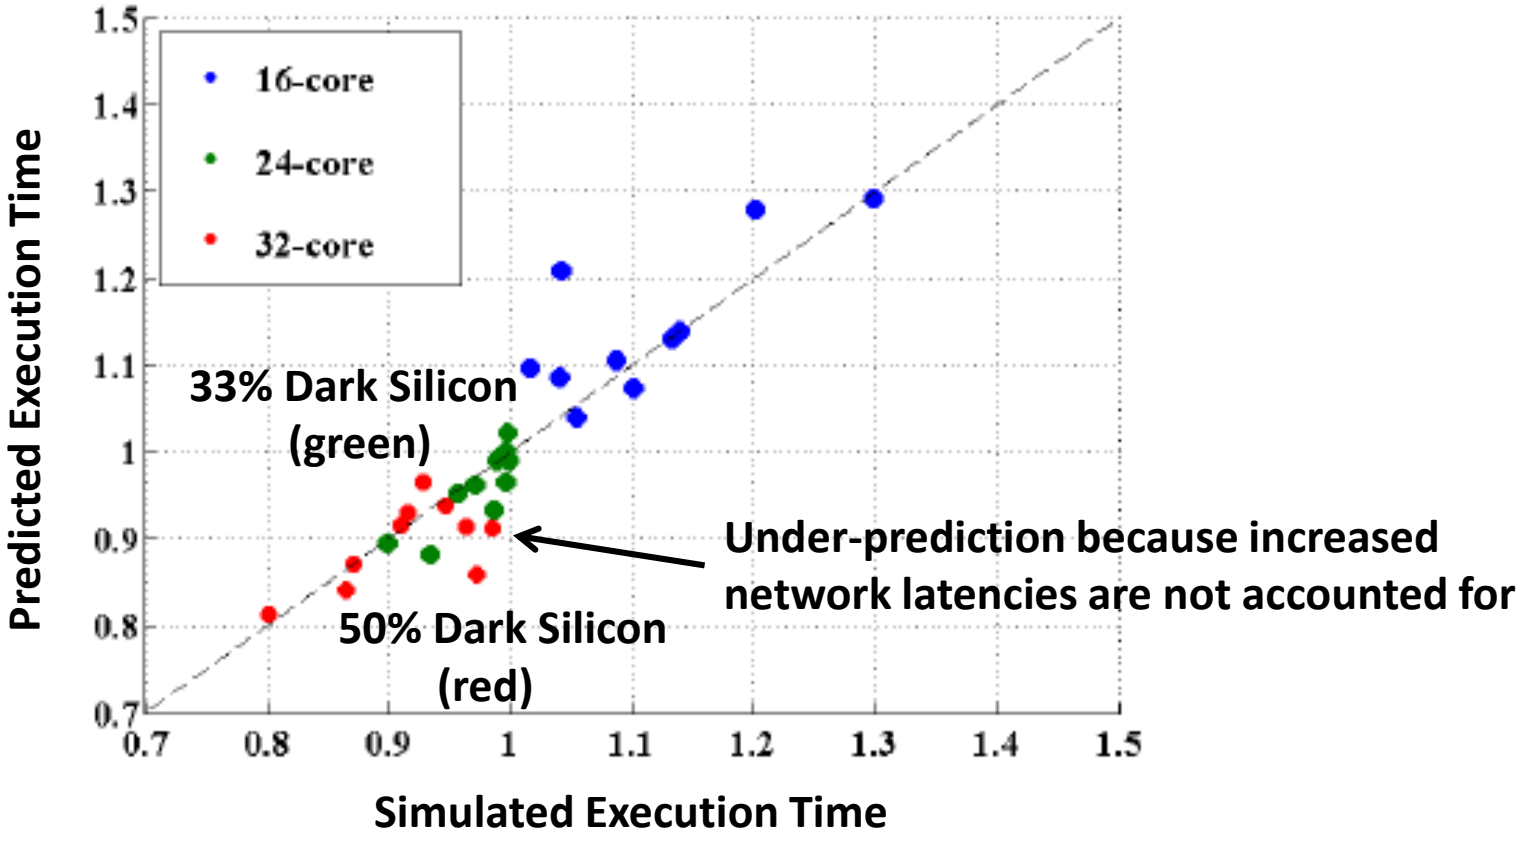

- **4.7% average error and 7.2% RMS error**

Performance Improvements

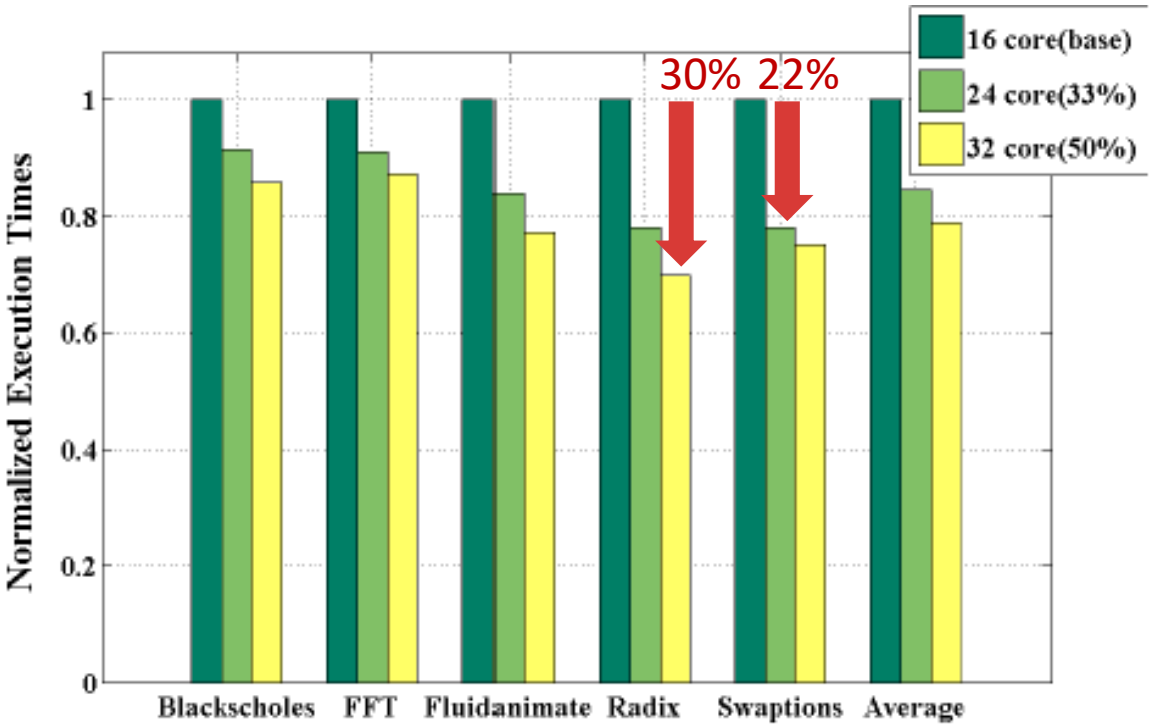

- **Averaged over 10 Monte Carlo experiments for each benchmark and each architecture**

Insight

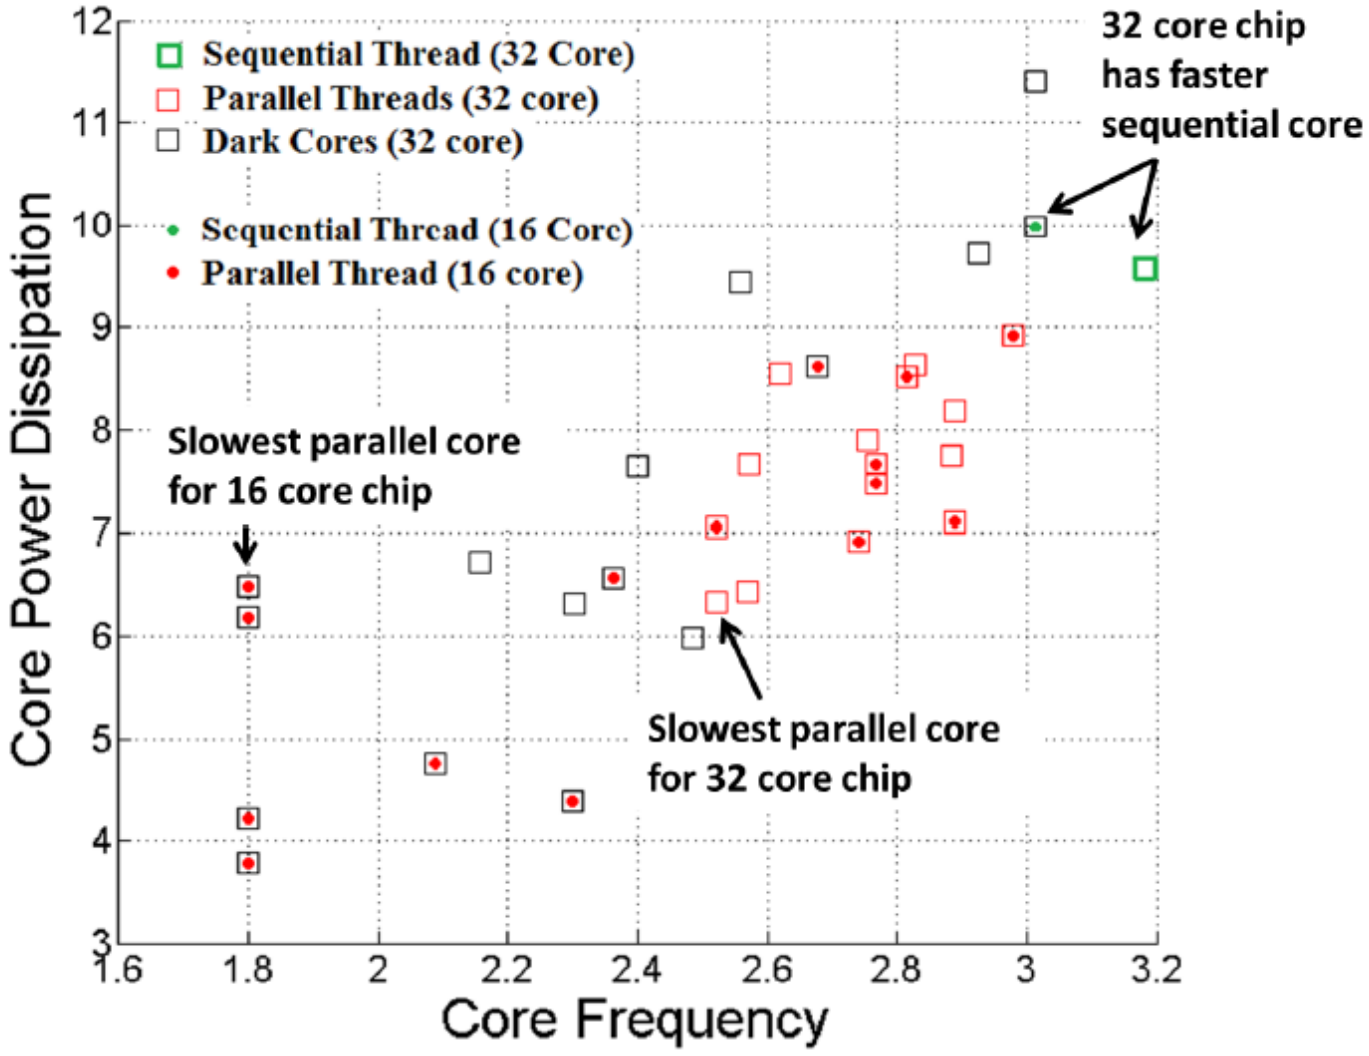
Sequential Thread (32 Core)
Parallel Threads (32 core)
Dark Cores (32 core)
Sequential Thread (16 Core)
Parallel Thread (16 core)
32 core chip has faster sequential core
Slowest parallel core for 16 core chip
Slowest parallel core for 32 core chip
Core Power Dissipation
Core Frequency
12
11
10
9
8
7
6
5
4
3
1.6
1.8
2
2.2
2.4
2.6
2.8
3
3.2

Discussion

- **Cherry picking proposes to pick the *best* subset of cores in a homogeneous dark silicon chip**
  - Power budget is met
  - Performance is maximized
  - Exploits process variations to create heterogeneity
- **Next generation dark silicon architectures might consist of a mix of architectural and process variation driven heterogeneity**
  - Replica accelerators

Upcoming

- **HaDeS: Architectural Synthesis for Heterogeneous Dark Silicon Chip Multi-Processors, DAC'13**
  - More sophisticated analytical performance models
  - Varying degrees of parallelism
  - Architectural heterogeneity