A Landscape of the New Dark Silicon Design Regime

Michael B. Taylor

UCSD Center for Dark Silicon

Associate Professor

University of California, San Diego

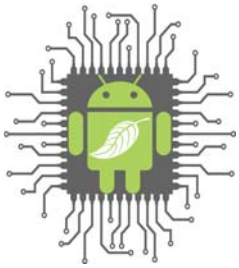

miketaylor.org

Compilers →
Kremlin
RawCC Parallelizing Compiler
Kismet

Architecture →
MIT Raw Tiled Multicore
Conservation Cores
GreenDroid Heterogeneous Proc
Scalar Operand Networks

Chips →

Tech Scaling →
Utilization Wall / Dark Silicon
Four Horseman

miketaylor.org

Compilers

Kremlin
RawCC Parallelizing Compiler
Kismet

Architecture

MIT Raw Tiled Multicore
Conservation Cores
GreenDroid Heterogeneous Proc
Scalar Operand Networks

Chips

Utilization Wall / Dark Silicon
Four Horseman

Tech Scaling

The Scaling Promise of Moore's Law

8 Years Ago

3.8 GHz

1 core

90 nm

The Scaling Promise of Moore's Law

8 Years Ago		Today, Xistors are:	
3.8 GHz	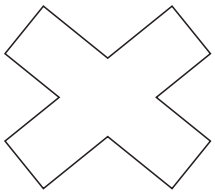	4x faster	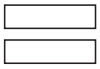
1 core		16x more plentiful	
90 nm		22 nm	

The Scaling Promise of Moore's Law

8 Years Ago		Today, Xistors are:		Today:
3.8 GHz	×	4x faster	=	**15.2** GHz
1 core		16x more plentiful		**16** Cores
90 nm		22 nm		

The Scaling Promise of Moore's Law

8 Years Ago		Today, Xistors are:	Today:
3.8 GHz	X	4x faster	= 15.2 GHz
1 core		16x more plentiful	16 Cores
			3.6 6
90 nm		22 nm	

The Scaling Promise of Moore's Law

8 Years Ago		Today, Xistors are:		Today:
3.8 GHz	X	4x faster	=	**3.6** GHz
1 core		16x more plentiful		**6** Cores
		↓		↓
		64X		5.7X

Dark Silicon: Transistors clocked far below their potential.

4 cores
1.8 GHz

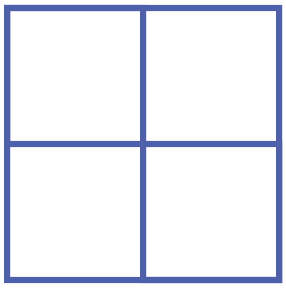

65 nm

8 cores
1.8 GHz

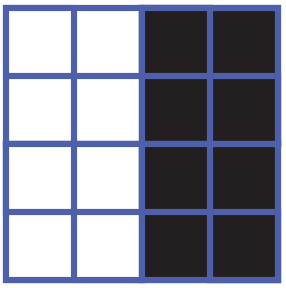

32 nm

2x cores per 2 generations, flat frequency

Dark or Dim Silicon (“uncore”)

This Talk

Explaining the Source of Dark Silicon: The Utilization Wall

The Four Horsemen of the Dark Silicon Apocalypse

Oct 2013 issue

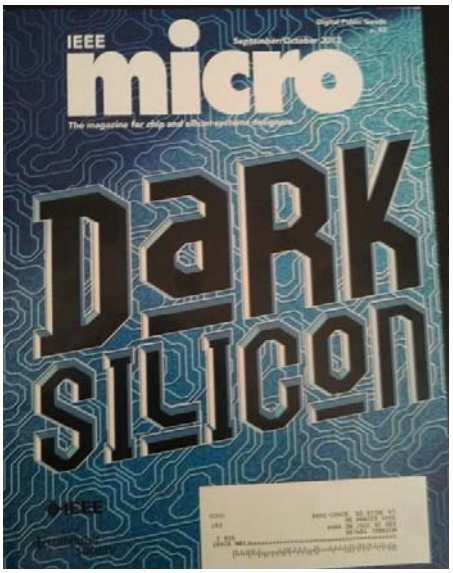

GreenDroid: An Architecture for the Dark Silicon Age

Where does dark silicon come from? And how dark is it going to be?

The Utilization Wall:

With each successive process generation, the percentage of a chip that can switch at full frequency drops **exponentially** due to power constraints.

[Venkatesh, Taylor, etc, ASPLOS ‘10]

Dennard:

"We can keep power consumption constant"

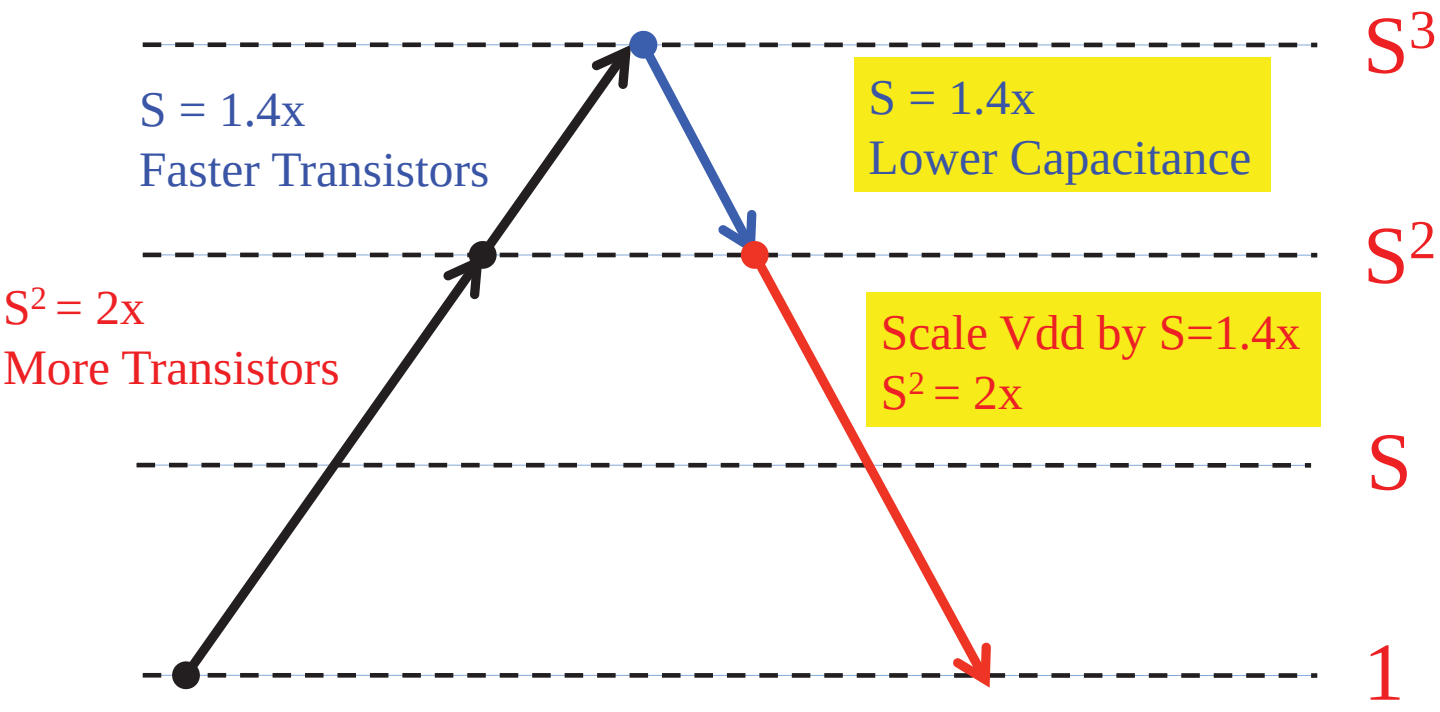

Fast forward to 2005:

Threshold Scaling Problems due to Leakage Prevents Us From Scaling Voltage

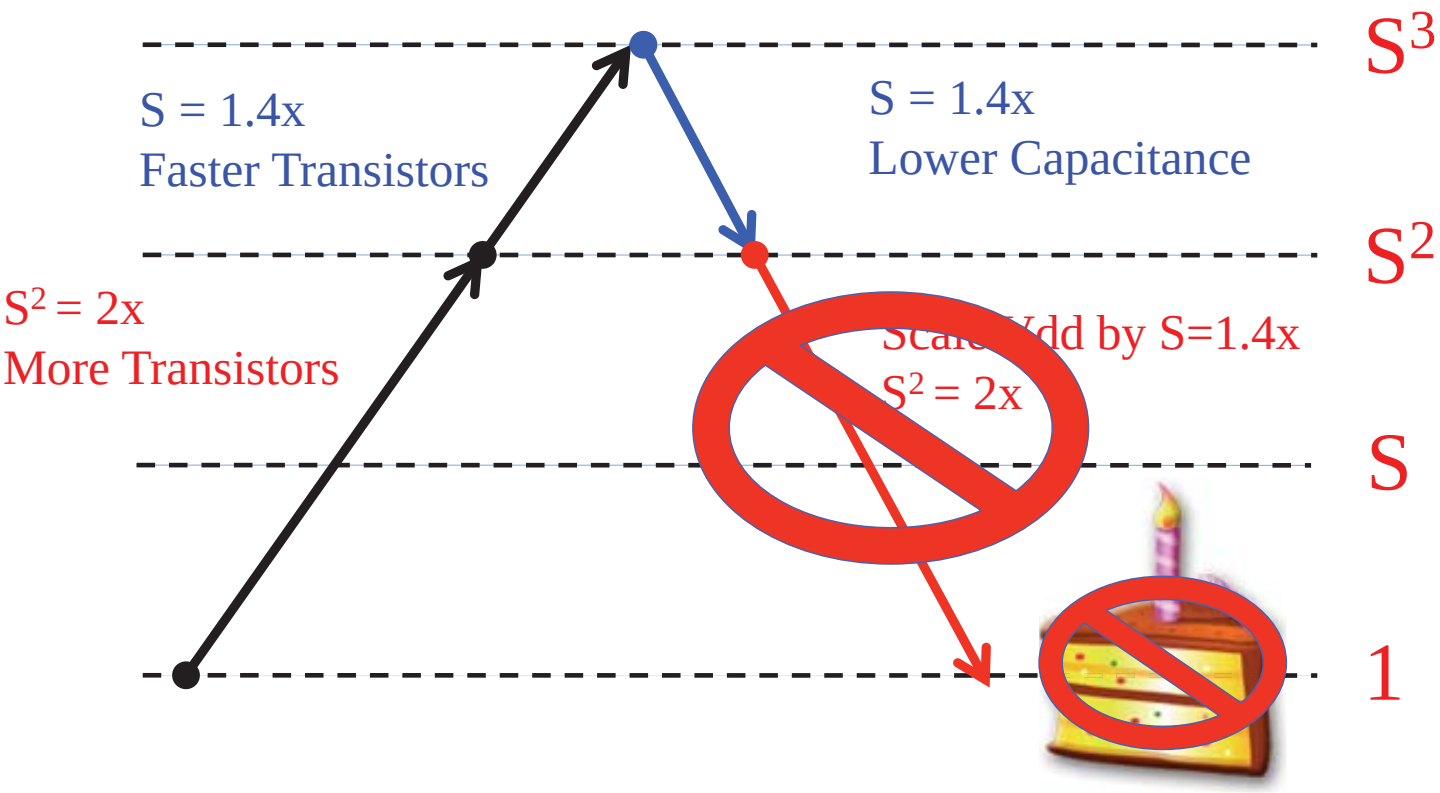

Full Chip, Full Frequency Power Dissipation
Is increasing exponentially by 2x with every process generation

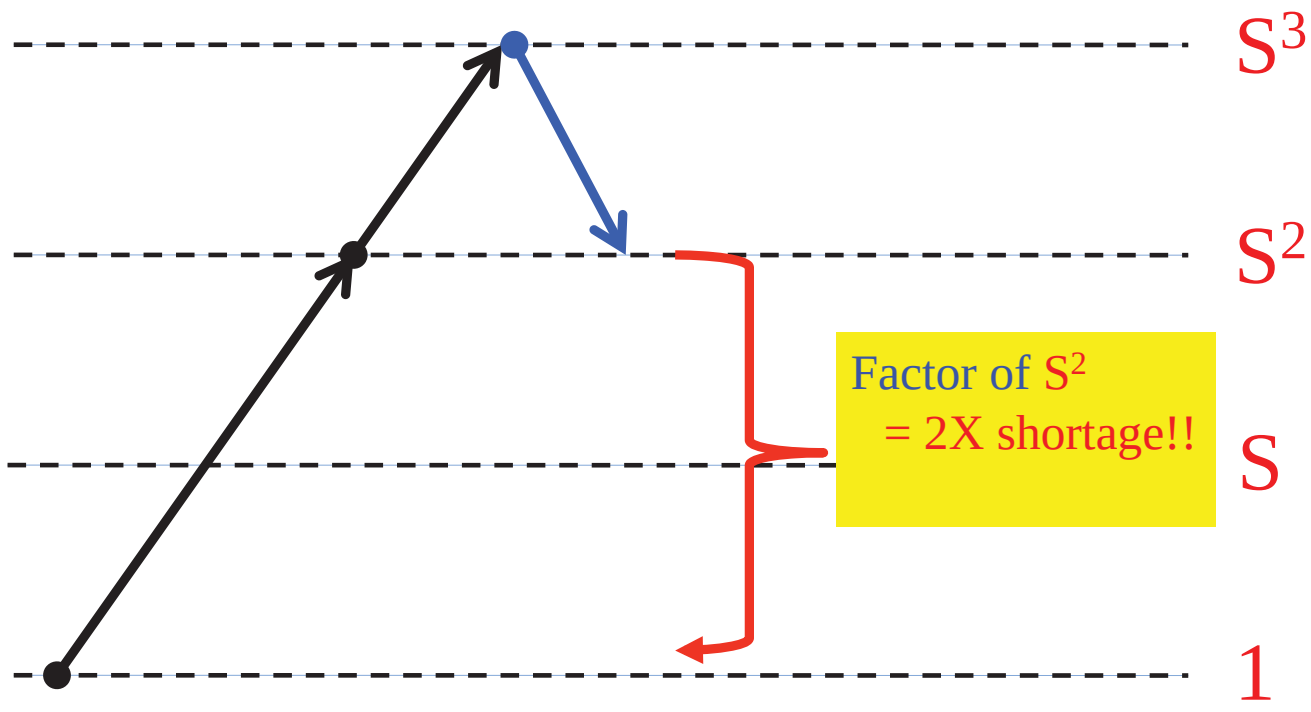

Multicore has hit the Utilization Wall

Spectrum of tradeoffs between # of cores and frequency

Example:
65 nm → 32 nm (S = 2)

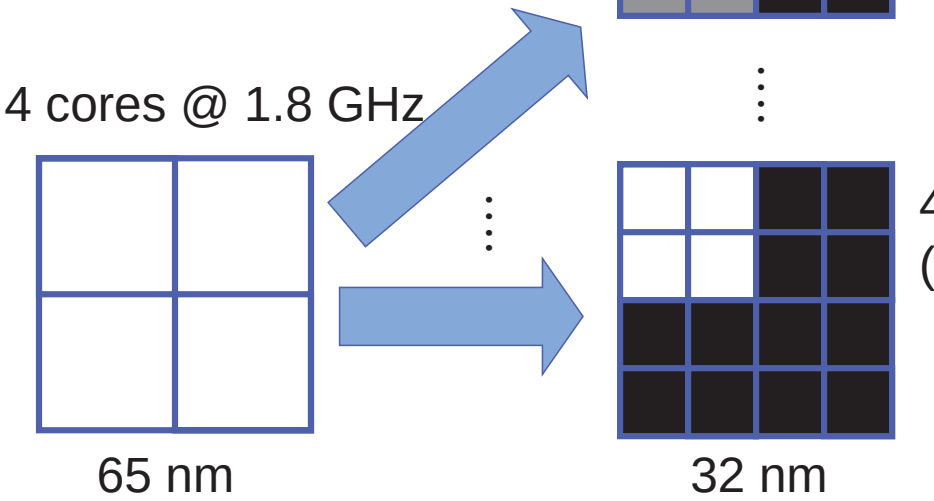

4x4 cores @ .9 GHz
(*GPUs of future?)*

2x4 cores @ 1.8 GHz
(8 cores dark, 8 dim)

(*Intel/x86 Choice, next slide*)

4 cores @ 2x1.8 GHz
(12 cores dark)

[Taylor, Hotchips 2010]
[Esmaeilzadeh ISCA 2011]

If multicore is not a solution to the dark silicon problem, what other potential avenues do we have?

→ Develop a more nuanced understanding of dark silicon.

This Talk

Explaining the Source of Dark Silicon

The Four Horsemen of the Dark Silicon Apocalypse

GreenDroid: An Architecture for the Dark Silicon Age

[Taylor, DAC 2012]

The Four Horsemen Taxonomy

Approaches for thriving in a future of dark silicon

A taxonomy of approaches for dealing with dark silicon. None is ideal, but each has its benefit and the optimal chip design probably incorporates all four…

I II III IV

The Shrinking Horseman (#1)

"Area is expensive. Chip designers will just build smaller chips instead of having dark silicon in their designs!"

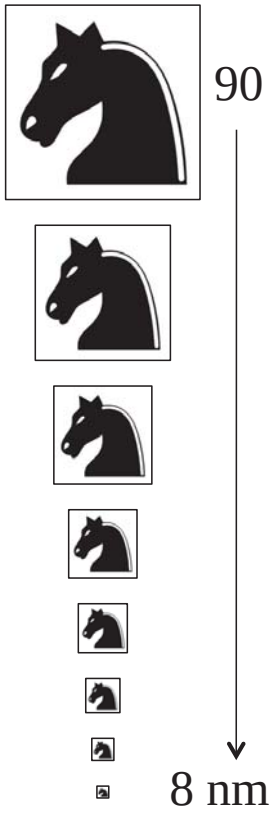

(if you work on Dark Silicon research, you will hear this a lot…)

The Shrinking Horseman (#1)

"Area is expensive. Chip designers will just build smaller chips instead of having dark silicon in their designs!"

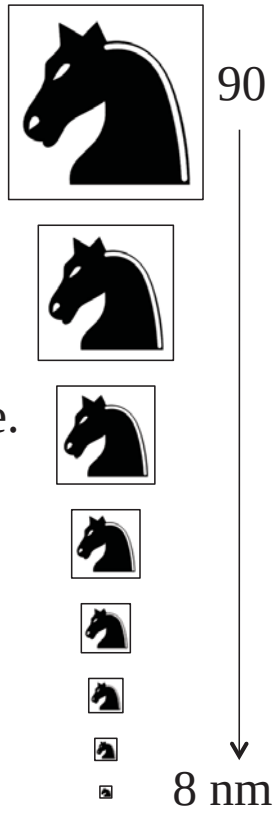

First, dark silicon doesn't mean *useless silicon*,
it just means it's under-clocked or not used all of the time.

There's lots of dark silicon in current chips:

- e.g. L3 cache cells are very dark
- un-core, unused SSE, etc.

The Shrinking Horseman (#1)

"Why not just build smaller chips!"

Possibly – but why didn't we shrink all of our chips before the dark silicon days? This too would be cheaper!

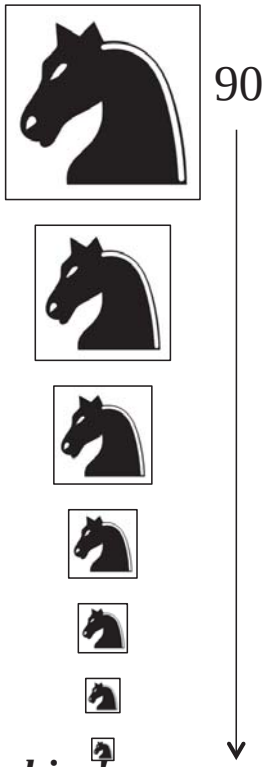

- **Exponential increase in Power Density**
  Exponential Rise in Temperature [Skadron]
- **Competition and Margins**
  - *If dark silicon buys you anything, you have to use it too, to keep up with the Jones.*
- **Diminished Returns**
  - *Fixed % of chip cost is in the silicon; (versus marketing, test, packaging, I/O pad area, etc); savings decrease exponentially w/ scaling.*
- **But, some chips *will* shrink**
  - *Especially chips where improved perform. does not matter*
  - ***Race to the bottom; except if we can sell exponentially more chips!***

The Four Horsemen Taxonomy

Explaining the Source of Dark Silicon

The Four Horsemen of the Dark Silicon Apocalypse

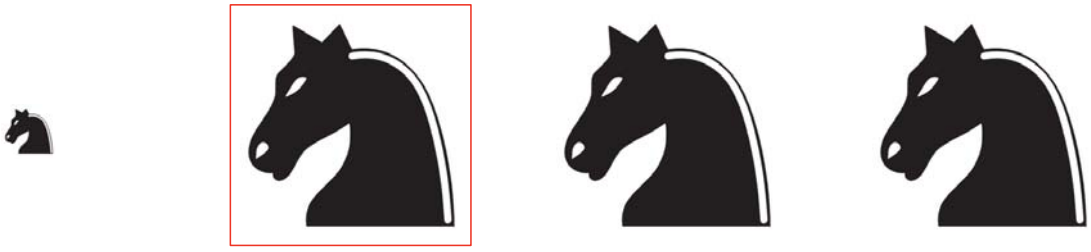

I II III IV

The Dim Horseman (#2)

"We will fill the chip with homogeneous cores that would exceed the power budget but we will:

- ***underclock them (spatial dimming)***
- ***use in bursts (temporal dimming)***

… "dim silicon".

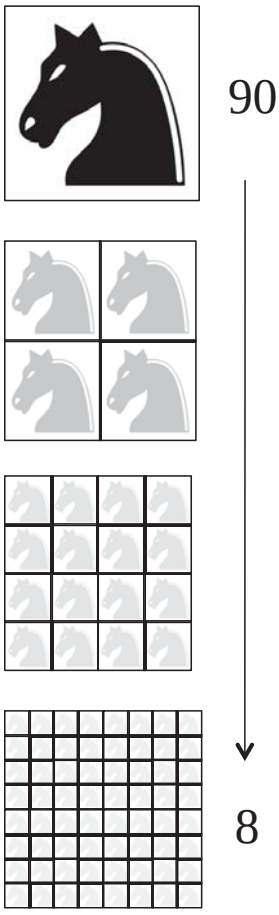

The Dim Horseman (#2)

Spatial Dimming

Multicores (higher core count → lower freqs)

"Upward DVFS" – Cubic Power Increase

Downward DVFS no longer decreases power cubically

Near Threshold Voltage (NTV) Operation

- Bad: Delay Increase >> Energy Gain
  e.g 5X more energy, 8X slower
  also, requires redesign of circuits for NTV operation!
- Good: *Energy per op is still lower*
- 8X more cores / parallelism → 1X perf, 5X lower energy
- 40X more cores / parallelism → 5X perf, same energy
- **But watch for Non-Ideal Speedups / Amdahl's Law**

NTV Examples:

- Manycore (e.g., UMich Centip3de [ISSCC 2012])
- SIMD (e.g., Synctium [CAL 2010]]
- x86 [Intel, ISSCC 2012]

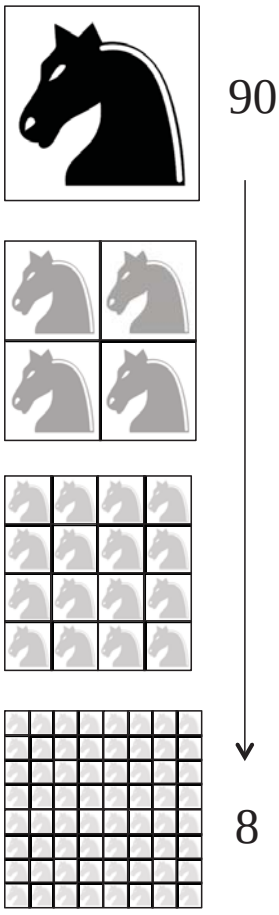

The Dim Horseman (#2)

Temporal Dimming: Computing in Bursts

- Battery Limited Systems
  - Active versus Standby mode
- Thermally Limited Systems
  - Turbo Boost 2.0 [Intel, Rotem et al., HOTCHIPS 2011]
    - Leverage Thermal Cap for DVFS – "overspend" if cold
  - Computational Sprinting, [Raghavan HPCA 2012]
  - ARM big.LITTLE in mobile phones or tablets [DAC 2012]
    - A15 power usage way above sustainable for phone → 10 second bursts at most

The Four Horsemen

Explaining the Source of Dark Silicon

The Four Horsemen of the Dark Silicon Apocalypse

I II III IV

The Specialized Horseman (#3)

"We will use all of that dark silicon area to build specialized cores, each of them tuned for the task at hand (10-100x more energy efficient), and only turn on the ones we need…"

[e.g., Venkatesh et al., ASPLOS 2010,
Lyons et al., CAL 2010,
Goulding et al., Hotchips 2010,
Hardavellas et al. IEEE Micro 2011]

90

8

This Talk

Explaining the Source of Dark Silicon

The Four Horsemen of the Dark Silicon Apocalypse

GreenDroid: An Architecture for the Dark Silicon Age

Conservation Cores

- Idea: Leverage dark silicon to “fight” the utilization wall
- Insights:
  - Specialized logic can improve energy efficiency by 10-1000x versus a general-purpose processor
  - Power is now more expensive than area

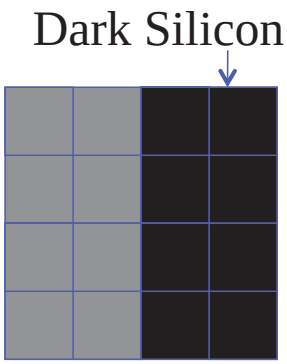

- Our Approach:
  - Fill dark silicon with *Conservation Cores*, or c-cores, which are specialized energy-saving coprocessors that save energy on common apps
  - Execution jumps from c-core to c-core
  - Power-gate c-cores that are not currently in use
- ***Conservation Cores provide an architectural way to trade dark area for an effective increase in power budget!***

Conservation Cores (C-cores)

"Conservation Cores: Reducing the Energy of Mature Computations," Venkatesh et al., ASPLOS '10

- Specialized coprocessors for reducing energy in irregular code
  - Hot code implemented by c-cores, cold code runs on host CPU;
  - C-cores use up to 18x less energy
  - Shared D-cache → Coherent Memory
  - ***Patching support in hardware***
- Fully-automated toolchain
  - ***No "deep" analysis or transformations required***
  - C-cores automatically generated from hot program regions
  - Design-time scalable
    - Emphasize Quantity over Quality!
    - *Simple conversion into HW buys us big gains, no need for heroic compiler efforts*.

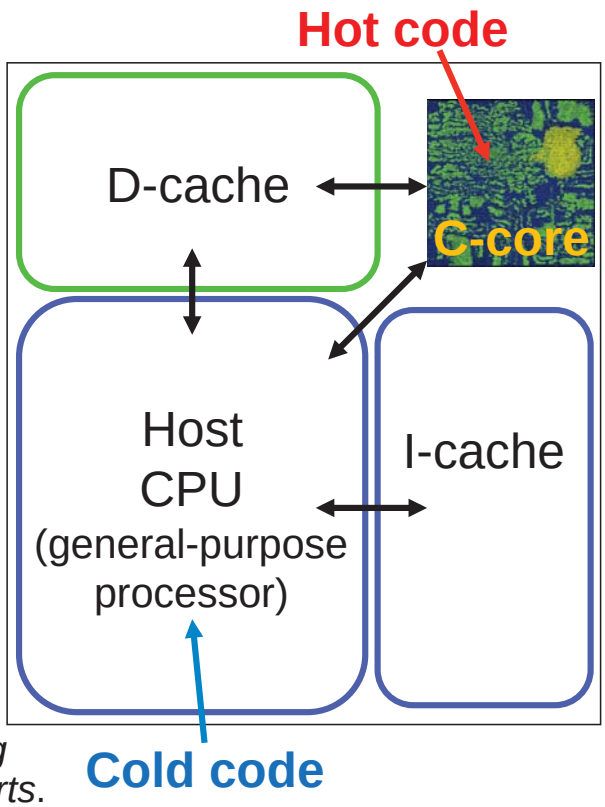

```
for (i=0; i<N; i++) {
  x = A[i];
  y = B[i];
  C[x] = D[y] + x+y+x*y;
}
```

C-core Generation

Start with ordinary C code. Irregular or regular is fine. Arbitrary control flow, arbitrary memory access patterns and complex data structures are supported.

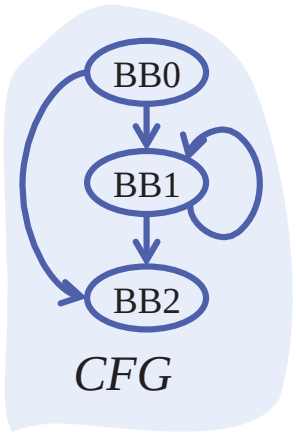

C-core Generation

Build a CFG; run ordinary compiler optimizations.

C-core Generation

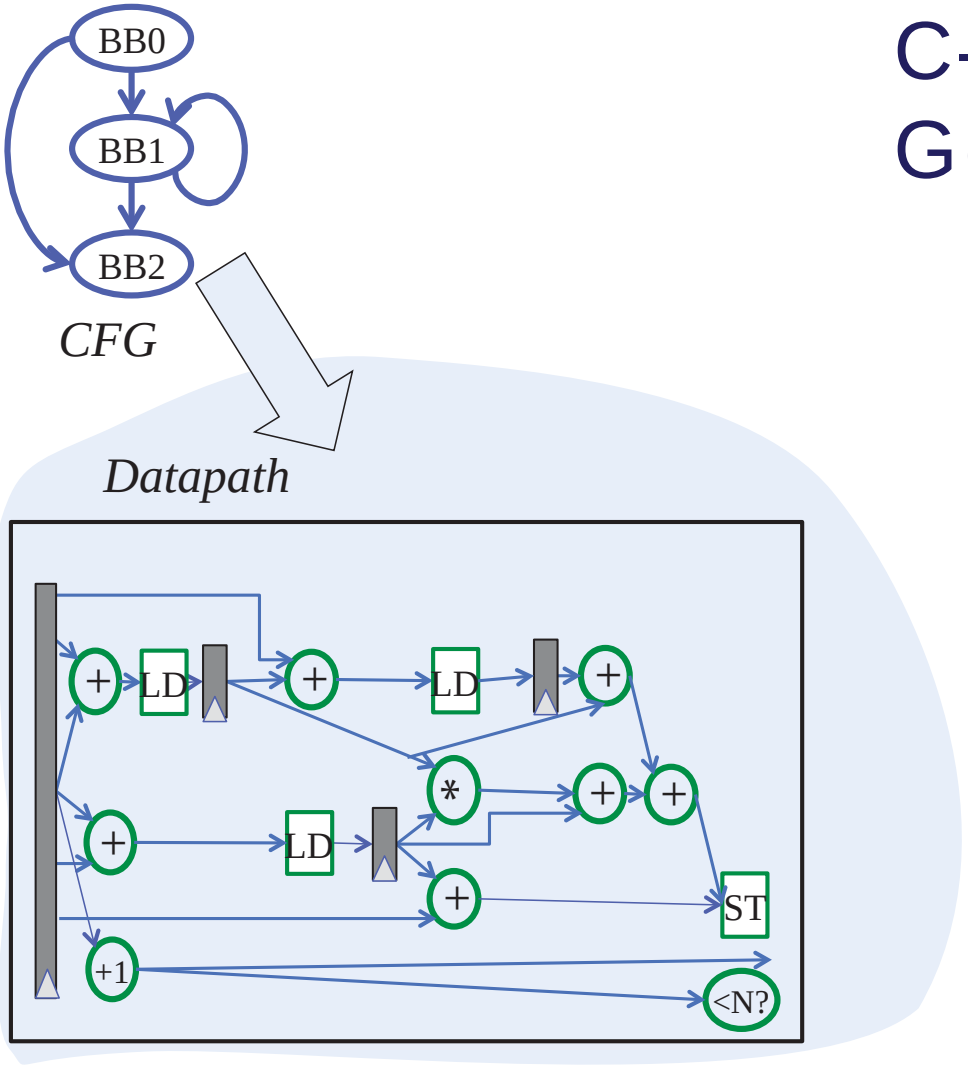

Each BB becomes a datapath; each operator turned into HW equivalent.

Memory ops mux'd into L1 cache.

Multiplier and FPUs may or may not be shared.

C-core Generation

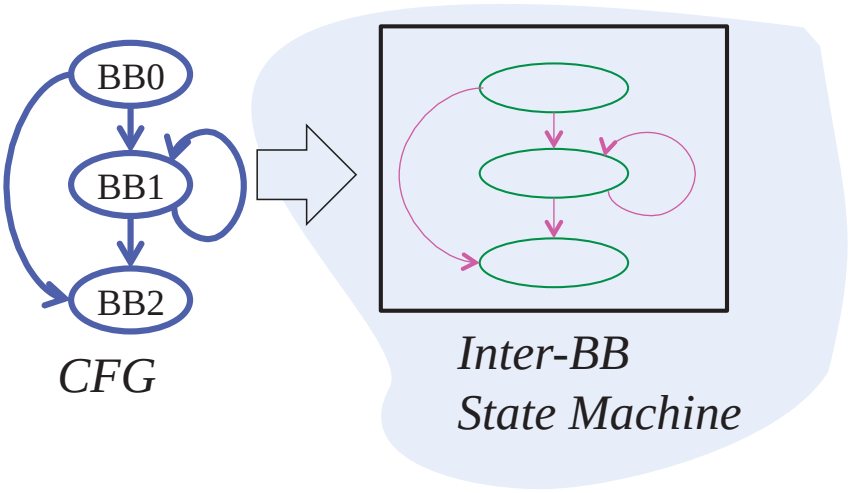

Datapath

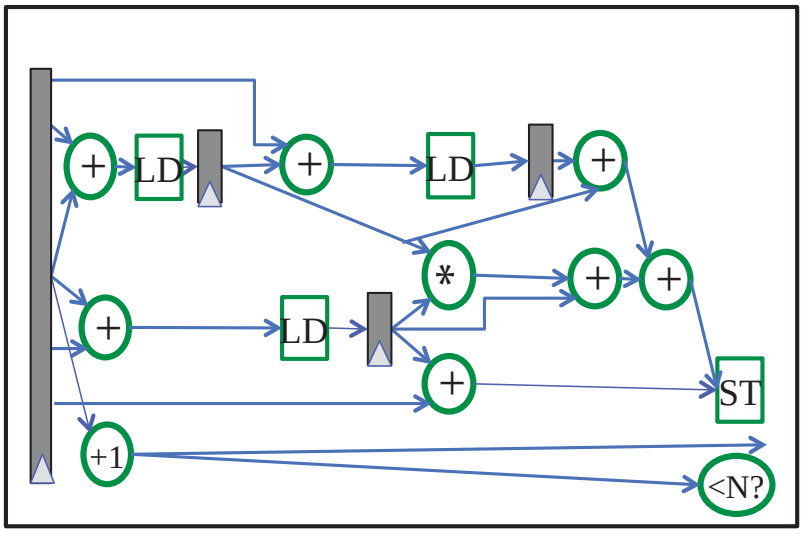

Create a state machine that determines which BB (datapath) is next.

C-core Generation

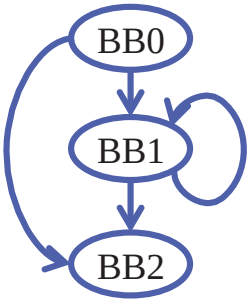

CFG

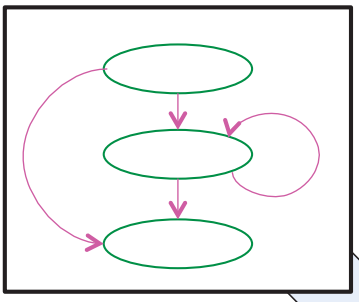

Inter-BB State Machine

.V

Via verilog, run through standard CAD flow.

Datapath

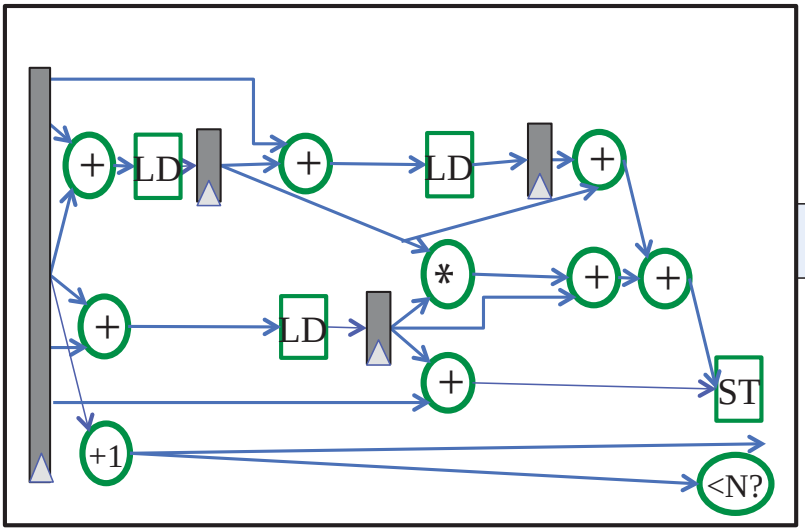

.V

Synopsys IC Compiler, P&R, CTS

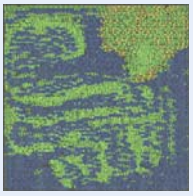

0.01 mm^2 in 45 nm TSMC runs at 1.4 GHz

Where do the energy savings come from?

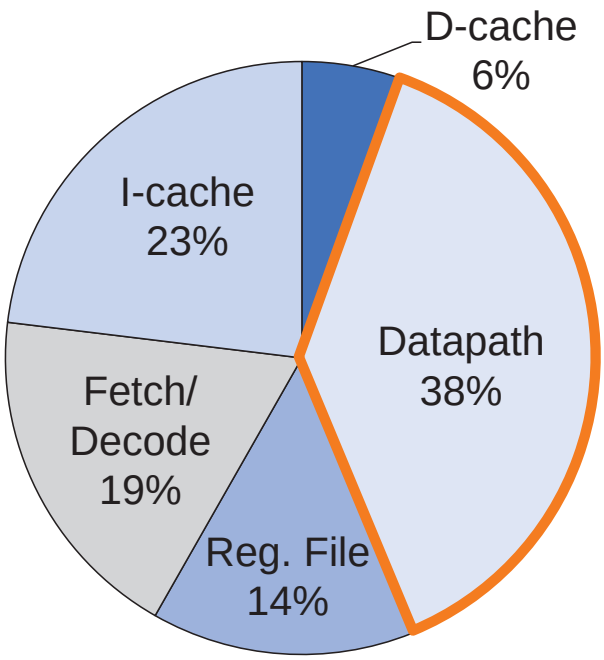

RISC baseline
91 pJ/instr.

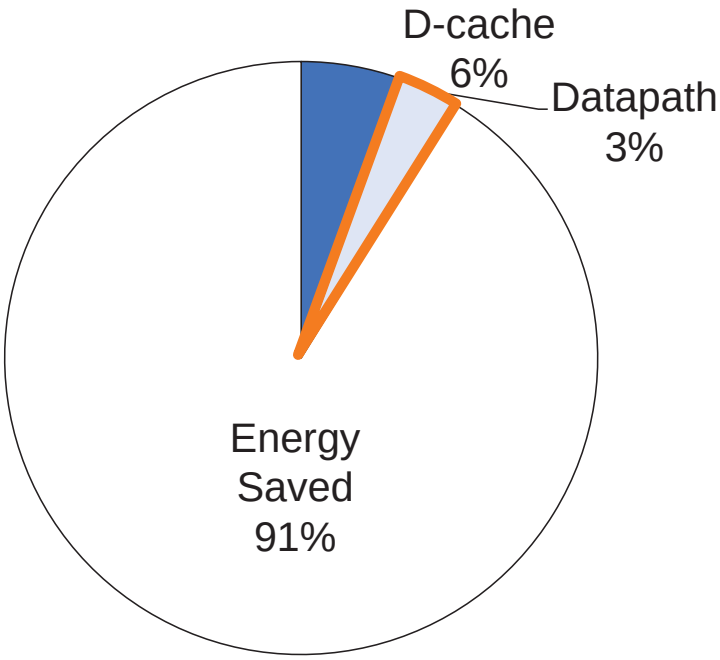

C-cores
8 pJ/instr.

Supporting Software Changes

- Software may change – HW must remain usable
  - C-cores unaffected by changes to cold regions
- Can support any changes, through *patching*
  - Arbitrary insertion of code – software exception mechanism
  - Changes to program constants – configurable registers
  - Changes to operators – configurable functional units
- Software exception mechanism
  - *State-tree* allows us to access any register in the C-core
  - Execute replacement code in processor
  - Write back values to c-core state-tree to resume execution

Android Time Highly Concentrated

61% of user time is spent in web + top 10 apps.
67% of user time is spent in web + top 20 apps.
73% of user time is spent in web + top 50 apps
growing at 2% for each additional 10 apps.

The average Android user spends almost an hour per day interacting with web and apps

Proportion of Time Spent on Web vs. Apps
Nielsen Smartphone Analytics, June 2011

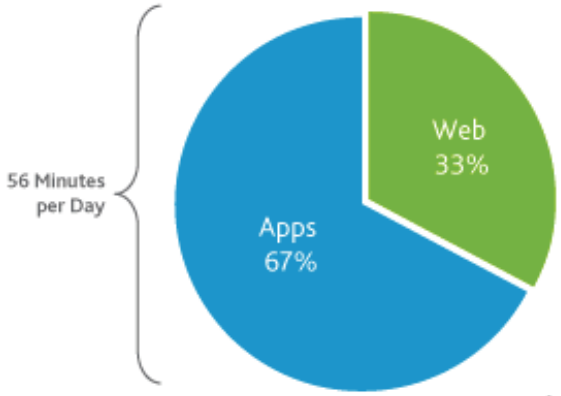

Source: Nielsen

nielsen

The top 50 apps by duration make up over 60% of all time spent on apps

Distribution of Time Spent in Apps
Nielsen Smartphone Analytics, June 2011

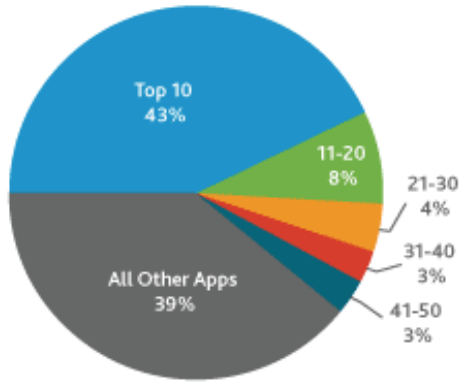

Source: Nielsen

nielsen

GreenDroid: Using c-cores to reduce energy in mobile application processors

"The GreenDroid Mobile Application Processor: An Architecture for Silicon's Dark Future,"
Goulding-Hotta et al., IEEE Micro Mar./Apr. 2011

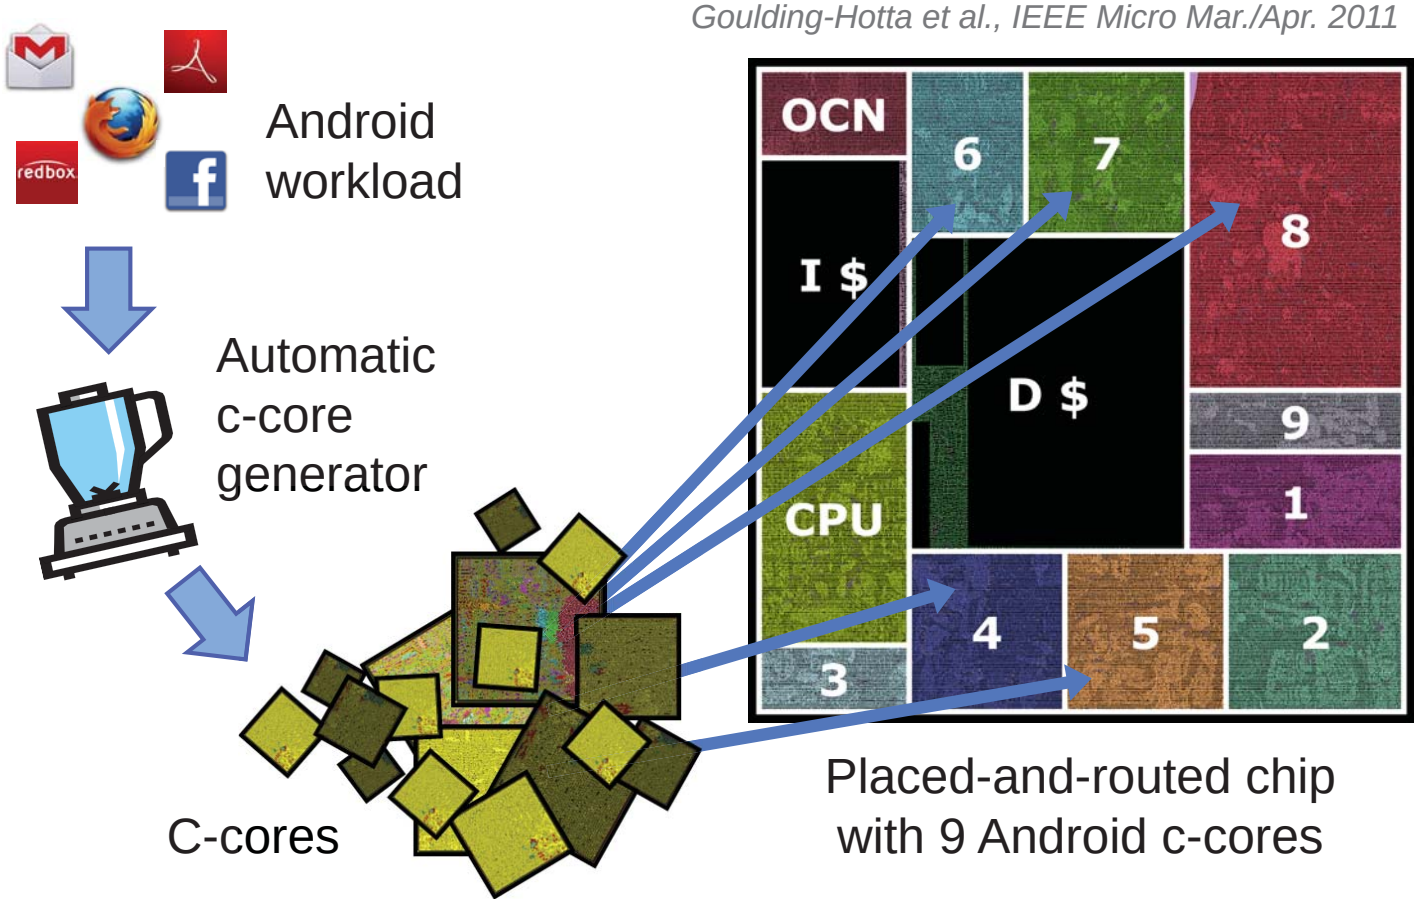

GreenDroid Tile Floorplan

- Norm to 45 nm:
  - 1.0 mm^2 per tile
  - 1.5 GHz
- 25% RISC core, I-cache, and on-chip network
- 25% D-cache
- 50% C-core “fill”

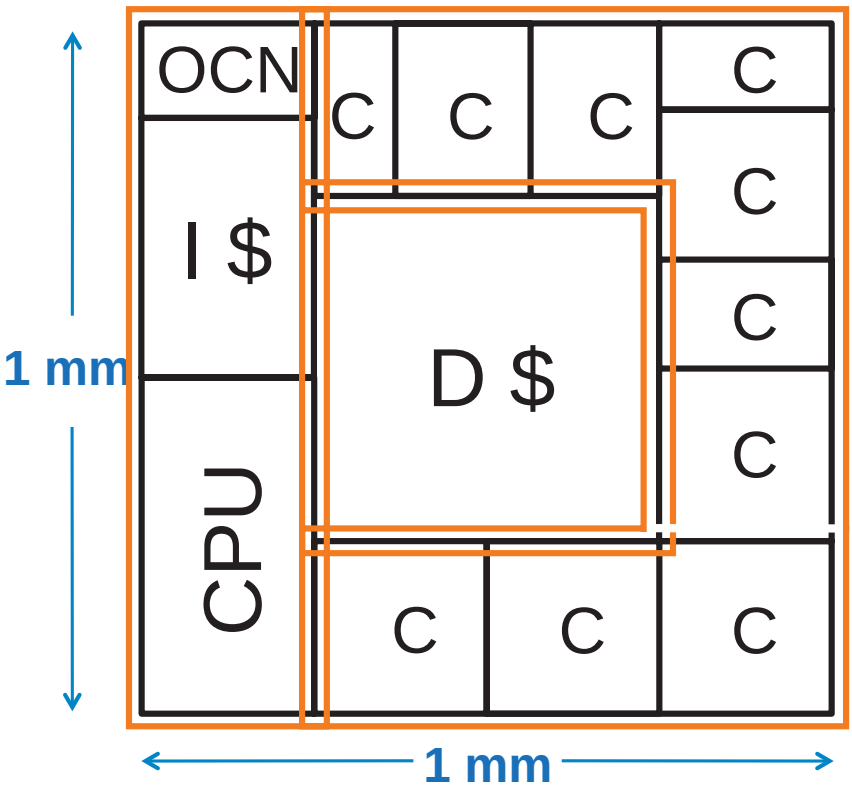

Quad-core GreenDroid Prototype

- Four heterogeneous tiles with ~40 C-cores.
- Synopsys IC Compiler
- 28-nm Global Foundries
- ~1.5 GHz
- 2 mm^2
- Multiproject Tapeout w/ UCSC

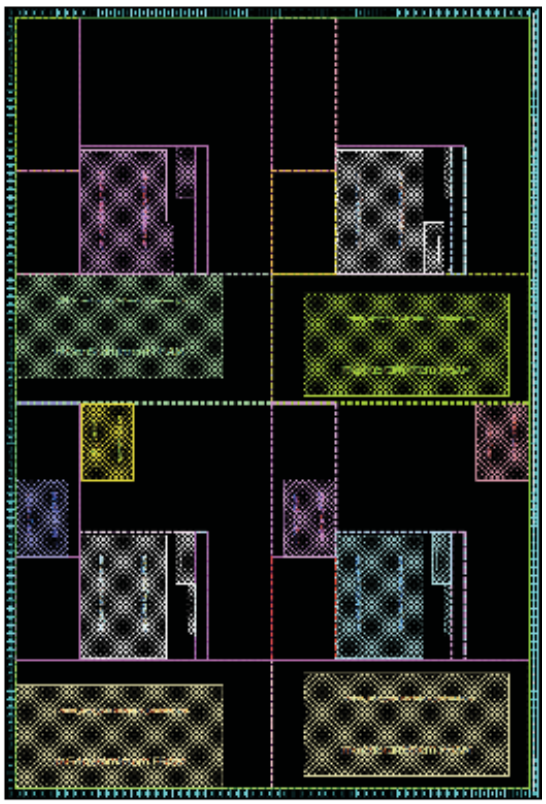

ASPLOS ‘10 → **HPCA ’11** → **MICRO ’11**

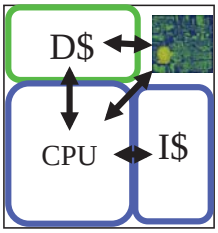

C-cores; Patching; Util. Wall

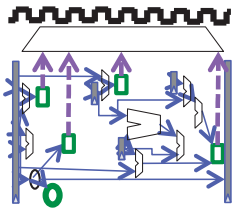

Selective Depipelining; Cachelets

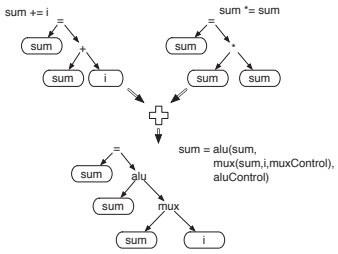

Automatic C-core Generalization

HOTCHIPS ‘10 → **IEEE MICRO ‘11** → **DAC ‘12**

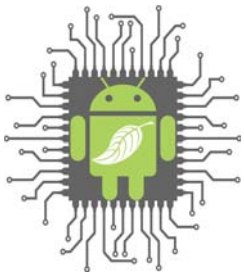

GreenDroid; Dark Silicon

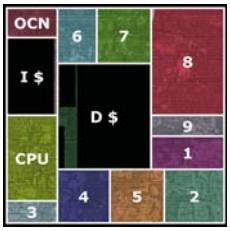

GreenDroid P&R Tile

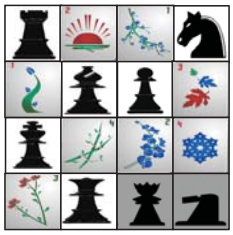

Four Horsemen

FCCM ’11 → **FPL ’11**

C-cores for FPGAs

Selective Depipelining for FPGAs

IEEE MICRO ‘13

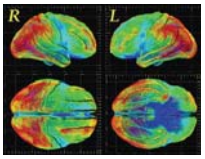

Landscape of Dark Silicon

The Four Horsemen

Explaining the Source of Dark Silicon

The Four Horsemen of the Dark Silicon Apocalypse

I

II

III

IV

The Deus Ex Machina Horseman

Latin [/*dayus ex makeena/*]
American [/duece ex mashina/]

deux ex machina */dayus ex makeena/*
A plot device whereby a seemingly unsolvable problem is suddenly and abruptly solved with the unexpected intervention of some new event, character, ability or object.

The Deus Ex Machina Horseman

"MOSFETs are the fundamental problem."

MOSFET variants
(FinFets, Trigate, High-K, nanotubes, 3D)

- **one-time improvements**
- **limited to 60 mV/decade subthreshold slope**
- **leakage is still there**

The Deus Ex Machina Horseman

Possible "Beyond CMOS" Device Directions (none are there yet, imho)

- **Nano-electrical Mechanical (NEMS) Relays**
  very low energy, physical connections, very slow [e.g, Spencer et al JSSC 2011]

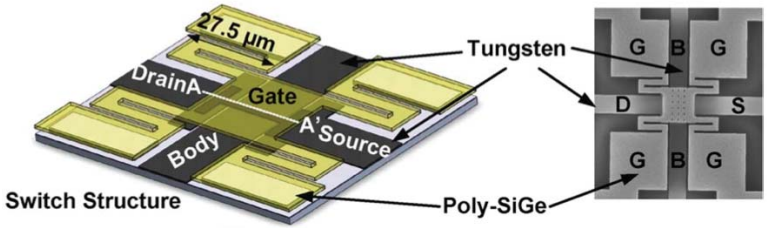

- **Tunnel Field Effect Transistors (TFETs)**
  use tunneling effects to overcome MOSFET limits better subthreshold slopes (~ 25 mV/decade) at lower voltages; not superior to MOSFETS at higher voltages

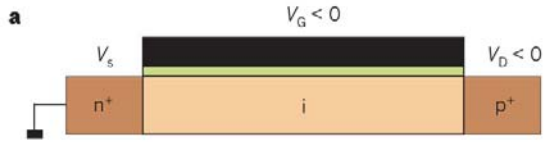

[e.g., Ionescu et al, Nature 2011]

DARPA / SRC / MARCO's $194M Starnet Investment

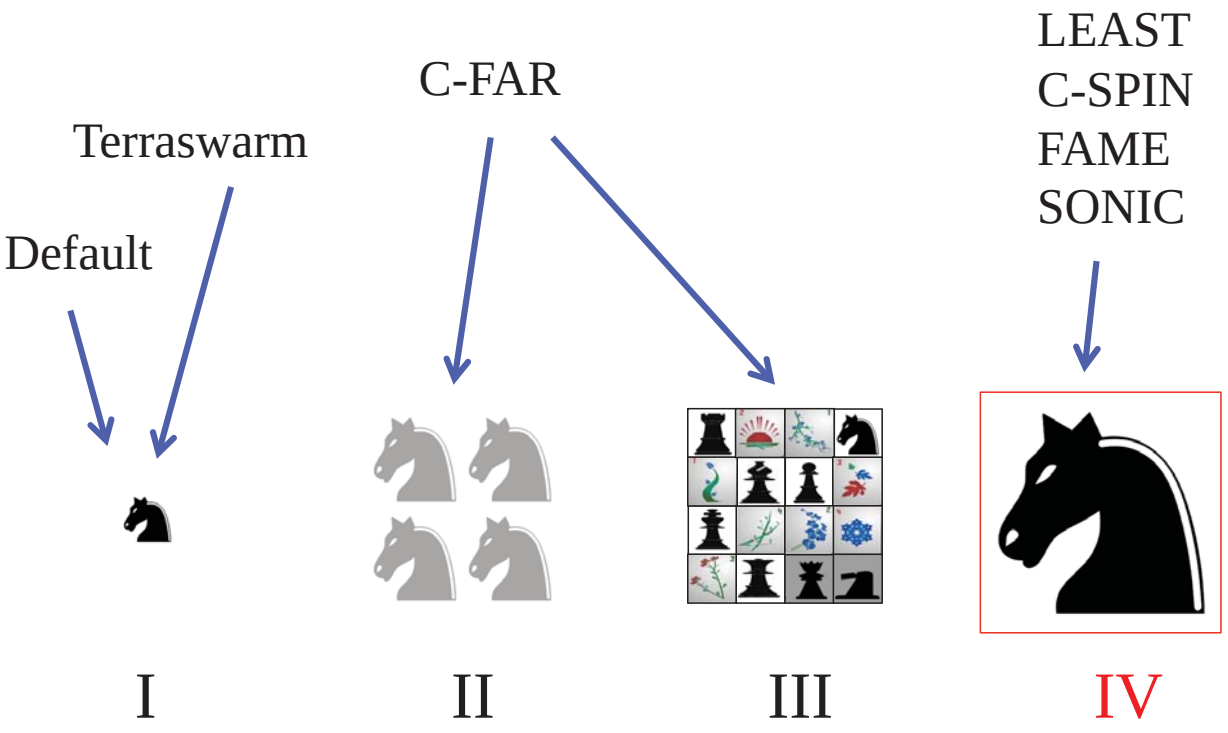

For more:

IEEE Micro 2013 Paper

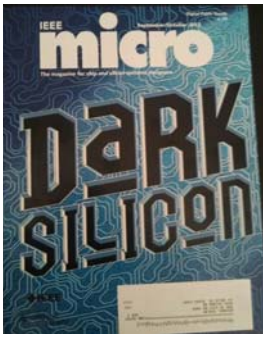

Explaining the Source of Dark Silicon

The Four Horsemen of the Dark Silicon Apocalypse

Dark Silicon Design Principles

A Truly Dark Computing Fabric: The Brain

Conclusion

- Dark Silicon is opening up a whole new class of exciting new architectural directions which many folks are starting to move into – which I have termed the “four horsemen”.
- GreenDroid is one interesting example of an architecture that explores one of these directions.

I II III IV

The UCSD GreenDroid Team

Prof. Taylor

darksilicon.org

Prof. Swanson