Power Constraints: From Sensors to Servers

Mike Muller
CTO

1983 The Rise of the PC and ARM Technology

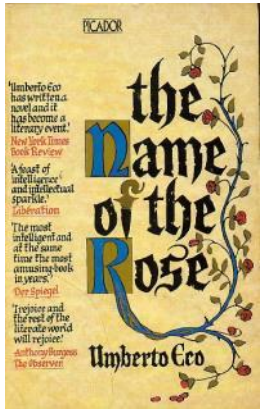

Vinyl

Compass

DynaTAC

Paper

Lotus123

Osborne

ARM®

1990

ARM Design Philosophy

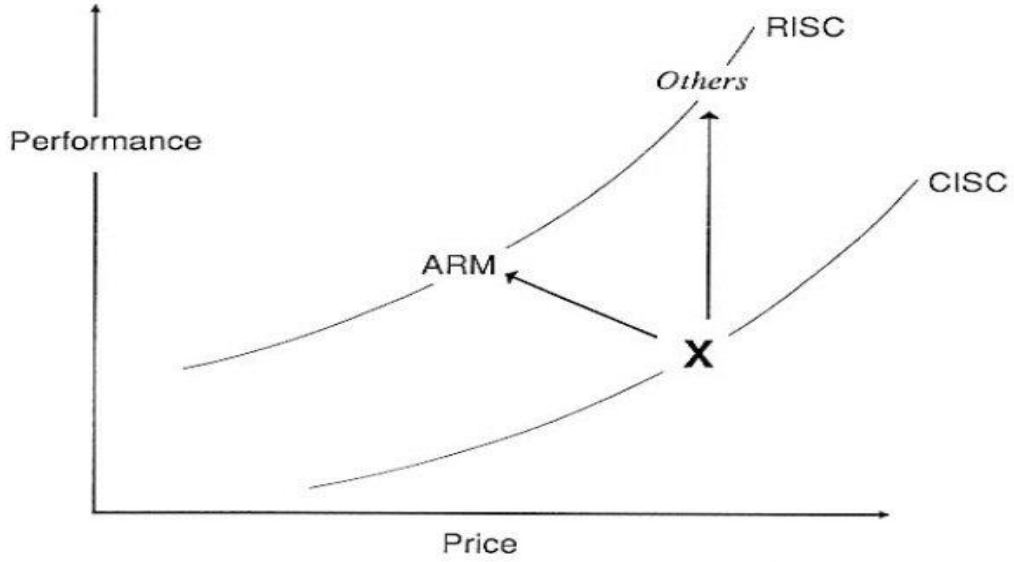

ARM Processors are:-
- Small in size
- Low cost
- Suitable as macrocells
- Efficient
  - low power consumption
  - low heat generation

ARM®

1992 HC04

1:45 – 3:15 **Session 3: Low Cost Processors**

Session Chair: John Mashey, Silicon Graphics

- **Highly Integrated SPARC Processor Implementation**
  S. Joshi, Sun Microsystems
- **The LR33020 GraphX Processor: A Single Chip MIPS-RISC Based X Terminal Controller**
  S. Desai, LSI Logic Corporation
- **The ARM600 Processor and FPA**
  M. Muller, Advanced RISC Machines Ltd.

3:15 – 3:45 **Break**

3:45 – 5:15 **Session 4: Low Power Systems**

Session Chair: Dave Ditzel, Sun Microsystems

- **A VLSI Chip Set for Personal Communications System**
  R. Scavuzzo, AT&T Bell Labs
- **SPARC90 - Chipset on a Chip**
  J. Pendleton, Sun Microsystems
- **Cold Chip Design Techniques**
  R. Broderson, A. Chandrakasan and S. Sheng, UC Berkeley

ARM®

1993 The Tipping Point

Mosaic v1.0

DVD

Game Boy

Zarlink Orion GPS

1st ARM GPS

1st ARM PDA

1st Consumer GSM

1st ThinkPad

ARM®

2003 Unplugged

2004

ARM®

2013 Mobile Computing

ARM®

Top Right

and Bottom Left

ARM®

ENABLING
INNOVATION

Moore

ARM7TDMI

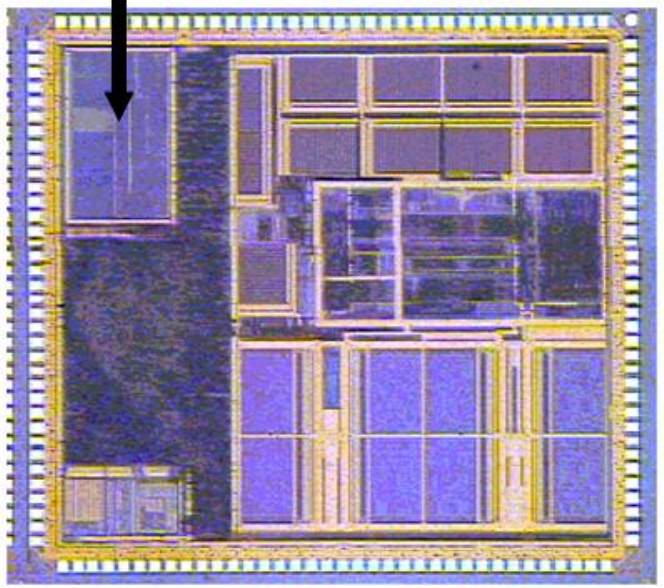

GSM Base-Band
0.5u.

Wireless Enabled Product
0.05u

2000

ARM POWERED

ARM®

Dark Silicon Source – ITRS 2008

Node	45nm	22nm	11nm
Year	2008	2014	2020
Area	1	x4	x16
Peak freq	1	1.6	2.4
Power			
@ 45nm freq	1	0.6	0.3
@ peak freq	1	1	0.6
Exploitable Si (equivalent power)		24%	10%

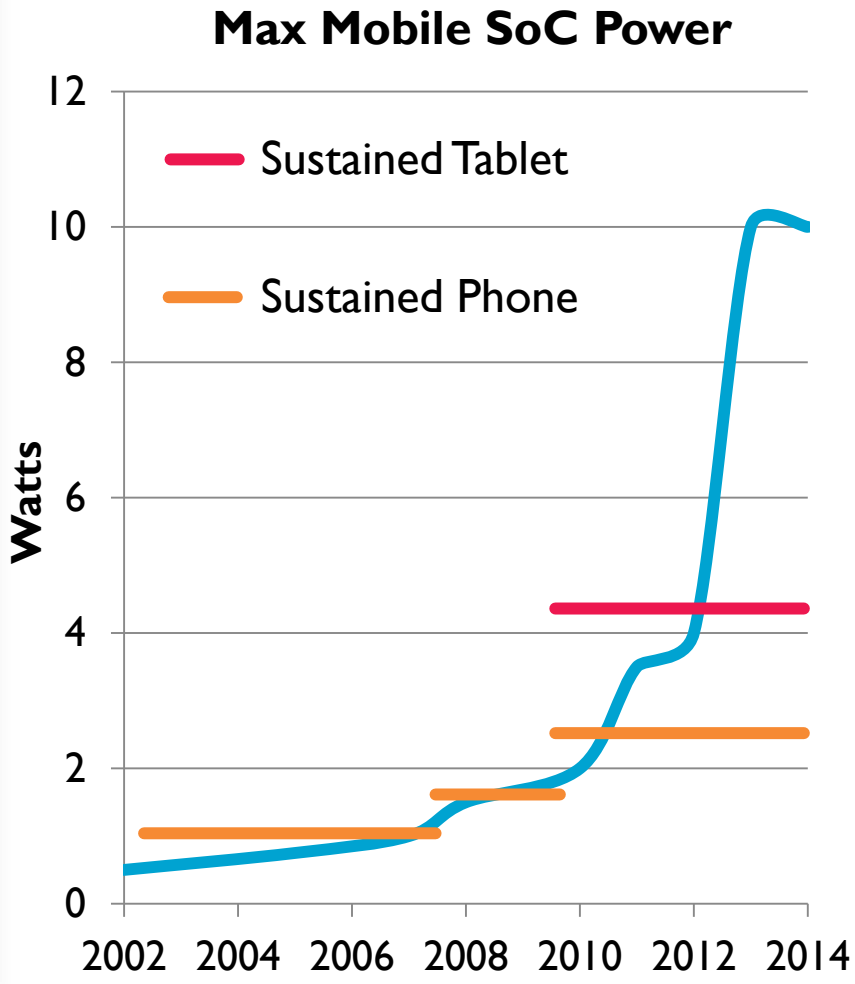

Source: ITRS 2008 & ARM

ARM®

Dark Silicon Source – ITRS 2008

Node	45nm	22nm	11nm
Year	2008	2014	2020
Area	1	x4	x16
Peak freq	1	1.6	2.4
Power			
@ 45nm freq	1	0.6	0.3
@ peak freq	1	1	0.6
Exploitable Si (equivalent power)		24%	10%

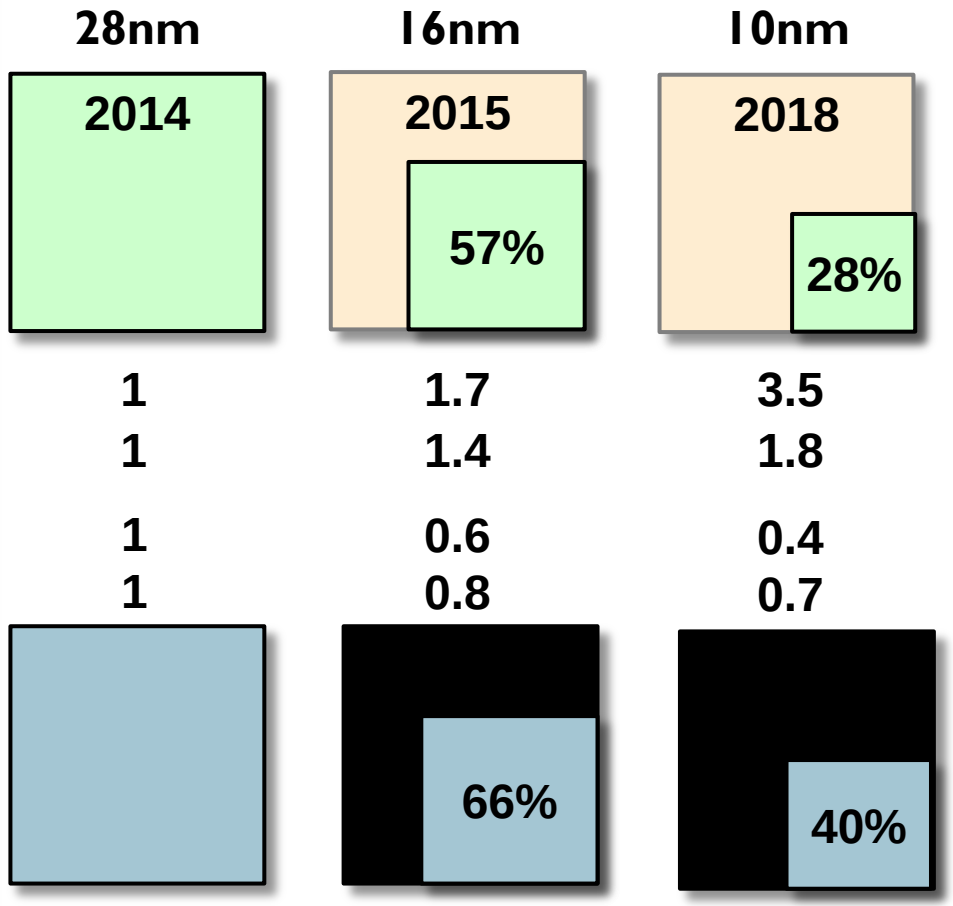

Source: ITRS 2008 & ARM

ARM®

It's Still All About the Power

ARM®

This Years Favorite Product

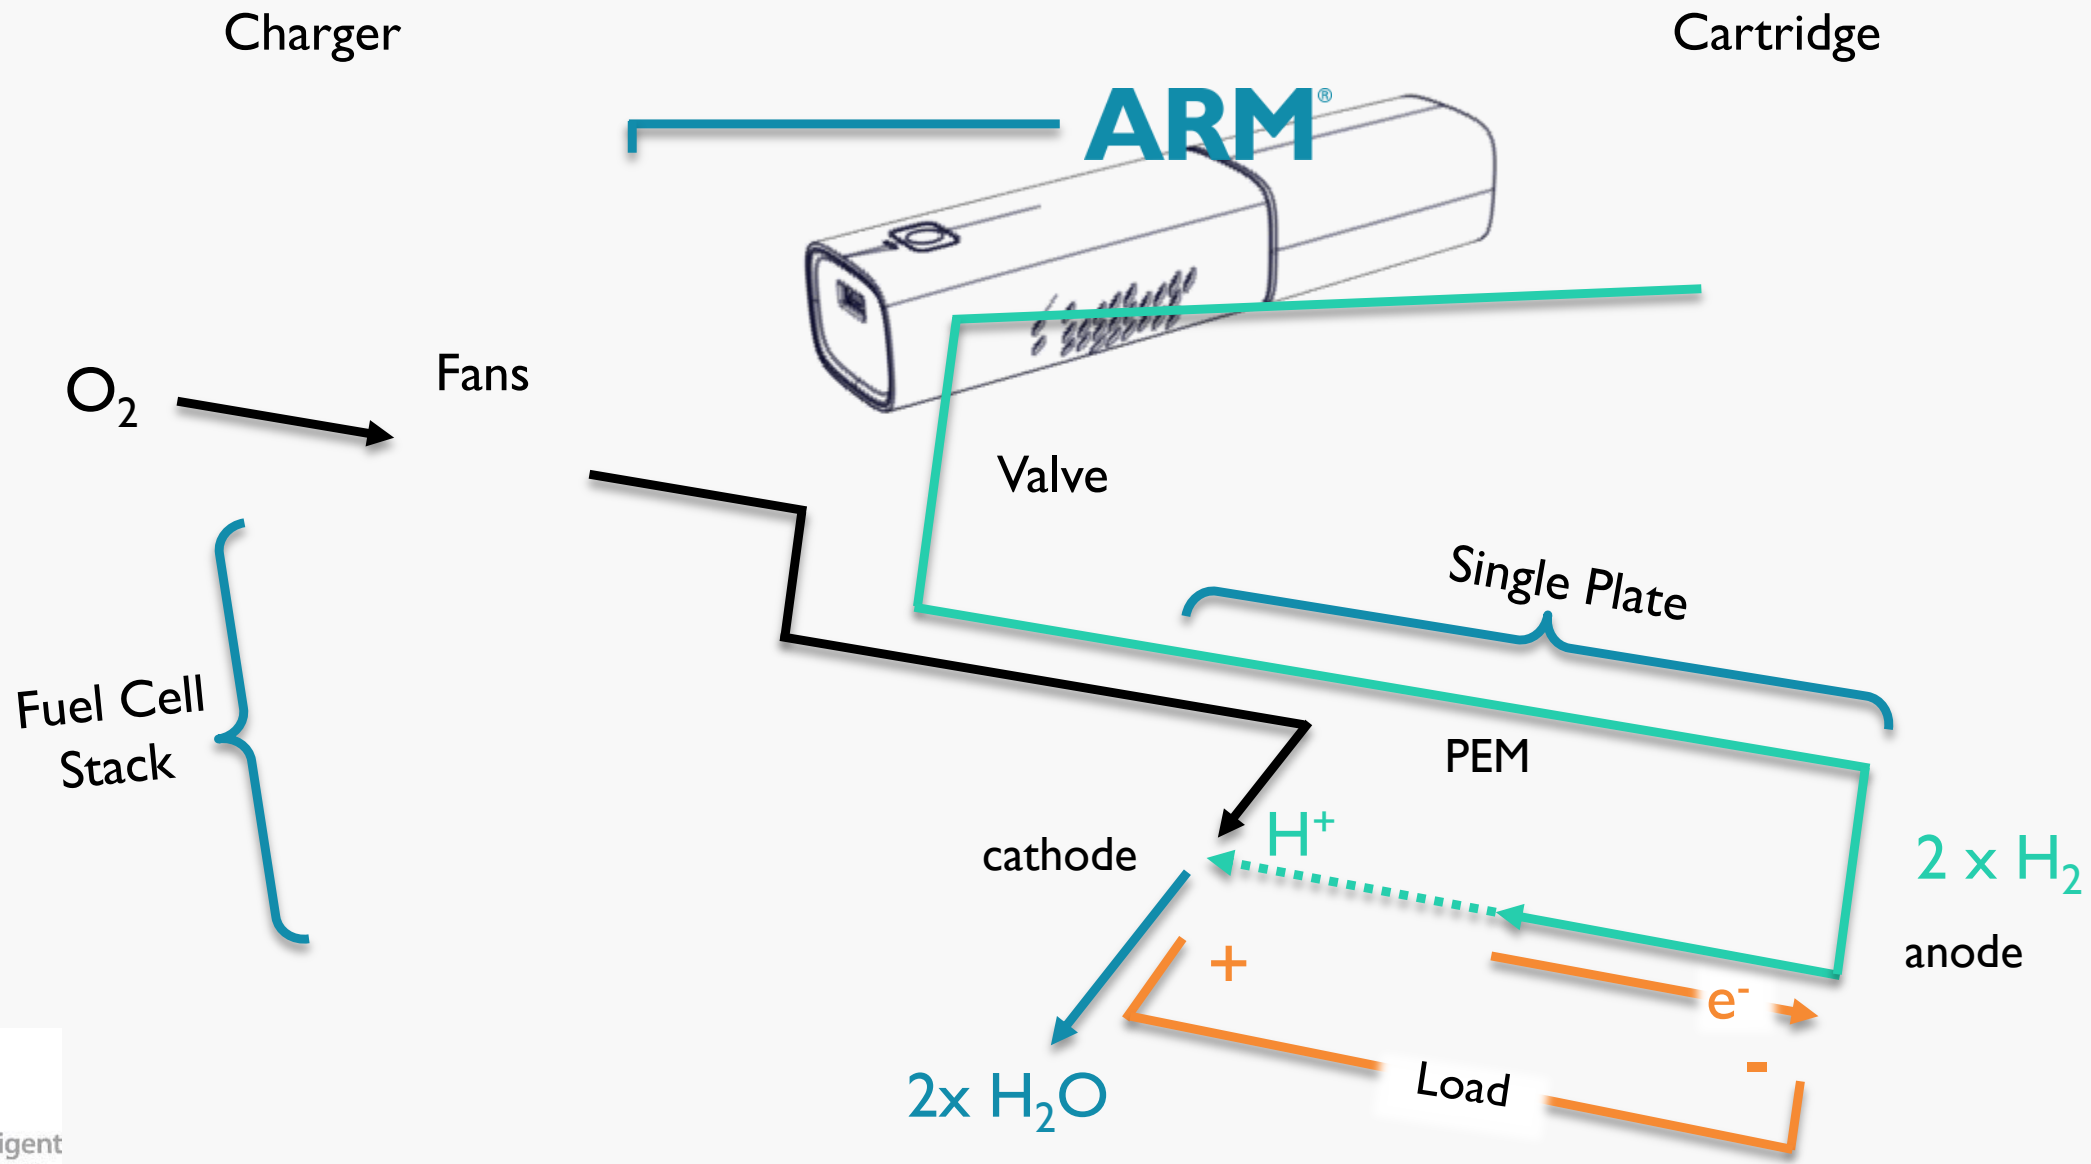

Intelligent Energy™

ARM

Heterogeneous Computing

Top Right

ARM + GPU

2013
Manual Partitioning
C++ & OpenCL/RenderScript

1998
Manual Partitioning
C & Assembler

ARM + DSP

Bottom Left

ARM®

Heterogeneous Computing Performance Drivers

Mobile Image Processing

- Expression detection
- Scene component aggregation
- Blemish reduction
- Should not interfere with GUI

- 2020: 100GOp/s @ 1W

Advanced Driver Assistance

- Collision avoidance
- Driver Assistance
- Enhanced Perception
- Driver monitoring

- 2020: 40GOp/s @ 1W, ASIL-D

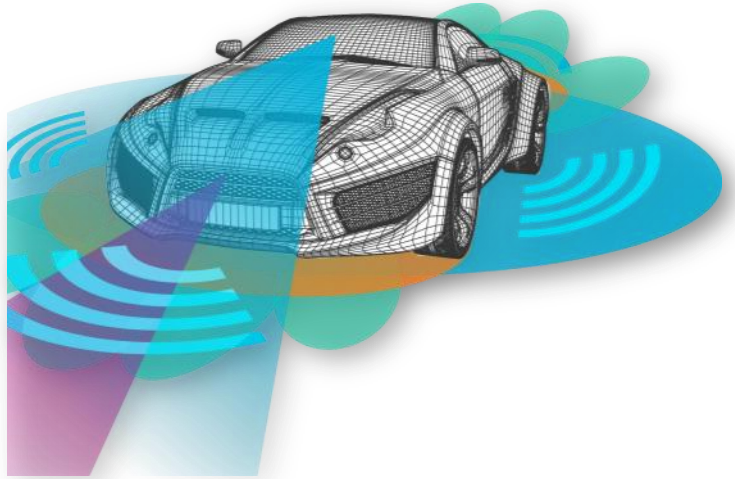

High Performance Computing

- Simulating atoms
- Predicting atmosphere
- Scalability
- Programmability

- 2020: One Exaflop @ 20MW

ARM®

What's Good About GPU style uArchitectures for Efficiency

- Relaxed single-threaded performance
  - No dynamic scheduling
  - No branch prediction
  - No register renaming, no result forwarding
  - Longer pipelines
  - Lower clock frequencies
- Multi-threading to obtain performance
  - Tolerate long latencies to memory
- Increasing the ALU/control ratio
  - Short-vectors exposed to programmers
  - SIMT based execution

ARM

But hard to optimize using GPU programing languages

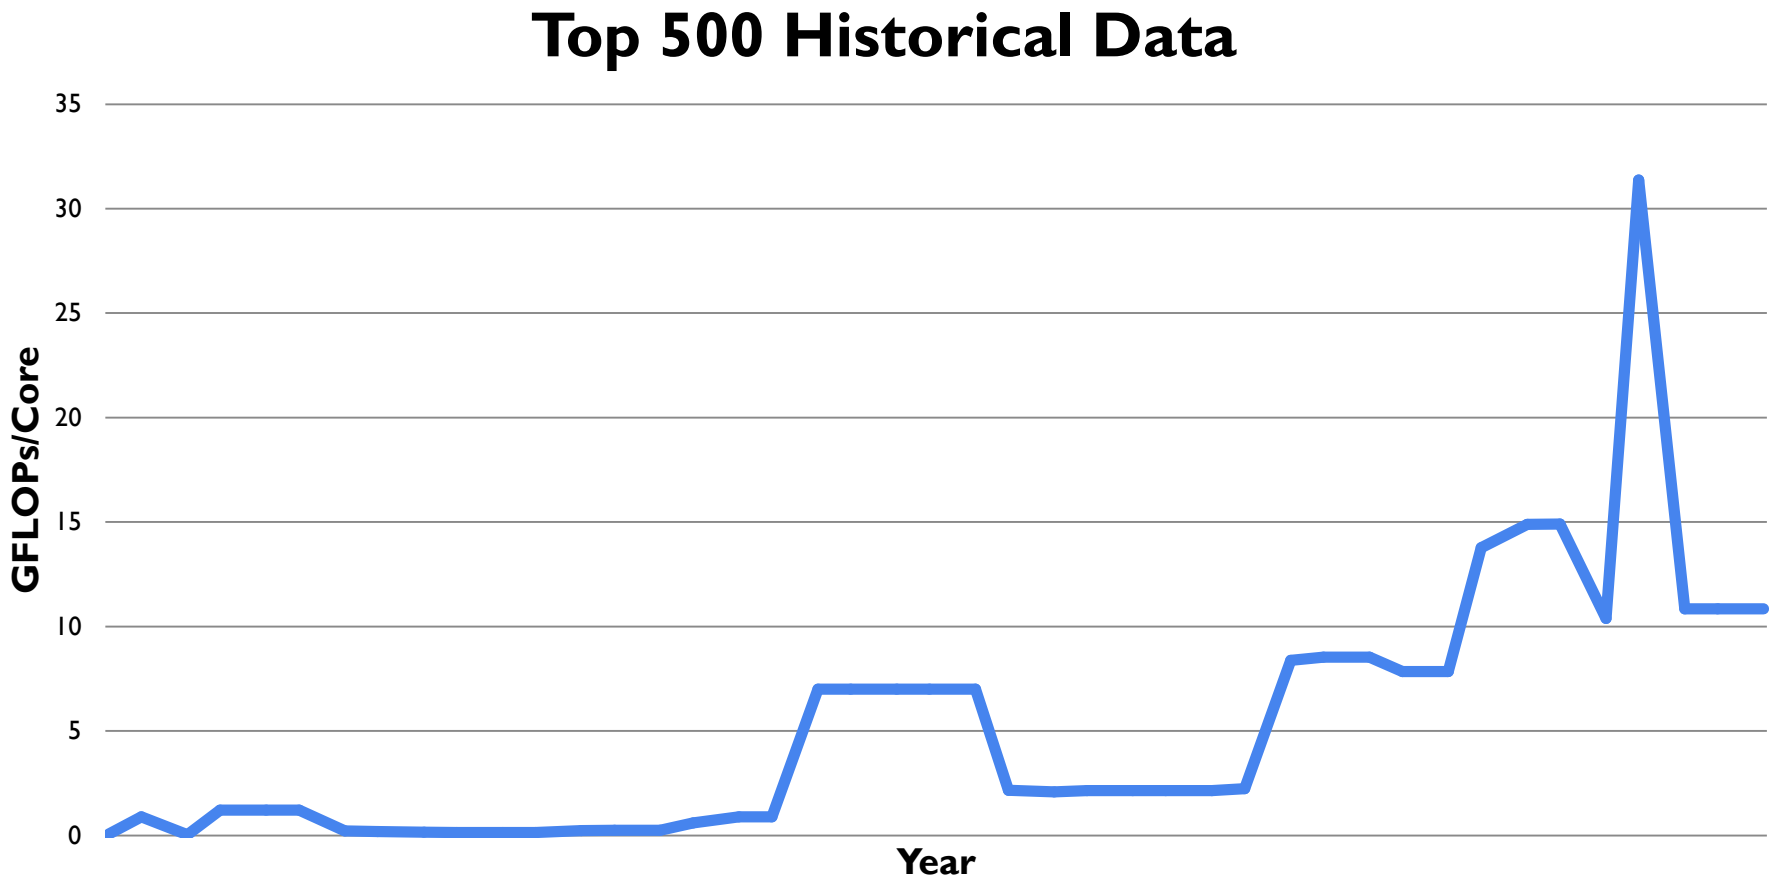

ARM®

Heterogeneous Compute

Homogeneous Architecture

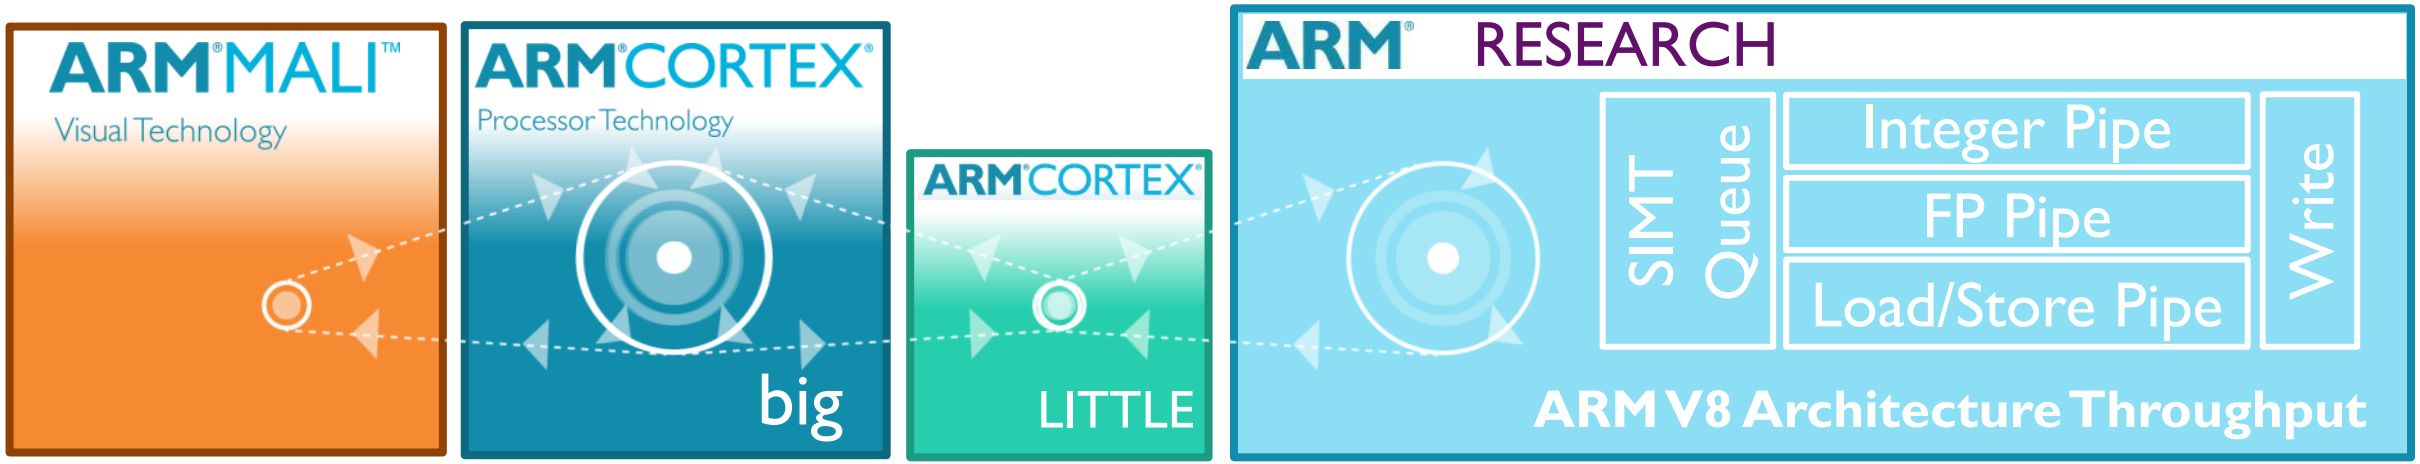

- C++ and OpenMP programming in the same language as the CPU
- Same binary, process, ABI, instruction set, mature familiar tools and ecosystem
- Not necessary to restructure the program for asynchronous execution
  - Very low latency offload
- Not necessary to restructure program to mark up memory

ARM

Performance vs Effort

- We've implemented SGEMM, a matrix-matrix multiplication benchmark, in various ways, to investigate the tradeoff between programmer effort and performance payoff
  - 1 core 500 MHz SIMT CPU vs 1 GHz Cortex-A15-like CPU

SGEMM version	Speedup	Effort
ARM in C	1x	Low

ARM®

Extending big.LITTLE MP for Thermal Management

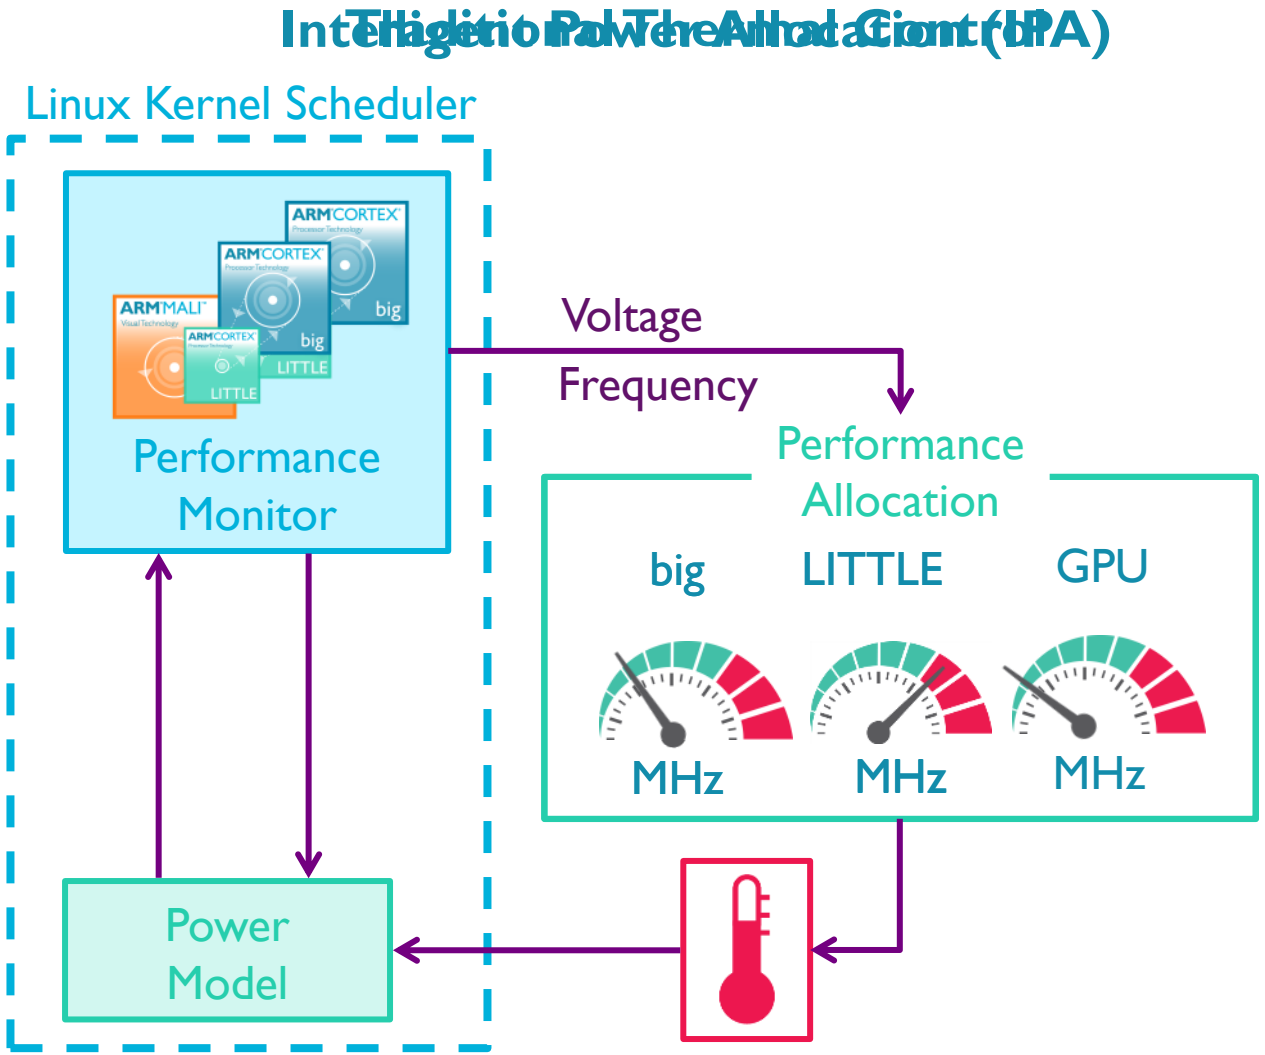

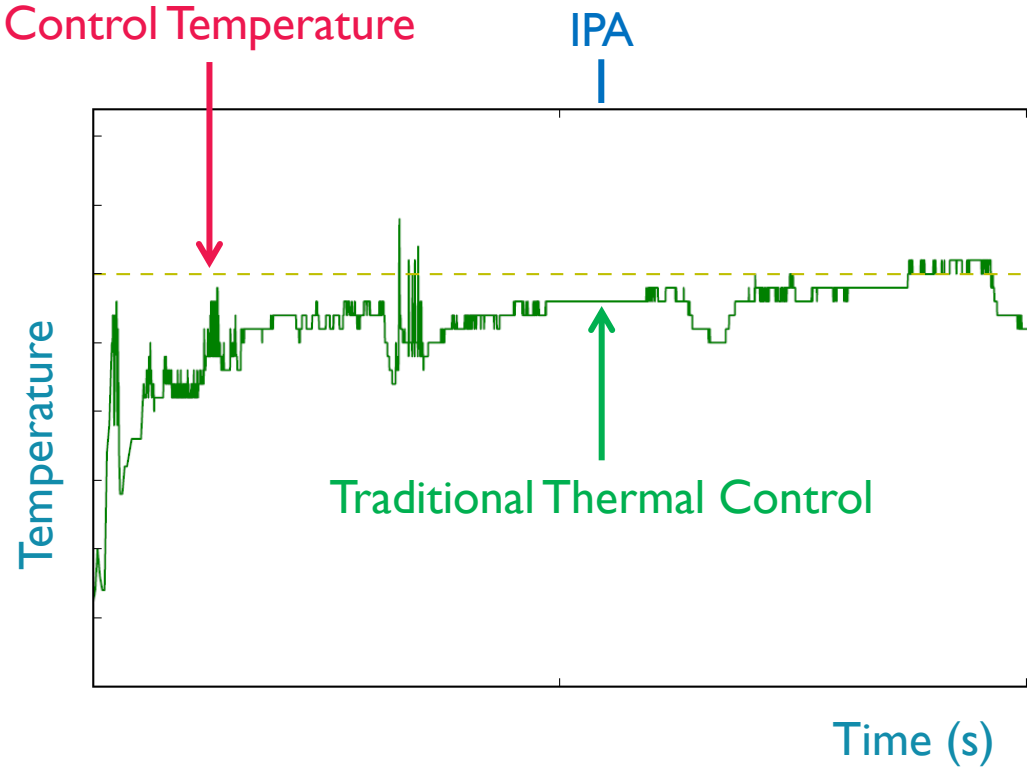

It's Still All About Heterogeneous Systems

ARM®

Aging/Wearout in Mobile Systems?

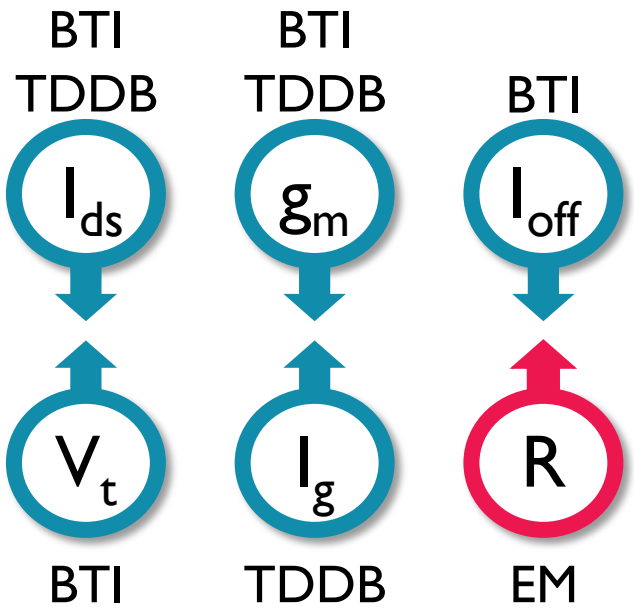

- Aging normal associated with high performance systems
- Degradation due to Bias Temperature Instability (BTI) aging is relevant for DVFS systems
  - High voltage and performance mode causes Vt increase over lifetime
  - Timing is critically affected in low voltage mode
- Workload based aging analysis flow
  - Mid-size ARM CPU: ~100K instances, ~10K flip-flops
  - 28nm library, Vdd = 0.9V, Temp = 105 oC, Lifetime = 3 years
  - Workloads: mp3, web-browse, 3D render, Dhrystone, Video H264
- Timing degradation from ~2.3% (mp3) to ~9% (Video H264)

Workload	Power-State (Active / Sleep) vs. Time	Average Active Time
mp3		0.03
web-browse		0.07
3D rendering		0.24
Dhrystone		0.40
video H264		0.54

ARM®

Is it Safe?

- Memory no longer dominant FIT contributor
  - ECC with physical interleaving reduces the FIT rate to ~0
  - Physical interleaving: multi cell error → single bit error
  - Temporal scrubbing to correct single bit errors
- Unprotected logic now dominates SoC FIT rate
  - No "cheap" way to protect - spatially distributed bits

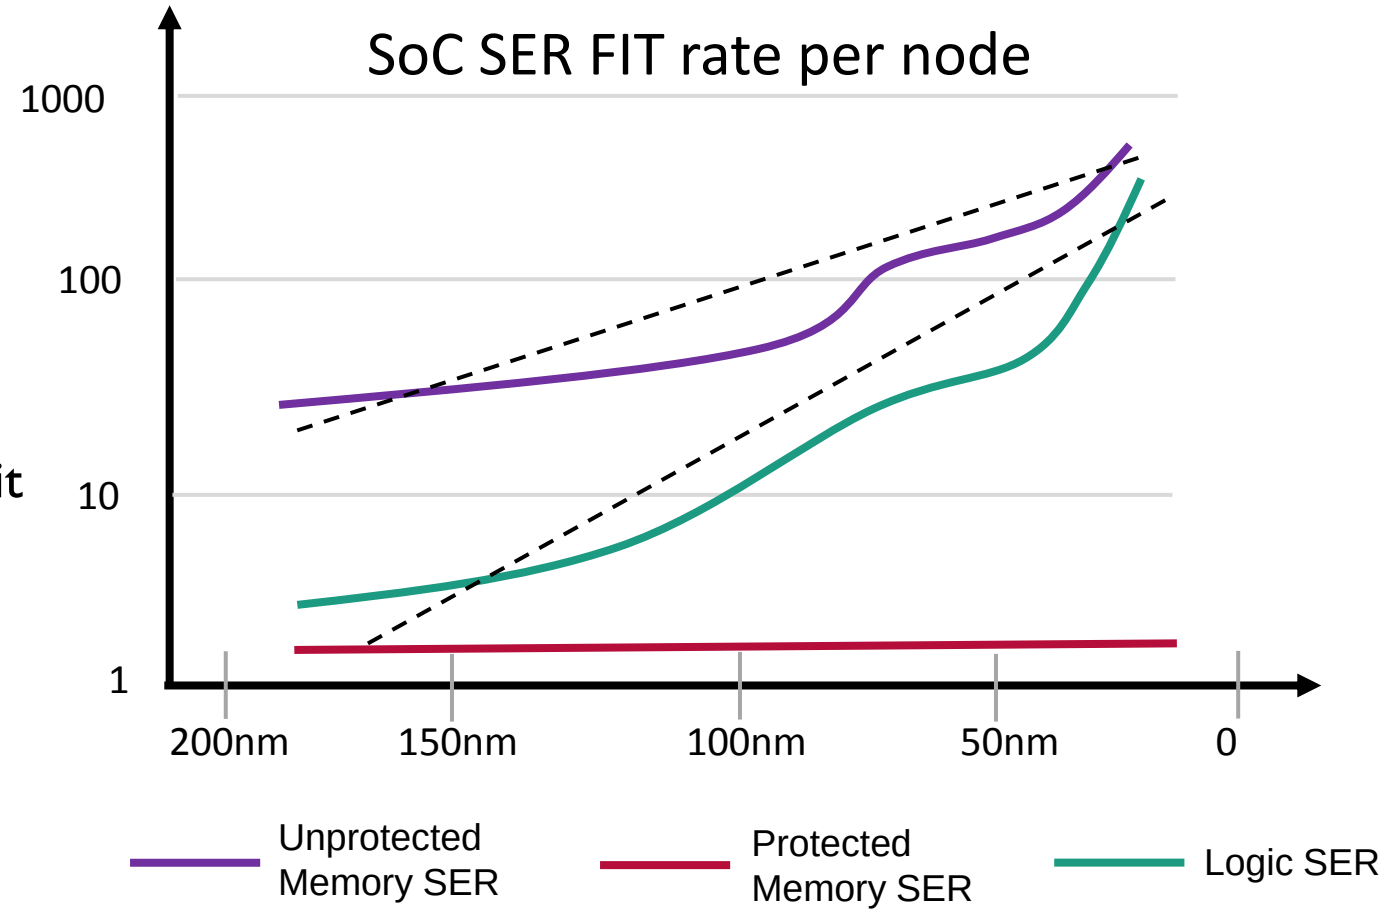

ARM®

Is it Safe?

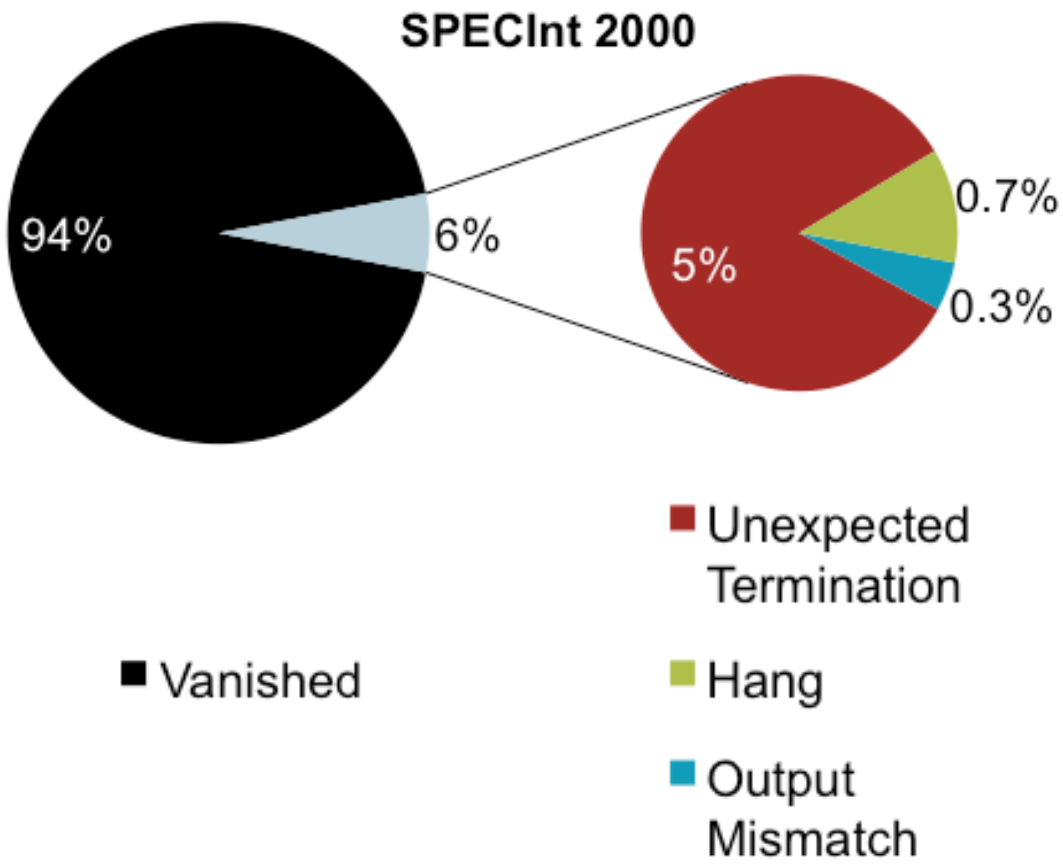

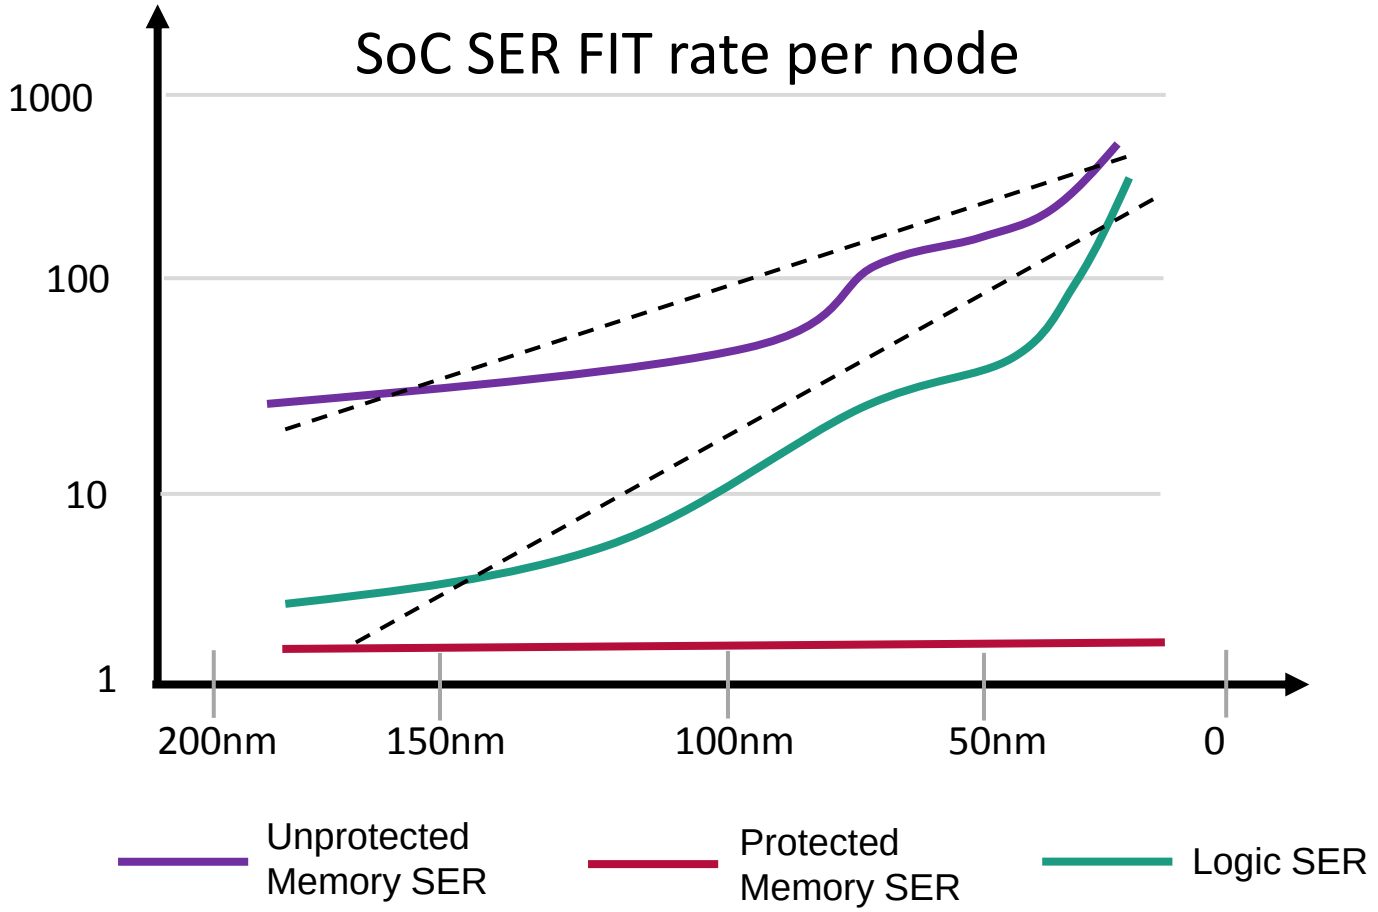

ARM®

Is it Safe?

- Protecting every flip-flop incurs significant area, performance and power penalty

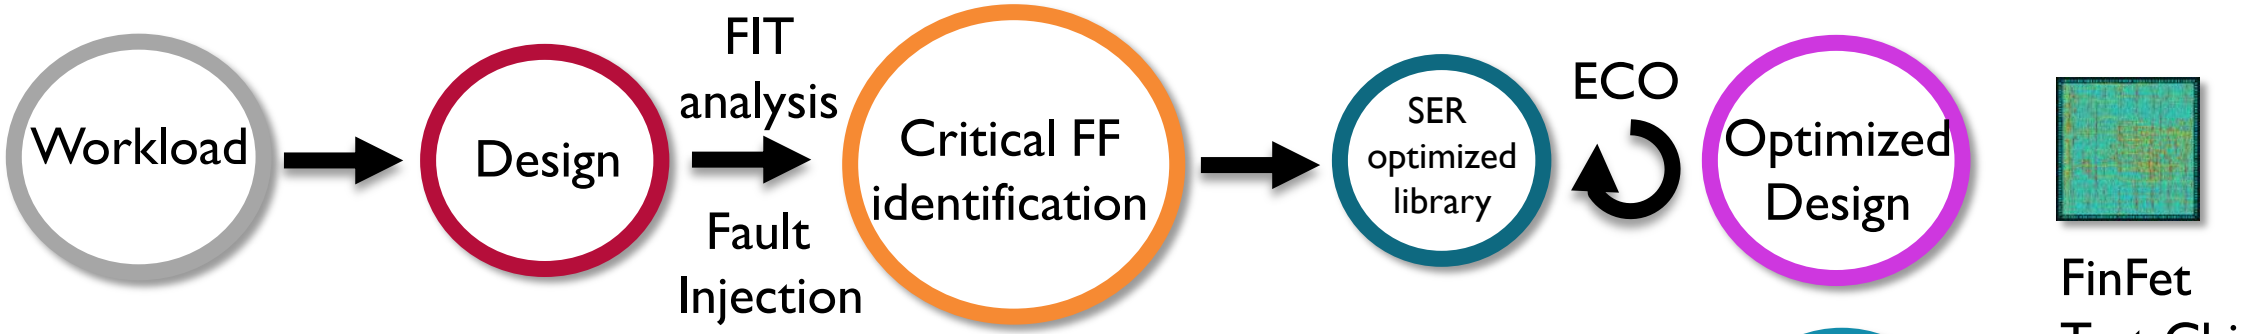

- Identify flip-flops which dominate design FIT rate
  - Fault injection provides each flip-flop's vulnerability
- Critical flip-flops are replaced by SER optimized flip-flops
- Design is iterated in ECO mode to achieve targets

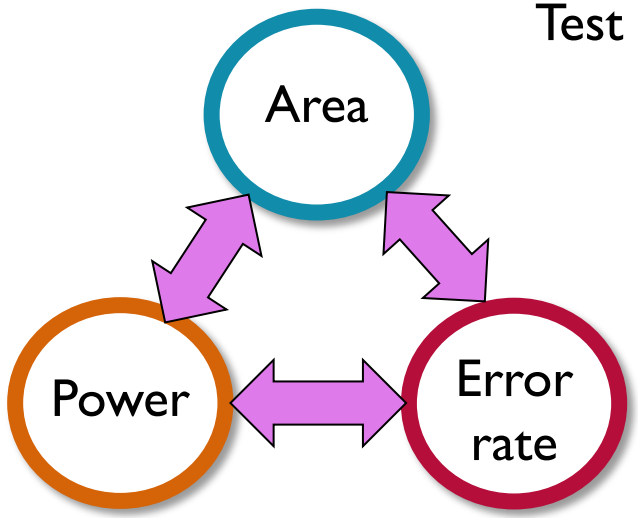

ARM®

It's All About the Memory Efficiency

Standards Don't Always Give You What You Want

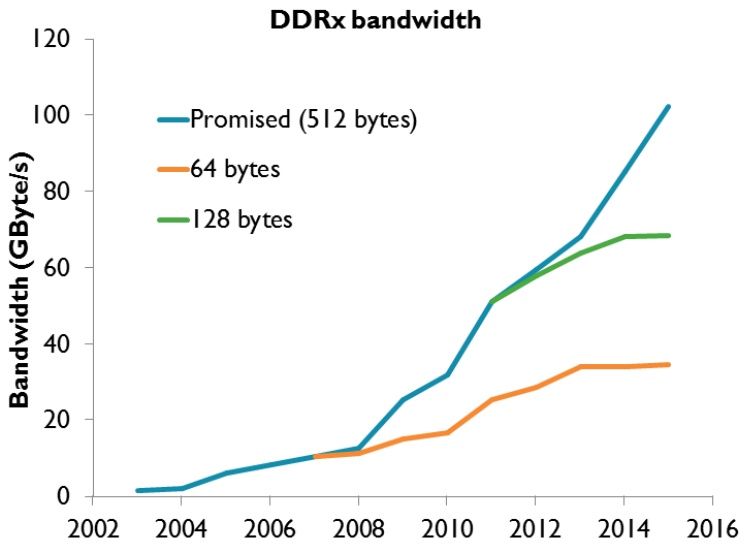

- Delivered bandwidth less than peak
  - 64 bytes natural cache fill size
- Power scales with peak not delivered
- Memory agnostic
  - DDRx becoming increasingly complex
  - Multiple technologies with unique features
- Abstract packet interface
  - Hiding microarchitecture idiosyncrasies
- Feature-scalable
  - Enabling differentiation
- Power-proportional
  - Link and physical layer

andreas.hansson@arm.com

- Engaging the Ecosystem
  - Wendy Elsasser @ ARM chairing Future Memory JEDEC Task Group
  - Continued participation in HMCC
- 2018 post DDR4/LPDDR4
  - Focus on useable bandwidth
  - Abstracted packet interface
- SoC targets

Mobile	>60GB/s	~1.5W
Server	>150GB/s	~10W
Networking	>300GB/s	~15W

ARM®

It's Still All About the Memory

ARM®

Top Right

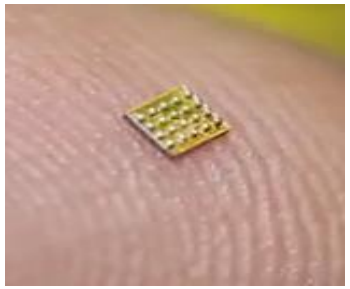

and Bottom Left

ARM®

Scalable Area, Performance and Battery Life

Mobile Implementation

Cortex-A7 MP2
1.2 GHz, 2.25 mm^2

High-end Wearable

Cortex- A7 MP2
500 MHz, 1.1 mm^2

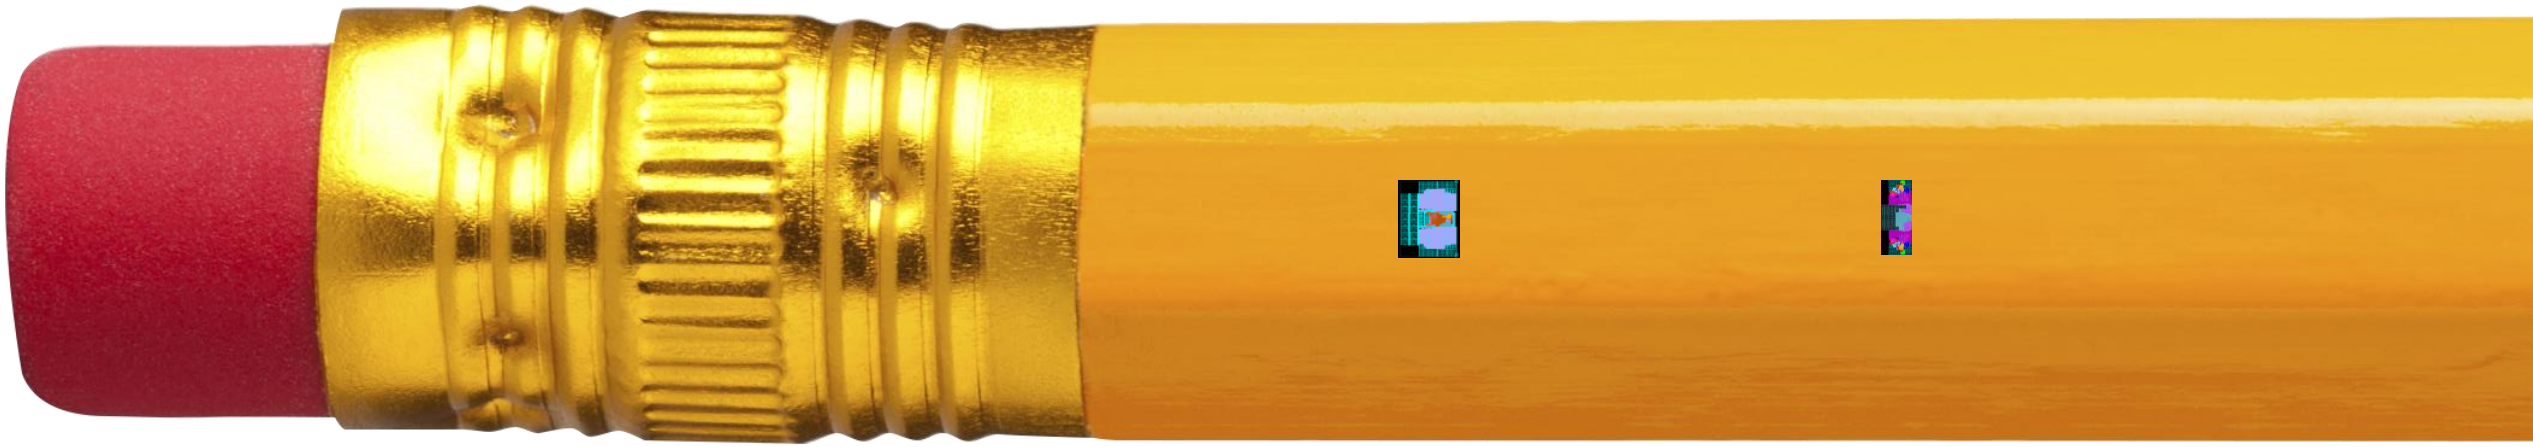

Area calculations based on 28HPM node, 9T

ARM®

Scalable Area, Performance and Battery Life

Mobile Implementation	High-end Wearable	Optimized Wearable	Optimized 40LP Wearable
Cortex-A7 MP2 1.2GHz, 2.25 mm^2	Cortex-A7 MP2 500MHz, 1.1 mm^2	Cortex-A7 UP 500MHz, 0.36 mm^2	Cortex-A5 UP 500MHz, 0.49 mm^2

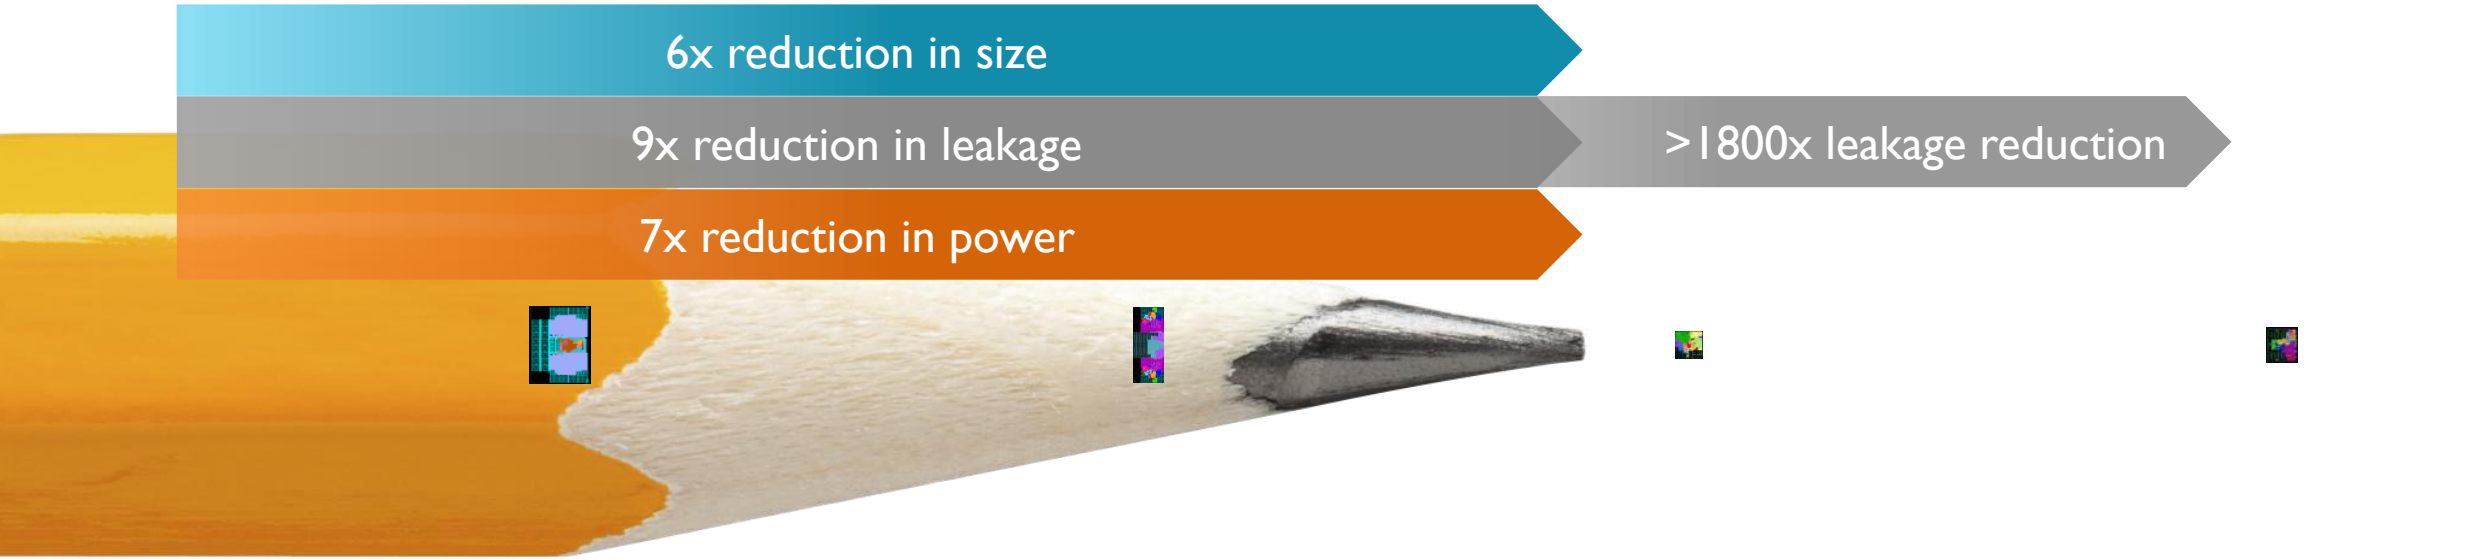

Area based on 28HPM node, 9T, unless otherwise indicated

ARM®

Scalable Area, Performance and Battery Life

Optimized Wearable

Cortex-A7 UP
500MHz, 0.36 mm^2

Optimized IOT

Cortex-M0
40MHz, 0. 05mm^2

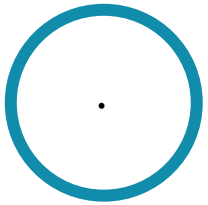

Area based on 28HPM node, 9T, unless otherwise indicated

ARM®

Bottom Left Efficiency

Applications
Cortex-A5 System
250MHz

Comms
Cortex-M0 System
40MHz

Sensor Hub
Cortex-M0+ System
1MHz – 20kHz

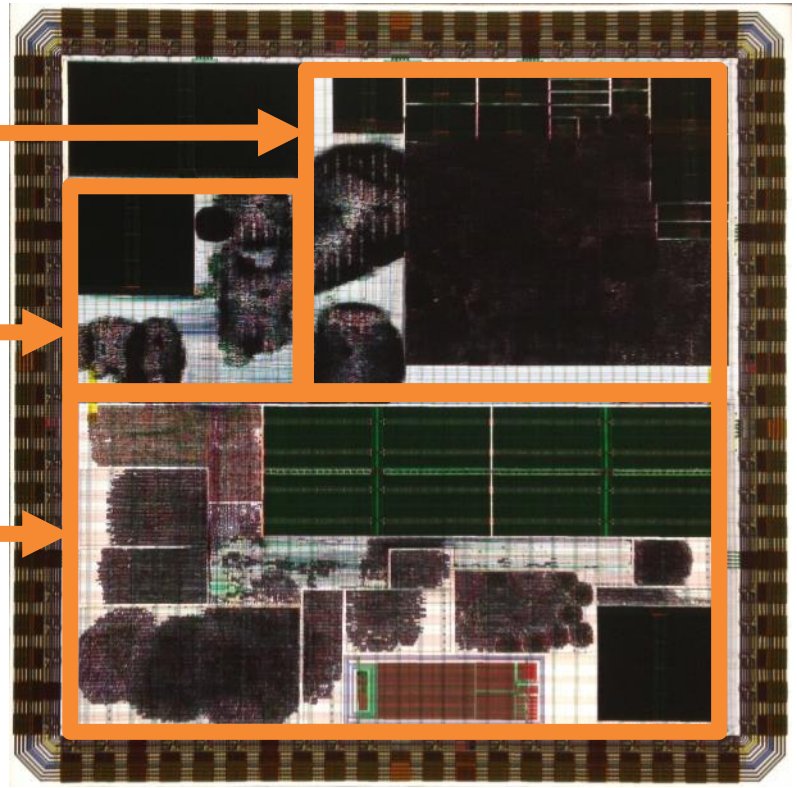

Bottom Left Efficiency 1.2V to 0.2V AM

65nm
1.2V to 0.2V
On Chip Regulators
Direct Battery operation

12 Power Domains

Test Logic	2kB Boot ROM
Debug Interface	16kB 6T SRAM 1.2V
Serial Interface	4kB 10T SRAM LV
Host Interface	4kB 10T SRAM LV
RTC	Bus Fabric
128bit AES H/W	Cortex-M0+

Sensor Hub
Cortex-M0+ System
1MHz – 20kHz

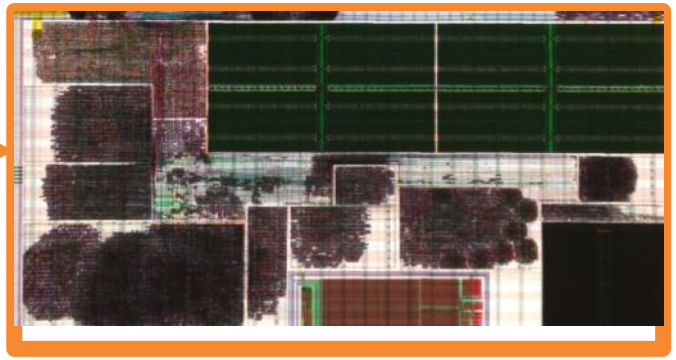

ARM®

Bottom Left Efficiency Implementation Matters

12 Power Domains

Test Logic	2kB Boot ROM
Debug Interface	16kB 6T SRAM 1.2V
Serial Interface	4kB 10T SRAM LV
Host Interface	4kB 10T SRAM LV
RTC	Bus Fabric
128bit AES H/W	Cortex-M0+

Sensor Hub
Cortex-M0+ System
1MHz – 20kHz

1.2V to 0.2V
On Chip Regulators
Direct Battery operation

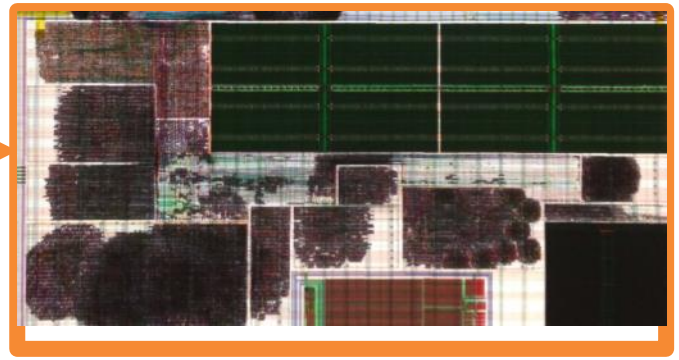

- SRAM contribution to system standby/retention power is >80%
- Level Shifters must be optimized for range & reduced variability
- Power gate voltage drop significantly impacts frequency and leakage
- EDA tools assume RC dominates but at low voltage gate-delay dominates

ARM®

Bottom Left Efficiency 1000 x Power

- Minimum Power Point
  - 200mV @ 20 kHz
  - 1/6000 x power
  - 1/3000 x frequency
  - Energy Scavenging
- Minimum Energy Point
  - 400mV @ 1MHz
  - 1/1000 x energy
  - 1/100 x frequency

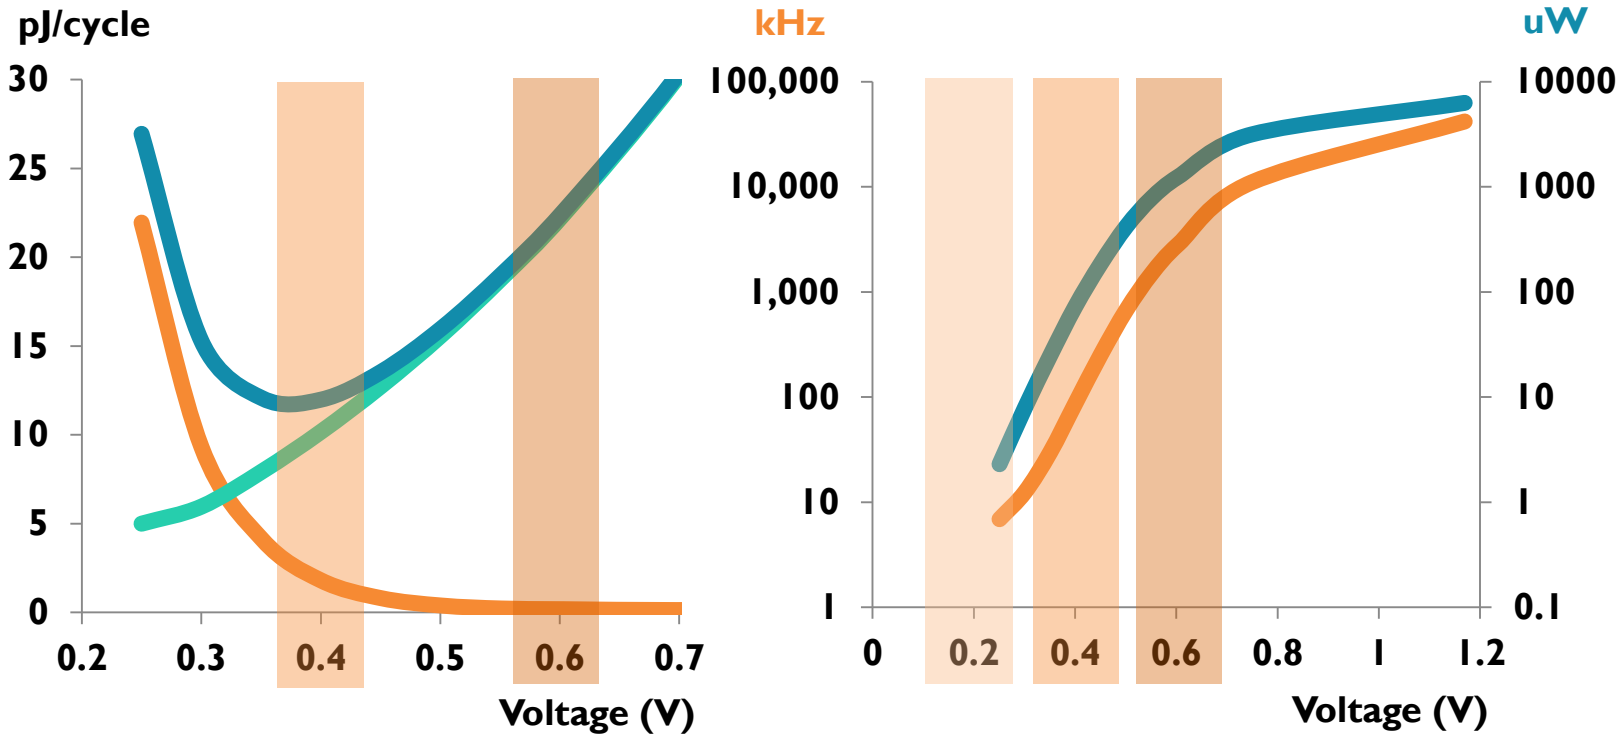

- Circuits can be made to work but the system level leakage needs great care
  - 200mV Min Voltage – useful for *power* limited energy harvesters
  - 400mV Min Energy – best for *energy* limited batteries
  - 600mV Near Threshold – half the energy savings but much lower complexity & risk

ARM®

Great, but what can you do with it?

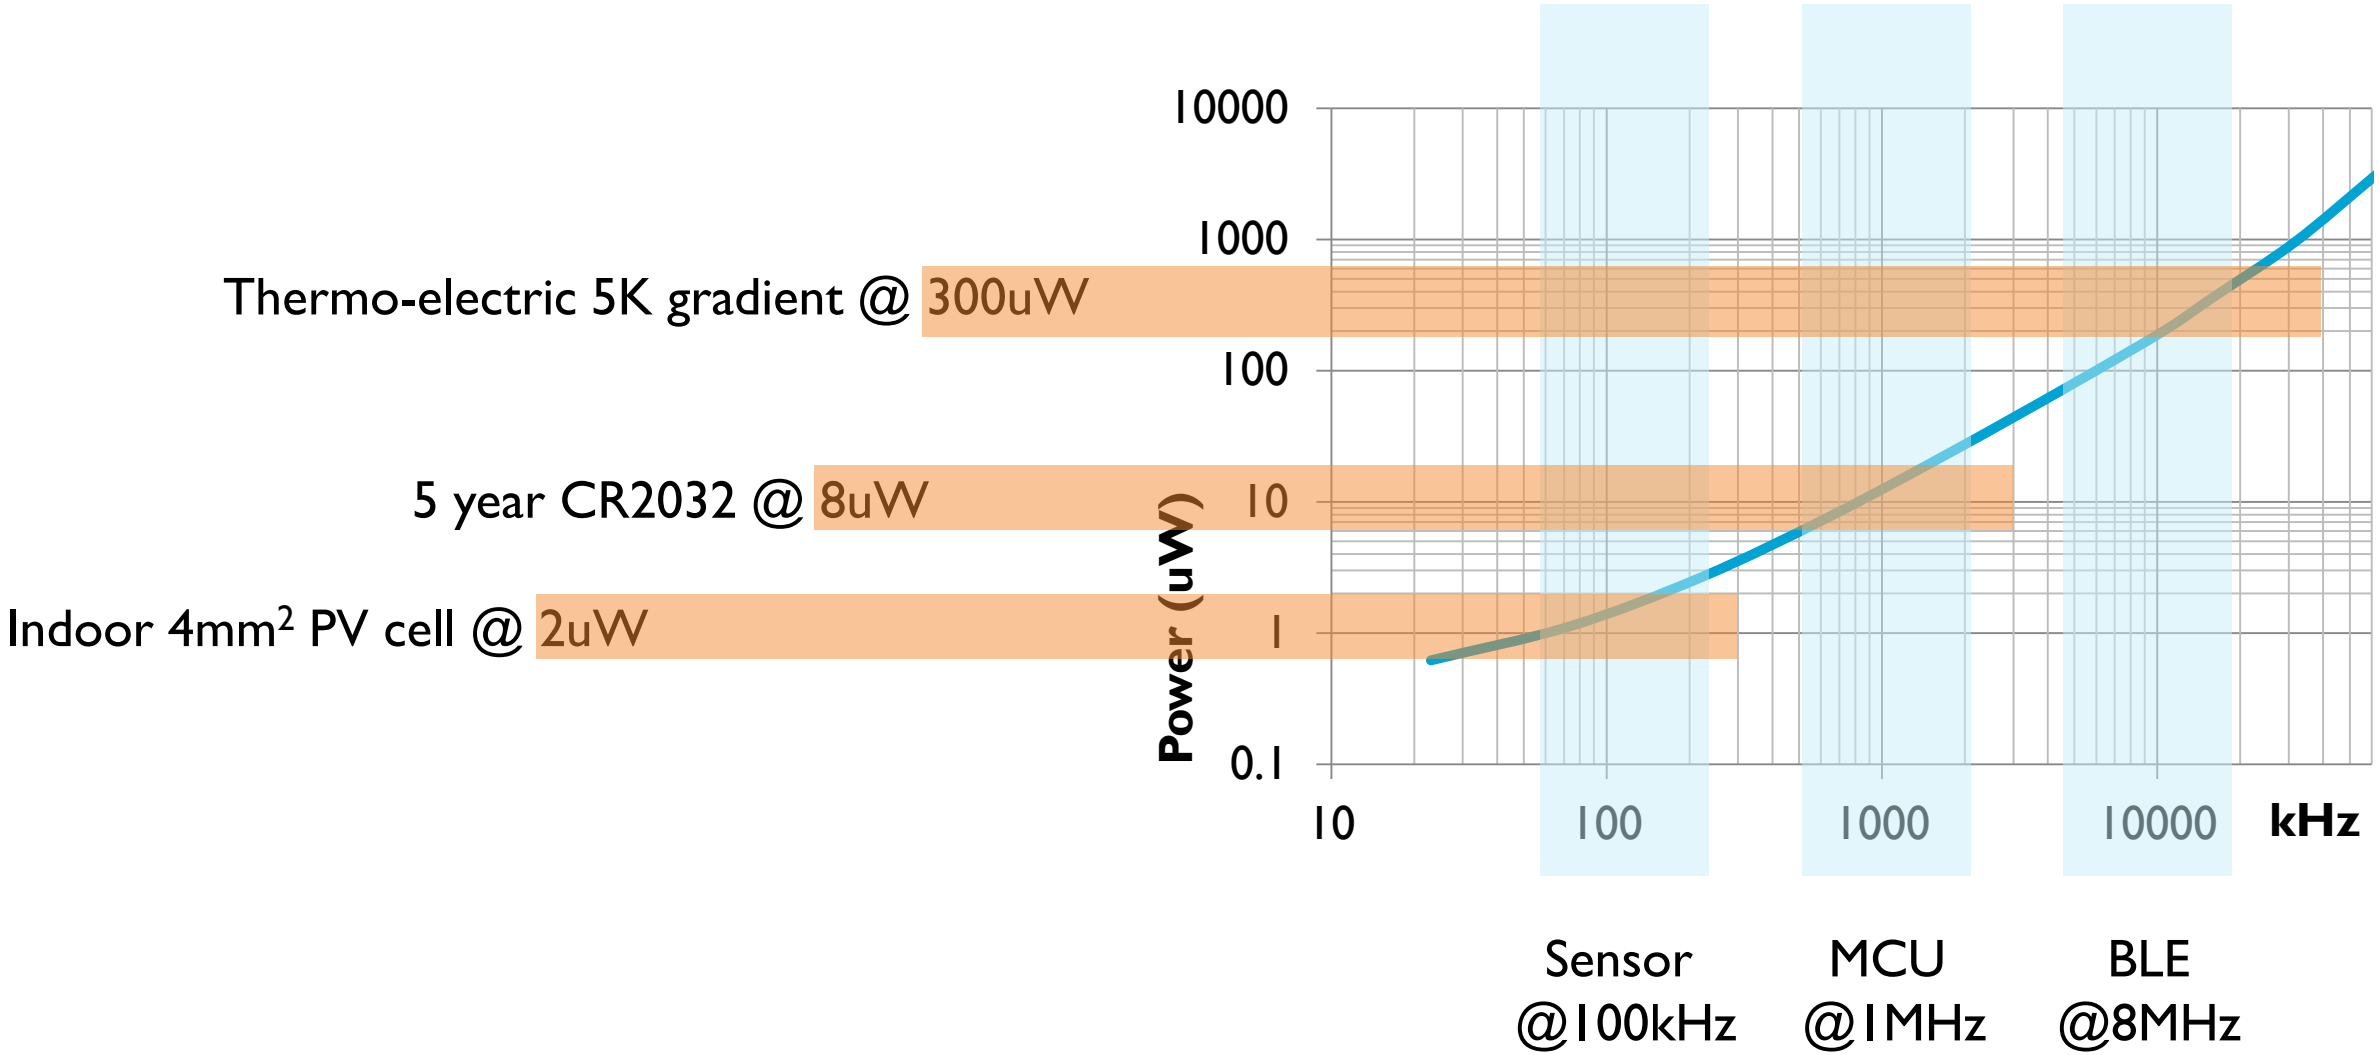

ARM®

It’s still all about the power

It’s still all about heterogeneous solutions

It’s still all about the memory

In the future the cloud drives device requirements

ARM®