A Breakthrough New CPU Architecture Revives IPC Scaling

Mohammad Abdallah
Founder, President and CTO

Linley Processor Conference
October 23, 2014

Soft Machines™

Introducing Soft Machines™

- ***Emerging from stealth mode***
- ***Developed new VISC™ Architecture***
- ***7 years, $125M R&D***
- ***~250 employees , 75+ patents filed***

The Death of CPU Scaling

Microprocessor Scaling Realities after 2004

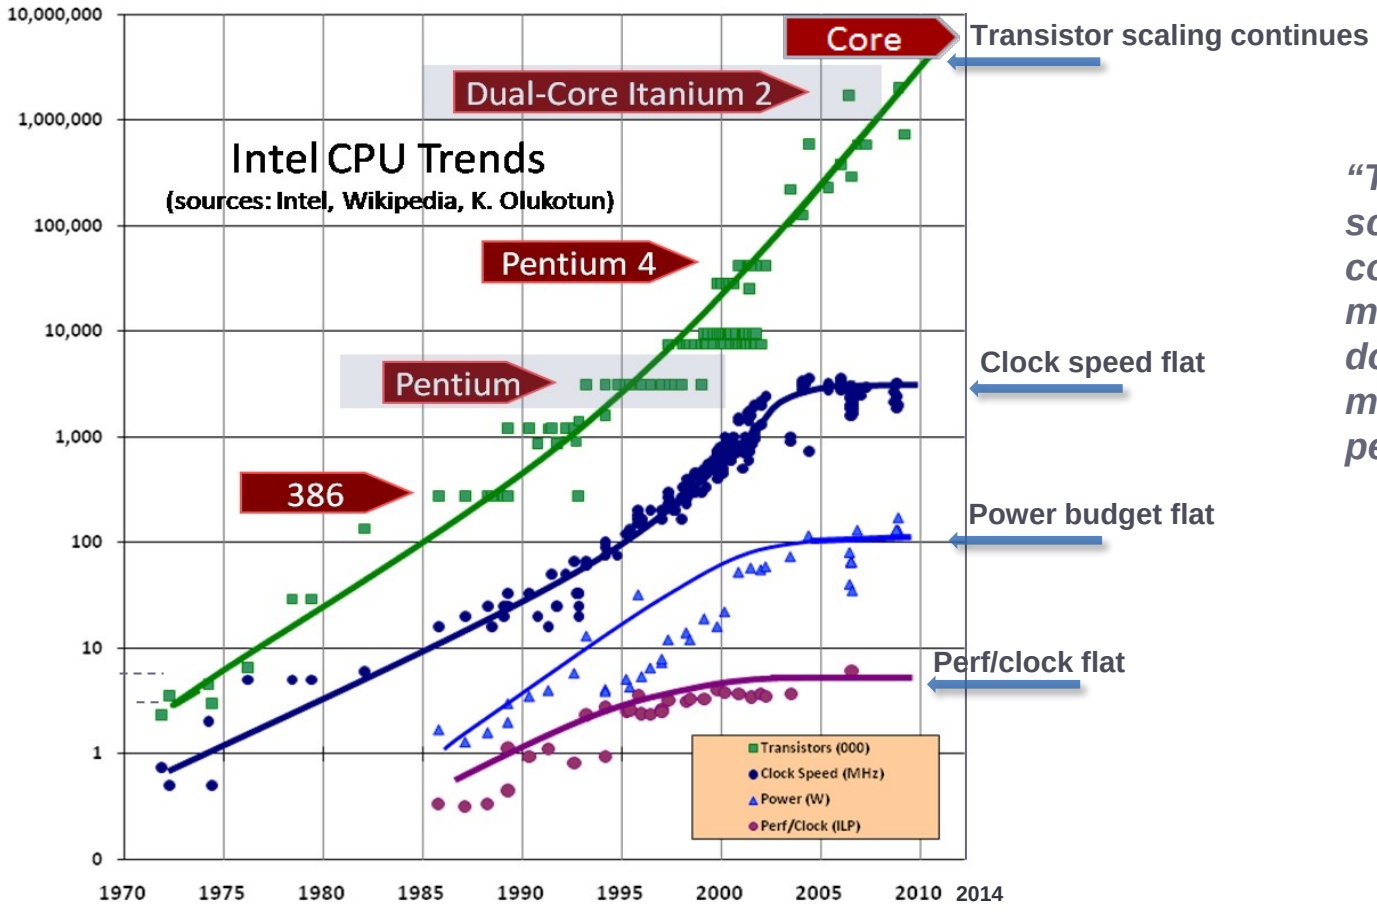

Transistor scaling continues

Clock speed flat

Power budget flat

Perf/clock flat

"The failure of CPU scaling after 30 years of continual improvements may have slammed the door on the easiest and most common type of performance scaling…"

The Death of CPU Scaling
ExtremeTech (2012)

Source: "The Free Lunch is Over", Herb Sutter

Industry Response: Multi-Core

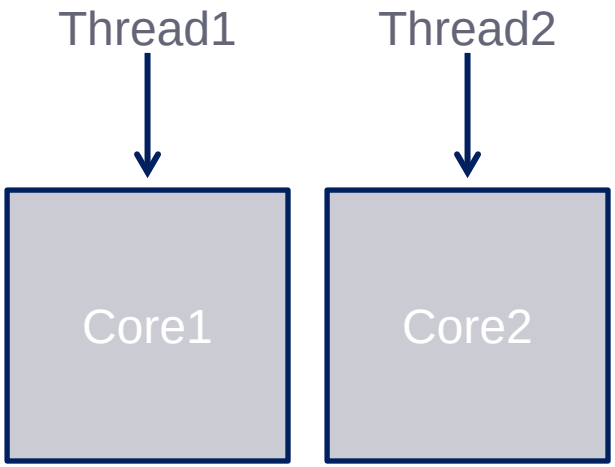

Advantages:

- Utilizes growing transistor budget
- Performance scaling for parallel code
- Improves throughput

Challenges:

- ST performance doesn't scale
- Threading/multicore coding complexity
- Amdahl's Law of diminishing returns
- Dark silicon

CPU Architecture Challenge

- ***Revive CPU performance scaling***
- ***Utilize Moore's Law transistor scaling***
- ***Mitigate dark silicon***
- ***Liberate ISA dependency***

VISC™ Architecture Wave

Software Scalability/Productivity

Assembly	Compilation	Concurrency Extraction

CISC	RISC	VISC
(IBM/Intel)	(MIPS)	(Soft Machines)

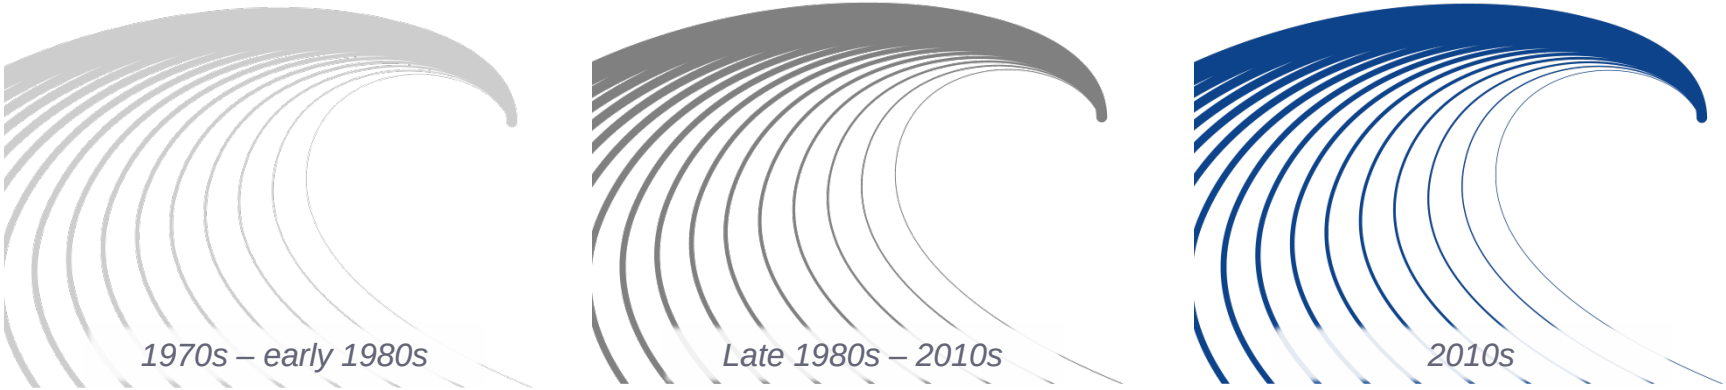

Short Pipeline	Deep OoO Pipeline	Virtual Cores/Threads
Code Memory size	Processor Speed	Processor Power

Device Physics Scalability

VISC Architecture scales on both physical and software productivity layers

VISC™ Processor Block Diagram

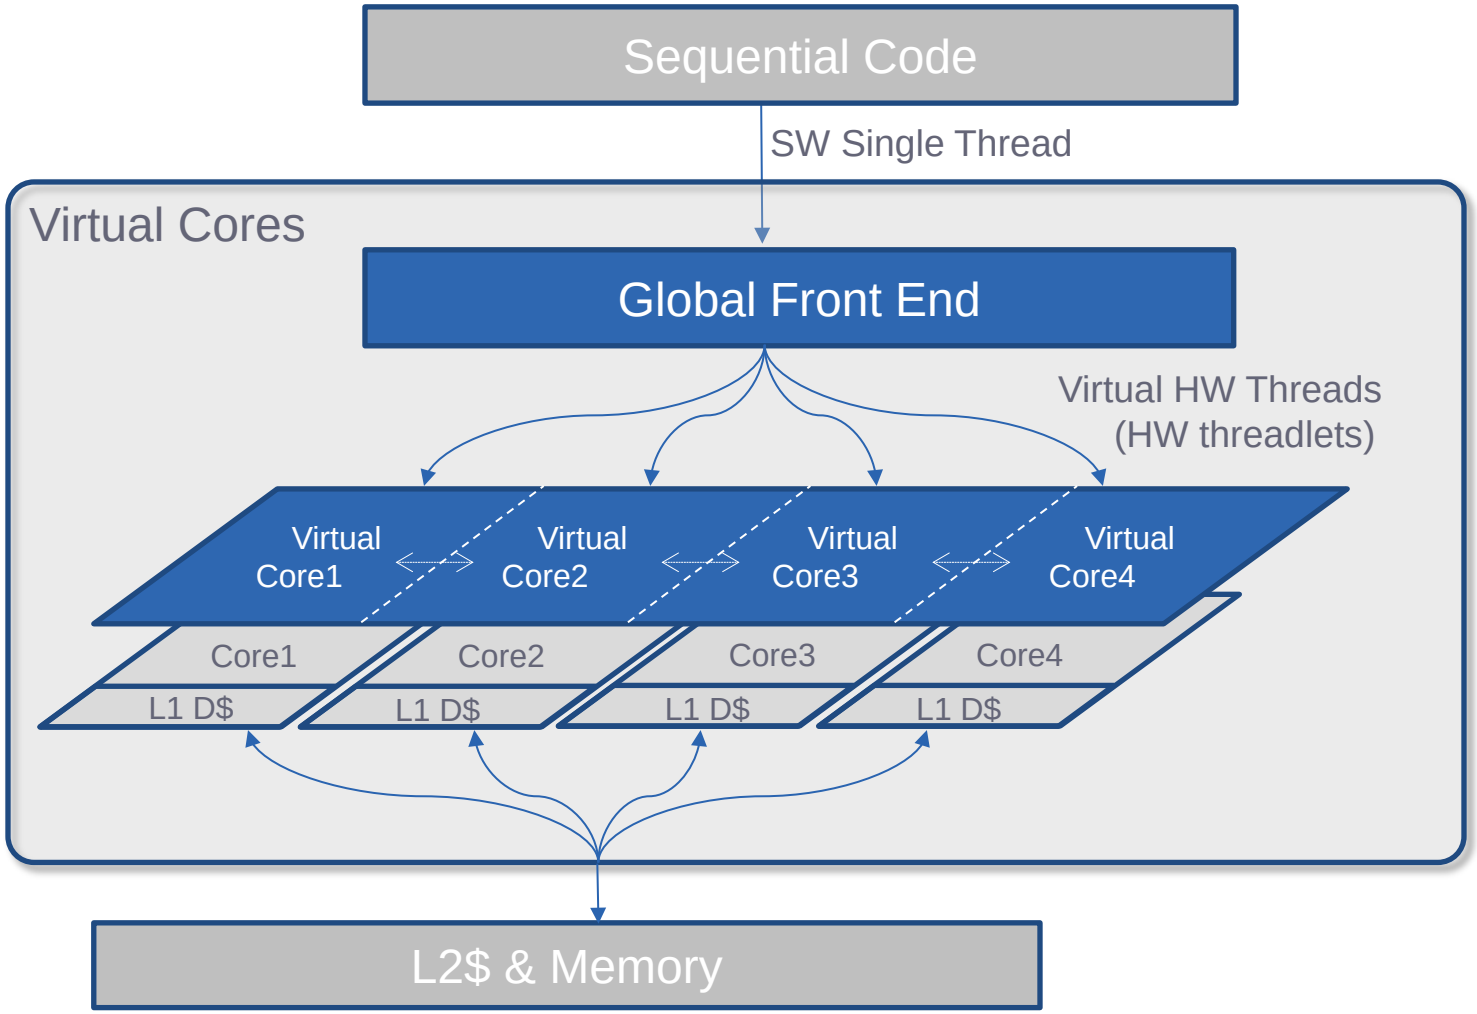

VISC™ CPU Usage Example

Single SW Thread

Heavy App

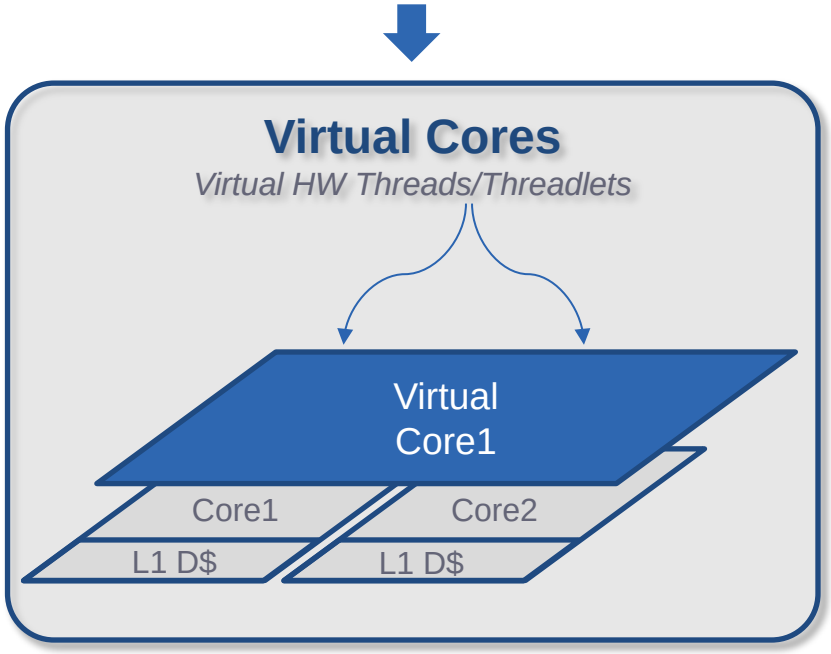

or

Dual SW Threads

Heavy App *Light App*

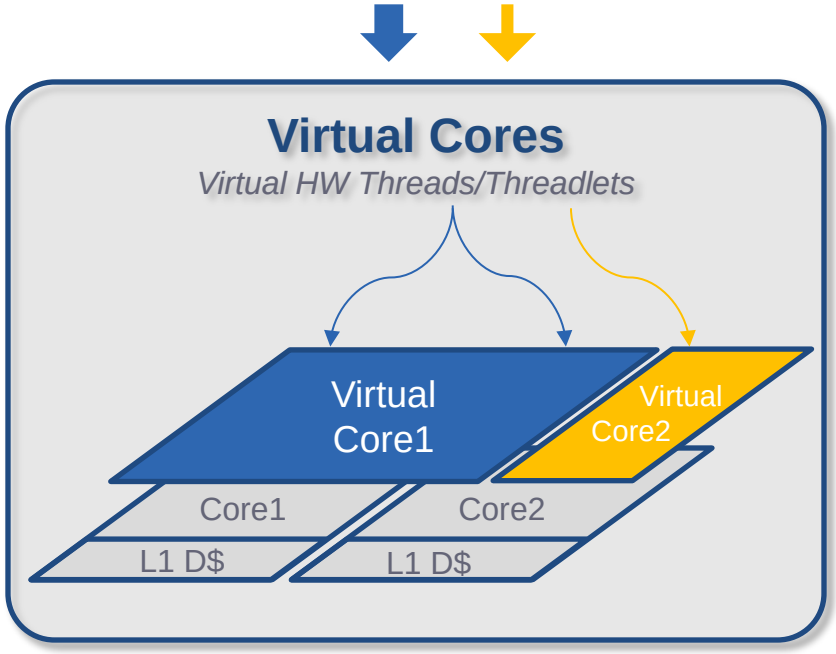

- VISC dynamically allocates resources across virtual cores based on individual application needs
- Performance/watt balanced for both single & multi-thread applications

VISC™ Architecture Prototype Pipeline

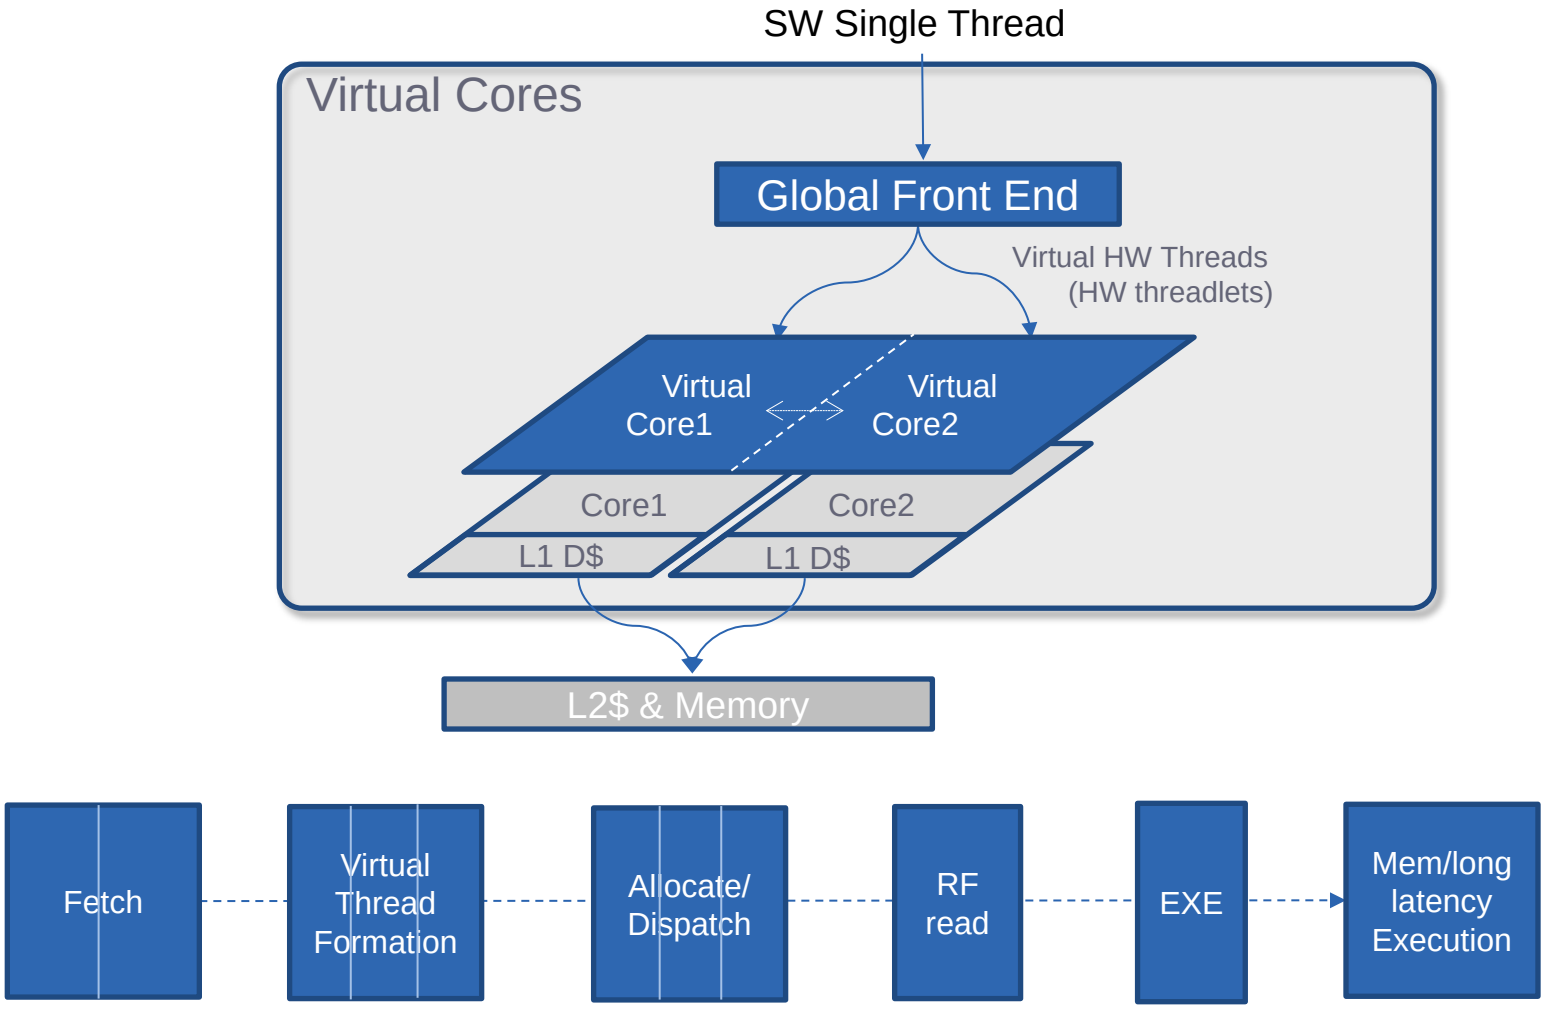

Pipeline of Virtual Threads Across the Virtual Cores

VISC™ Revives IPC Curve

Mobile CPU designs are pursuing higher ARCH/µARCH complexity

	ARM A15 1C	Intel Atom 1C	Soft Machines 2VC Proto	Apple A7 1C	ARM A57 1C	Intel Haswell 1C
Compiled Code	32-bit	32-bit	32-bit	32-bit	32-bit	64-bit
Cache	1M	2M	1M	1M+4M	2M	2M
Pipeline	Moderate	Moderate	Shallow	Moderate	Moderate	Deep
IPC(SPEC 2006)*	**0.71**	**0.69**	**2.1**	**1.0**	**.87**	**1.39**

* Company conducted benchmark tests and projections, using industry-standard Compiler GCC 4.6 or equivalent

VISC™ Concurrency Extraction
Linear vs. Quadratic Complexity

- Extracting ILP has significant complexity
- OoO complexity increases quadratically with machine width
- VISC complexity increases linearly with number of virtual cores
- VISC Performance/Watt utilizes linear scaling

System Energy Approach: DRVFS

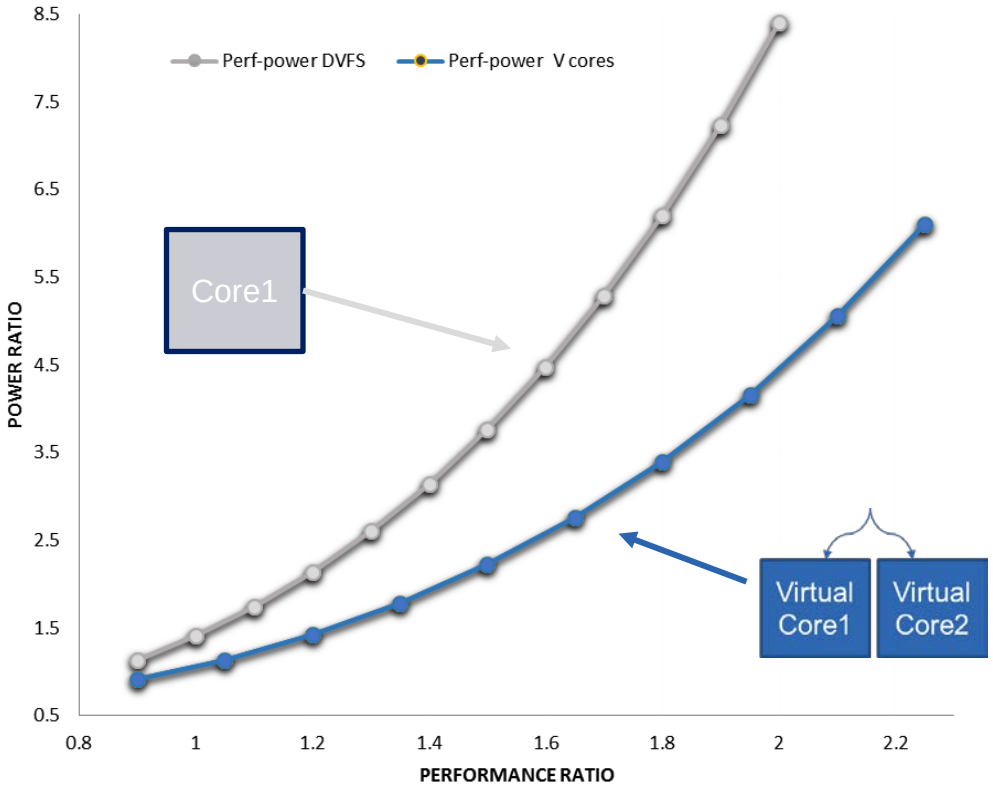

Use Case: Rush to low power mode (boosting performance or response time)

Physical Cores – DVFS

- DVFS: quadratic increase in power
  - $P \propto V^2 * F$
- Lower Perf/MHz requires DVFS scaling UP

Virtual Cores – D**R**VFS

- DRVFS: linear increase in power
  - $P \propto$ No. of virtual core resources
- Higher Perf/MHz enables DVFS scaling DOWN

VISC™ Single Thread SPEC/Watt

Soft Machines™

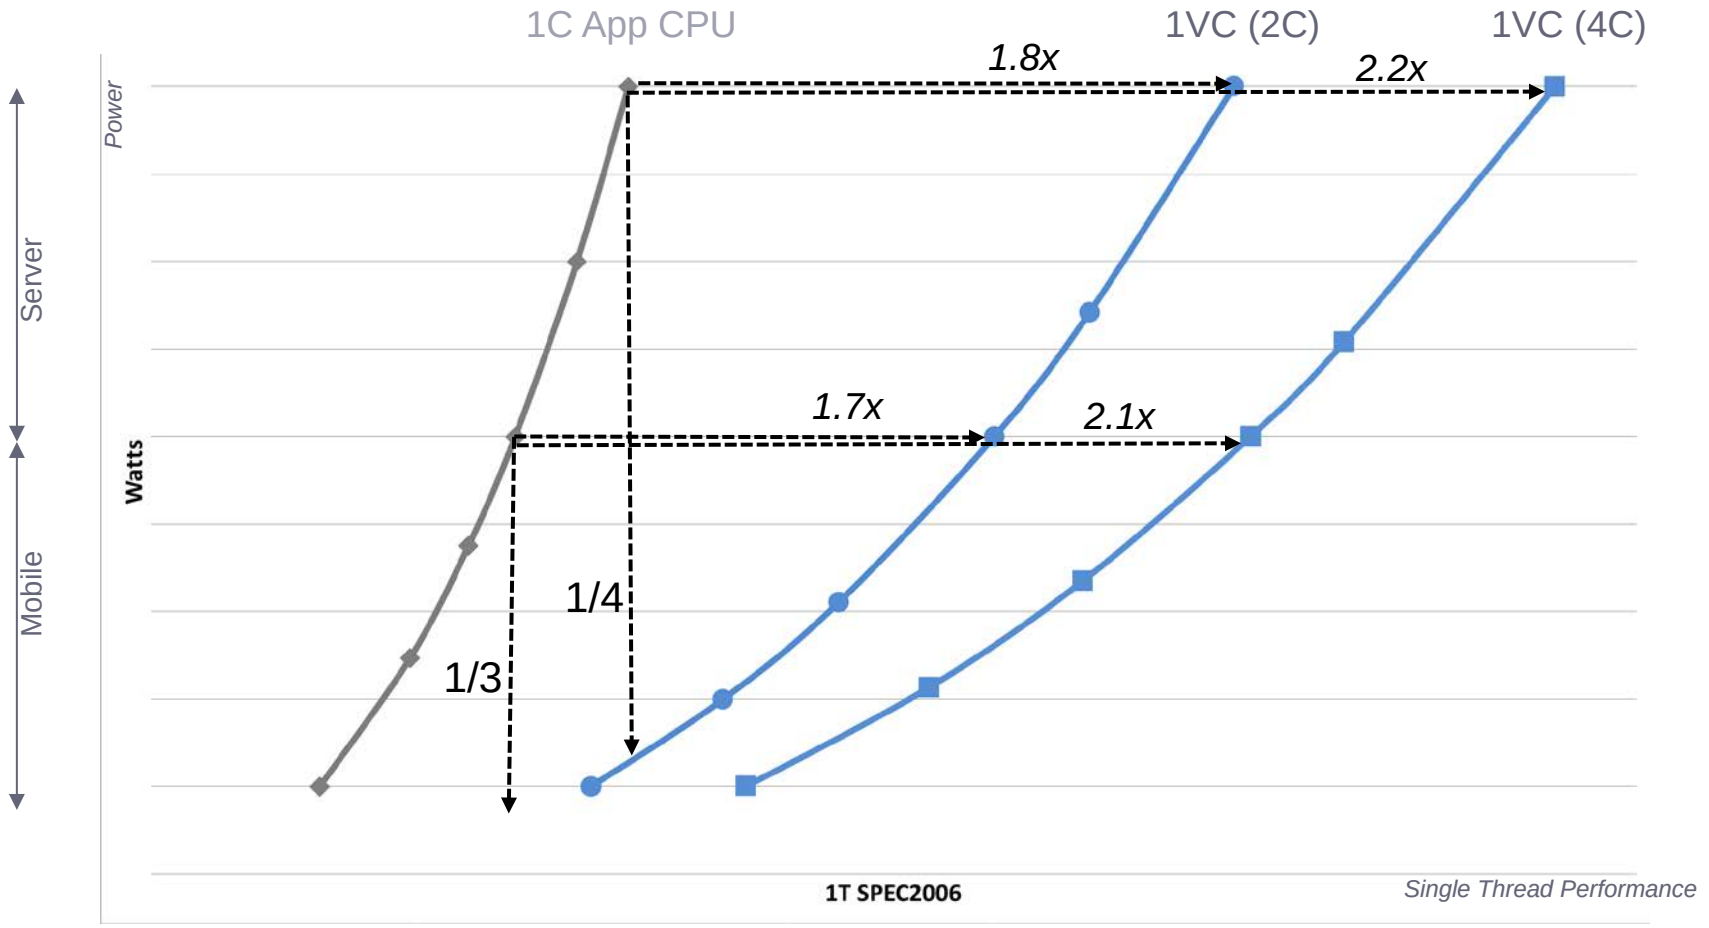

Same performance in 1/4-1/3rd power or 1.7-2.2x perf at the same power*

* Company conducted benchmark tests and projections for 28nm

VISC™ Dual Thread SPEC/Watt

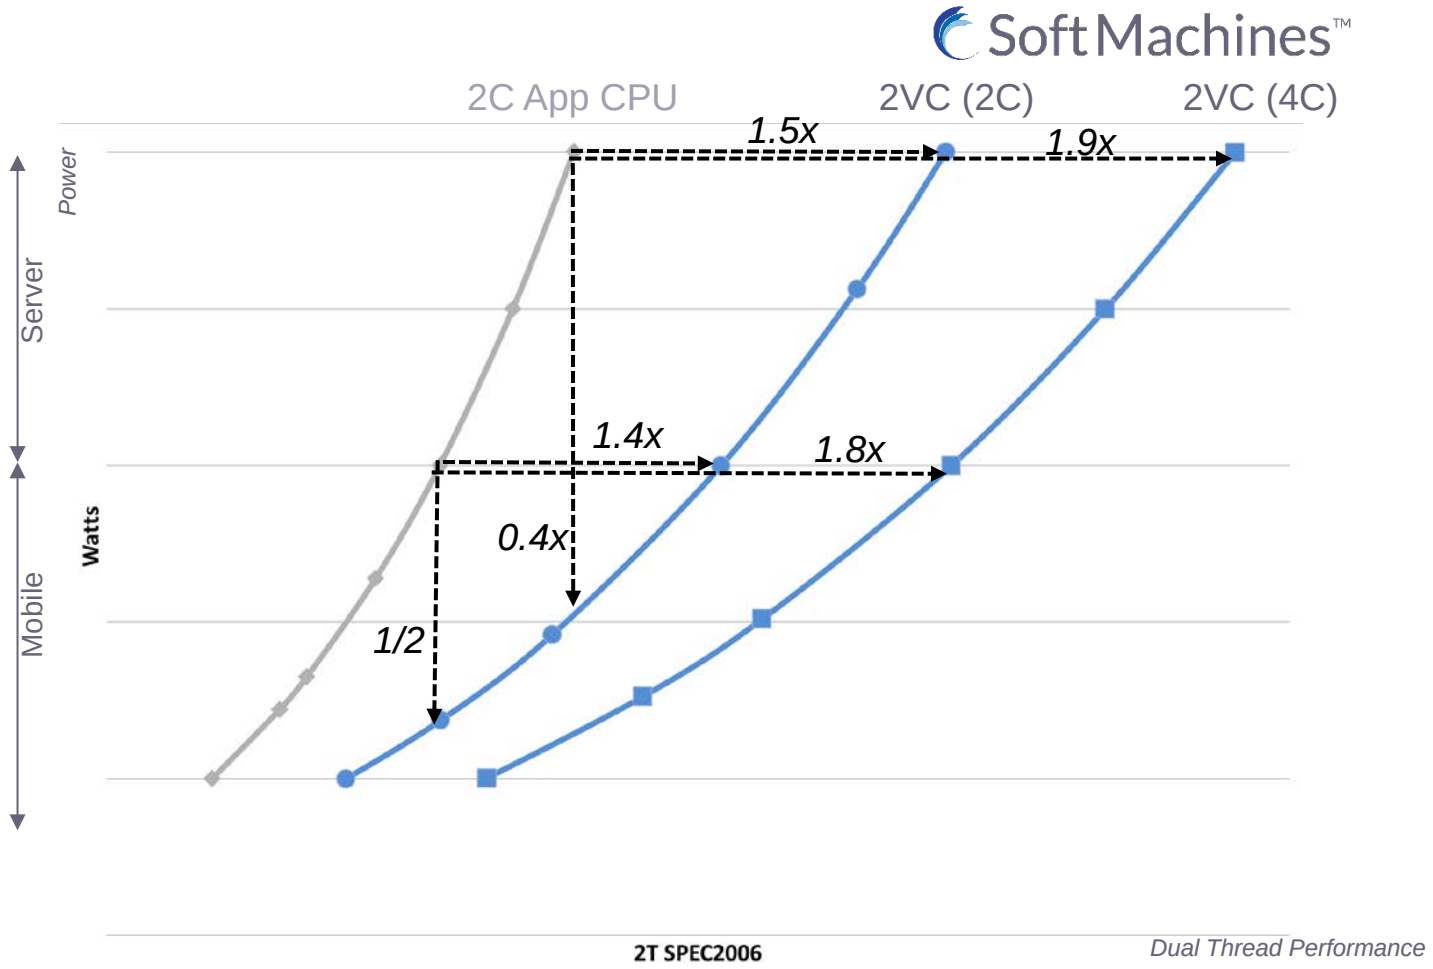

Same performance in 0.4 to 0.5x of power or 1.4 - 1.9x perf at the same power*

* Company conducted benchmark tests and projections for 28nm

VISC™ Technology Prototype

- VISC Processor Proof-of-Concept Prototype
  - IPC scalability
  - VISC architecture
  - Software efficiency
- Full Platform
  - VISC Dual Virtual Core Processor
  - SoC with 3D, Video, DRAM controller, HD video….
- Full System functionality
  - Linux OS
  - UEFI BIOS
  - Benchmarks running on Linux
  - Android ICS booting

Working Silicon

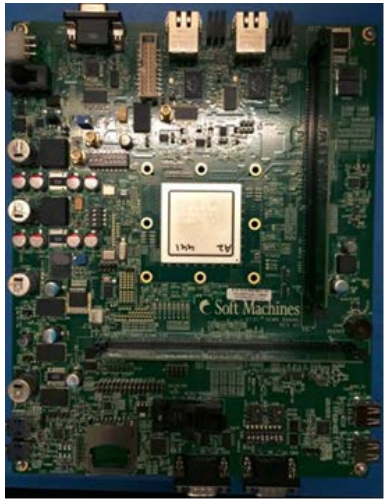

Silicon Results: Performance/MHz

Dual Virtual Core/A15 IPC Ratio

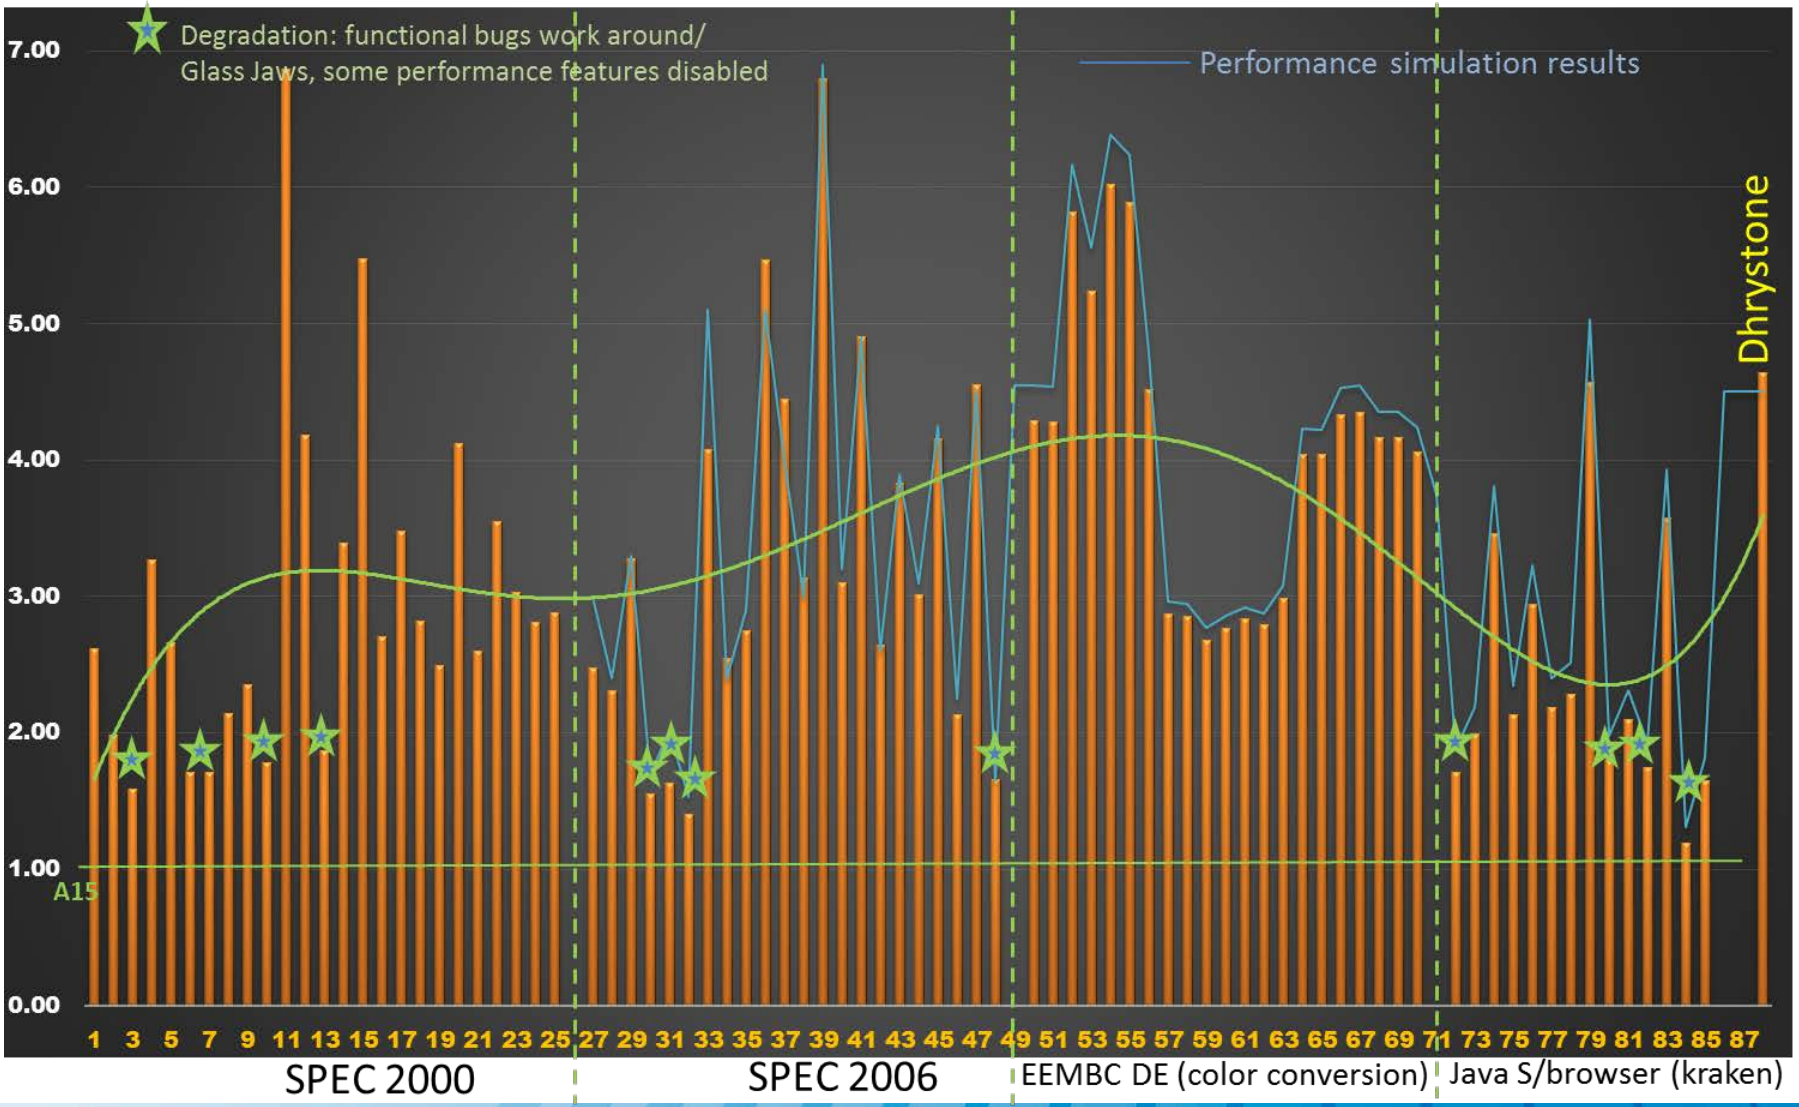

VISC™ Architecture

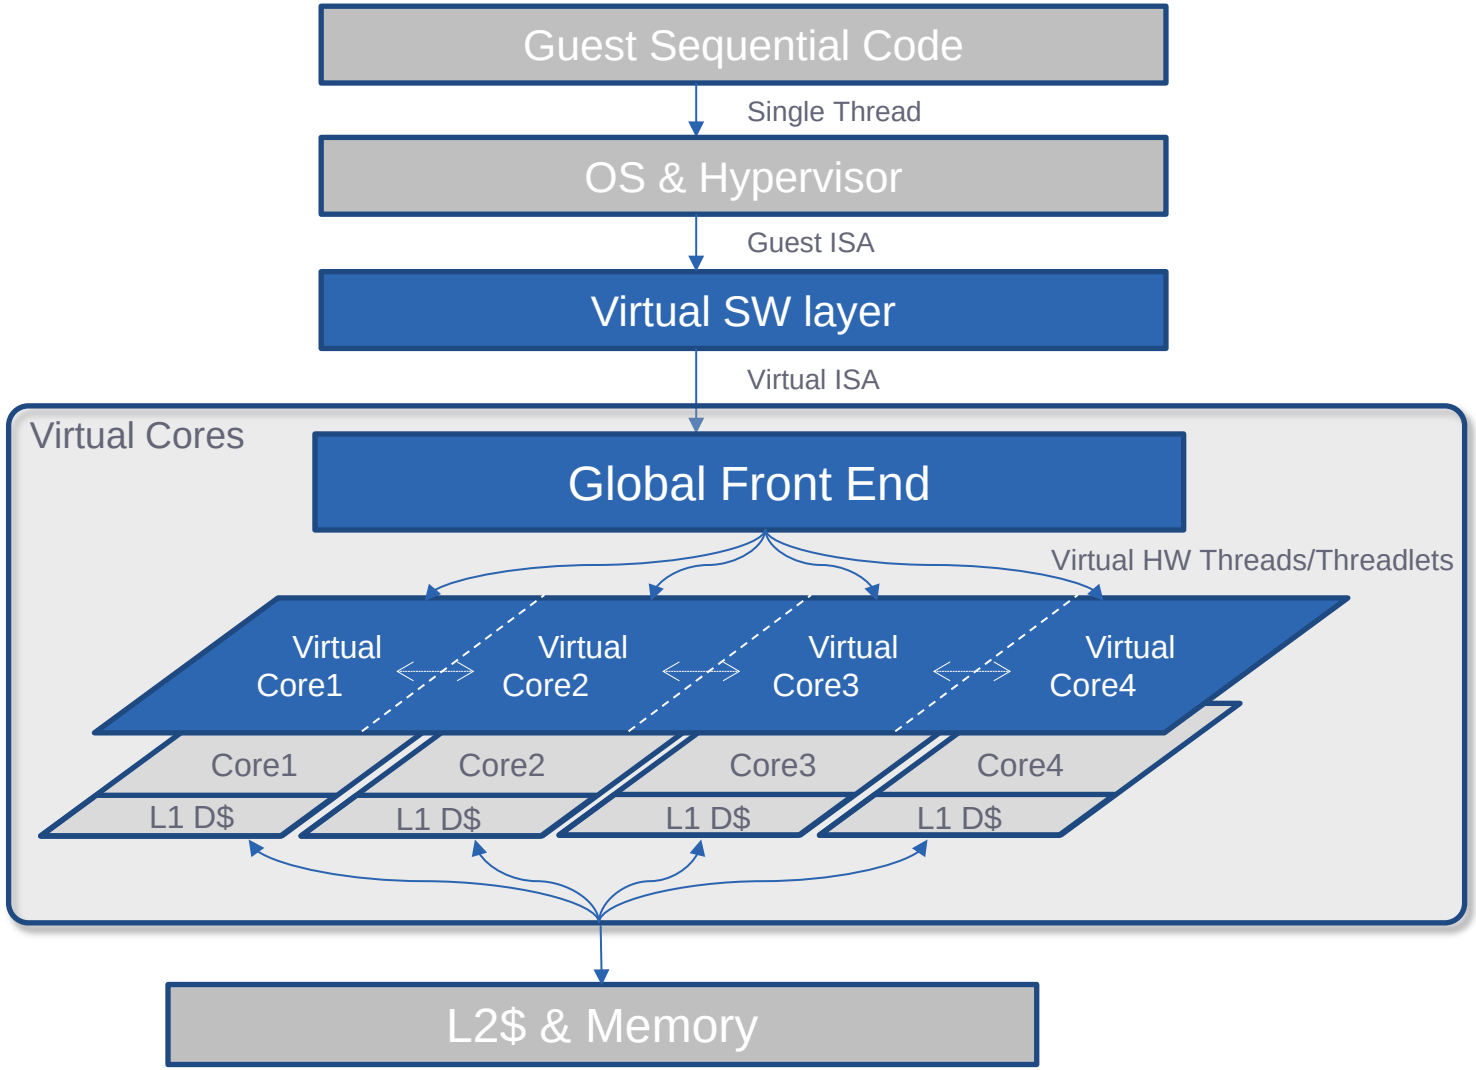

VISC™ Run-time SW Architecture

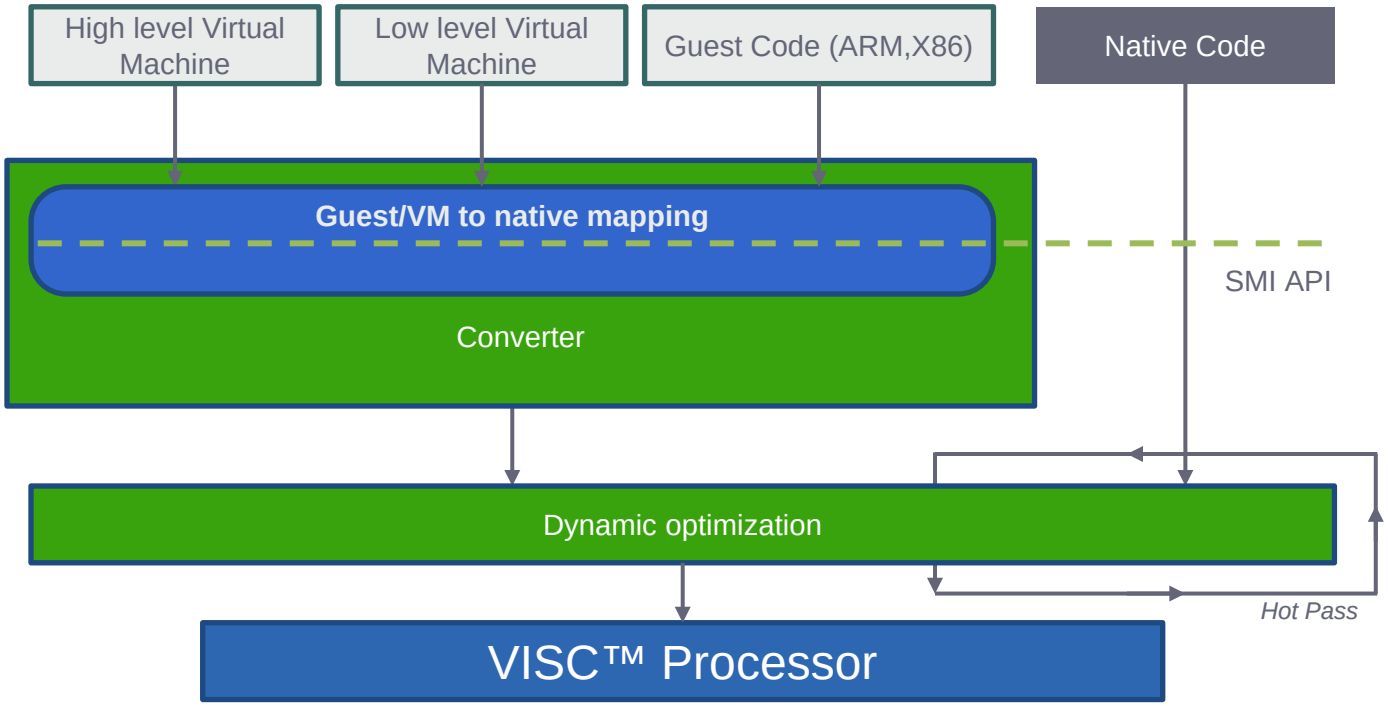

Summary

- Silicon proven VISC™ architecture delivers 3-4x IPC advantage on single and multi-threaded applications without software changes
  - Resulting in ~2-4x performance/watt advantage
- VISC architecture is scalable from IoT to mobile to servers due to its modularity and symmetry
  - Number of virtual cores, virtual threads, and virtual instruction layer
- VISC virtual instruction layer provides ISA agnostic and optimized run-time platform capabilities