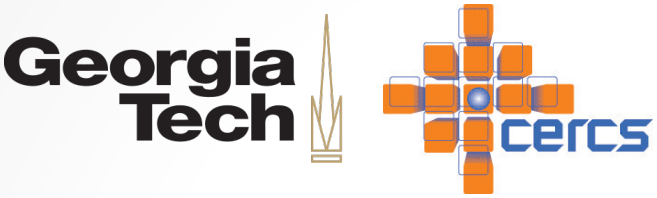

The Forgotten 'Uncore': On the Energy-Efficiency of Heterogeneous Cores

USENIX ATC 2012

(Boston, MA)

Vishal Gupta*, Karsten Schwan (College of Computing, Georgia Tech)

Paul Brett, Scott Hahn, David Koufaty, Dheeraj Reddy (Intel Labs)

Ganapati Srinivasa (Intel Corporation)

Mobile Devices

Battery life critical for usability; end-user demands for high performance as well (e.g., gaming)

Heterogeneous Cores

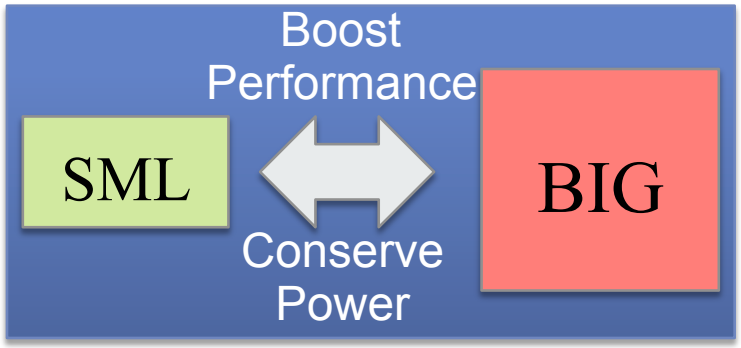

Use low-powered small (SML) and high-performance big (BIG) cores together to improve performance/ energy-efficiency

Why heterogeneity?

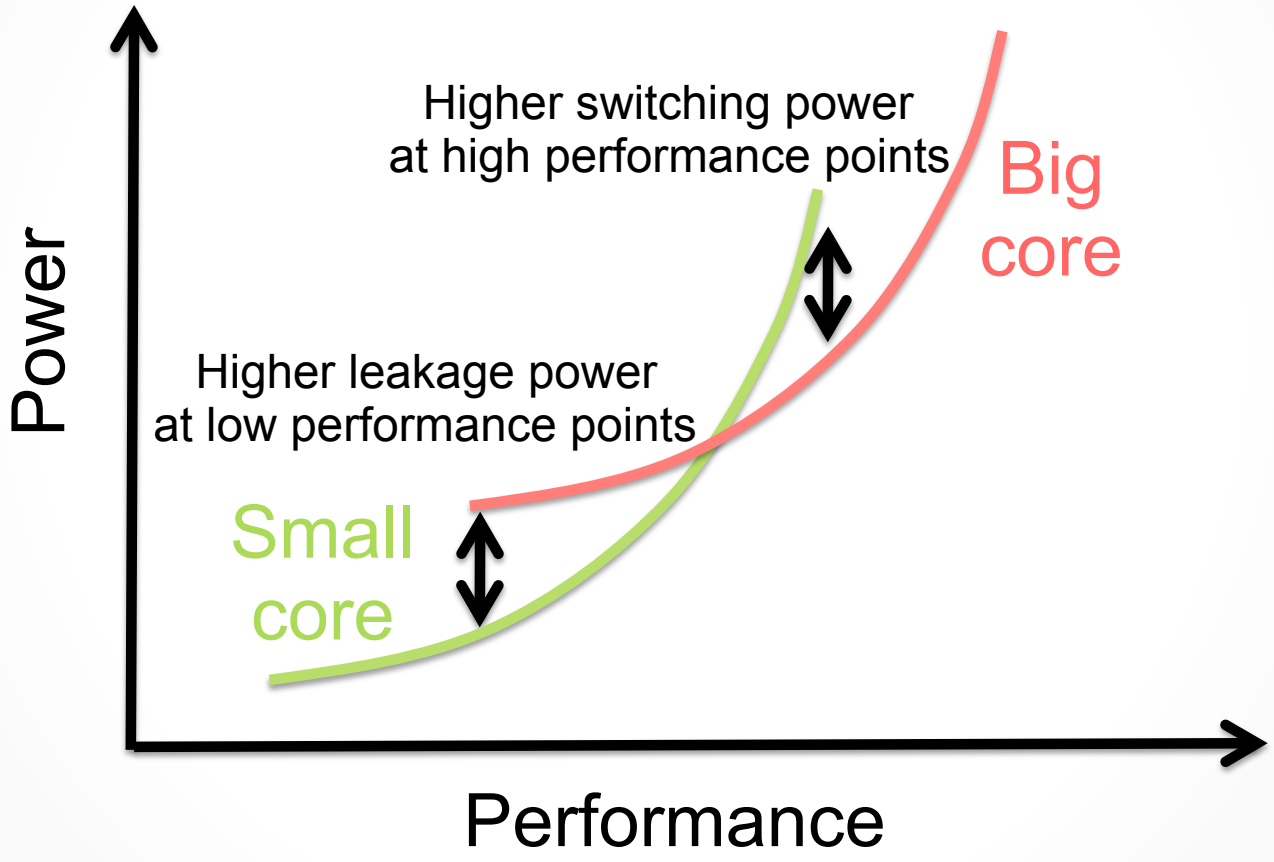

Lots of previous work…

[1] R.Kumar et al. Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction. In Proceedings of the 36th *MICRO* 2003.

[2] P. Greenhalgh. Big.LITTLE processing with ARM CortexTM-A15 & Cortex-A7. White paper, ARM, Sept 2011.

[3] D. Andersen et al. FAWN: a fast array of wimpy nodes. In *Proceedings of the SOSP* 2009.

[4] Variable SMP: A multi-core CPU architecture for low power and high performance. White paper, NVIDIA 2011.

[5] A. Fedorova et al. Maximizing power efficiency with asymmetric multicore systems. *Commun. ACM* Dec. 2009.

Beyond Core: Uncore

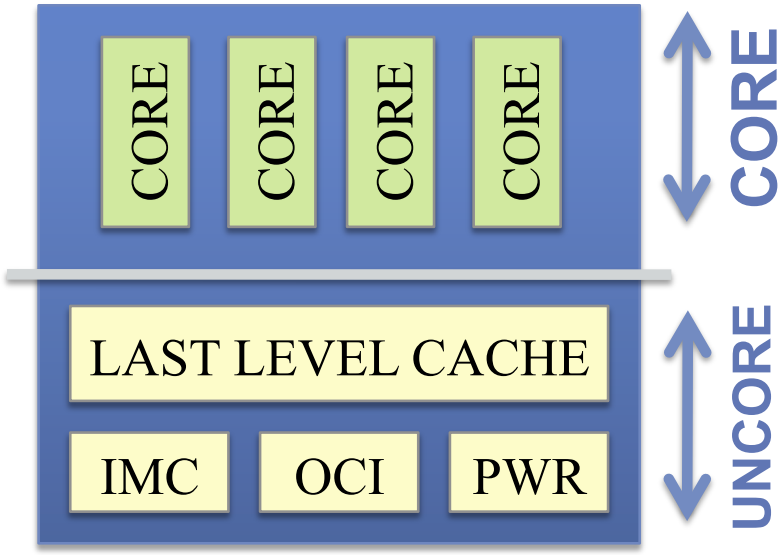

"Uncore" subsystem becoming a significant contributor in total SoC power consumption

Core/Uncore Idle-state Coordination

Core and uncore idle (sleep) states to save power

Uncore		Core 1	
		Active	Idle
Core 0	Active	Active	Active
	Idle	Active	Idle

Uncore remains active as long as any core is active

Impact of Uncore

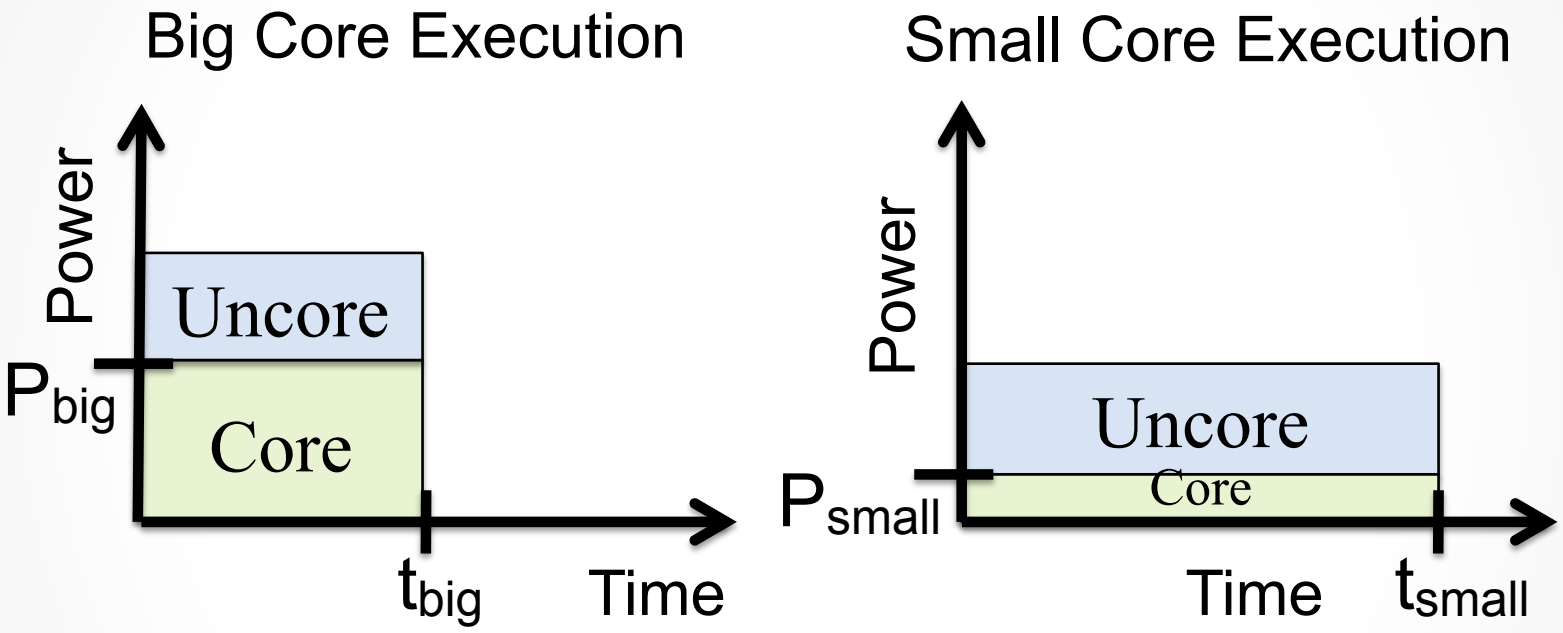

Small cores keep uncore active for longer time, increasing the uncore power component

Uncore Scaling

Impact of uncore analyzed using two configurations:

a) **Fixed uncore**: use the same fixed uncore during both big and small core execution

b) **Scalable uncore**: use a low-powered uncore to save power when execution switches to the small core

Experimental Heterogeneous Platform

Intel quadcore Sandybridge Client Processor
Proprietary tools to create heterogeneity
Shared uncore

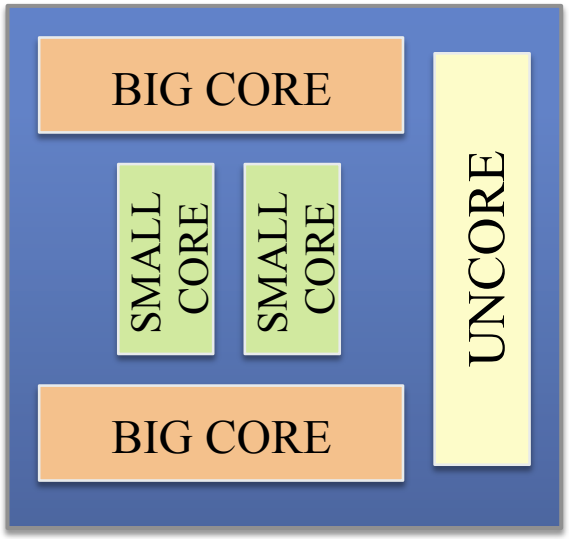

Power Models

- **Core Power**: uses IPC to model capacitance, also incorporates core idle state residency
- **Uncore power**: LLC accesses to estimate uncore activity along with package idle state residency
- Uncore power scaled to half for a scalable uncore

Client Workloads

Workload	Description	Metrics
browse	Web-page rendering	Load time
javascript	Javascript operations	Load time
palbum	Photo album viewer	Load time
mplayer	Media playback	FPS
mytube	Web-based video playback	FPS
openarena	3D shooting game	FPS
strike	Web-based 2D game	FPS
7zip	File archiver	Time
eclipse	IDE environment	Time
filesan	File system scan	Time
gmagick	Image editor	Time
x264	Media encoder	Time

Client Workloads: browse

Loads a series of web-pages from a web-server at 3s interval to emulate user's think time (metric: load-time)

1) Amazon

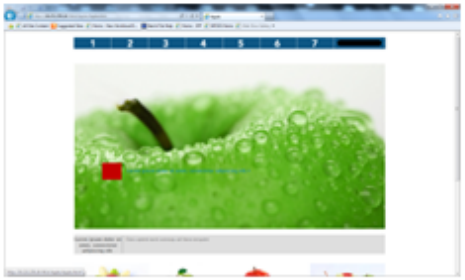

2) Apple

3) Atom

4) Baidu

5) CNN

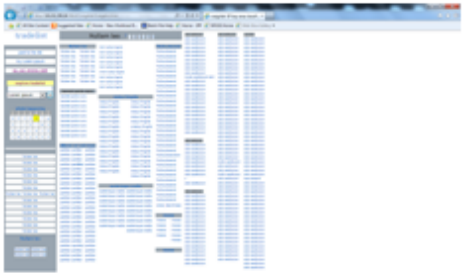

6) Craigslist

Client Workloads: openarena

Plays a benchmarking demo (120s) from a 3D first- person-shooter game (metric: FPS)

Experimental Results

Performance Impact (Big vs. Small)

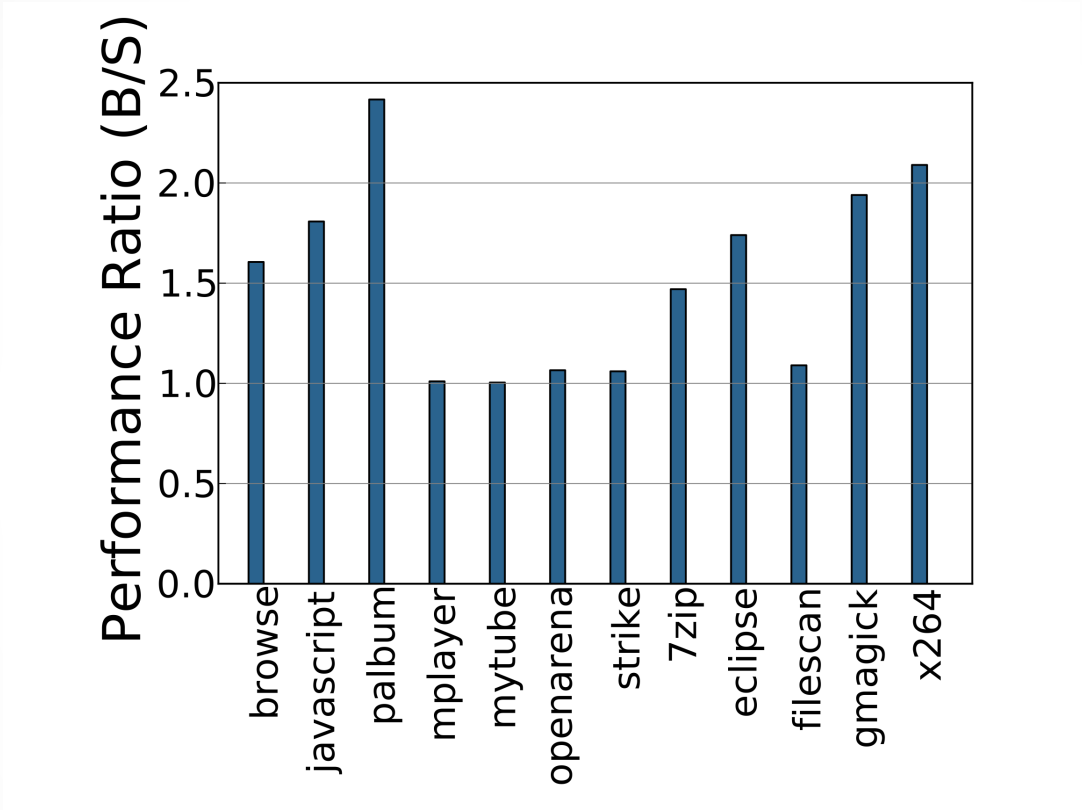

Big core improves performance for several applications.

Small core provides comparable performance for graphics/media and I/O intensive applications.

Energy Savings (Small vs. Big)

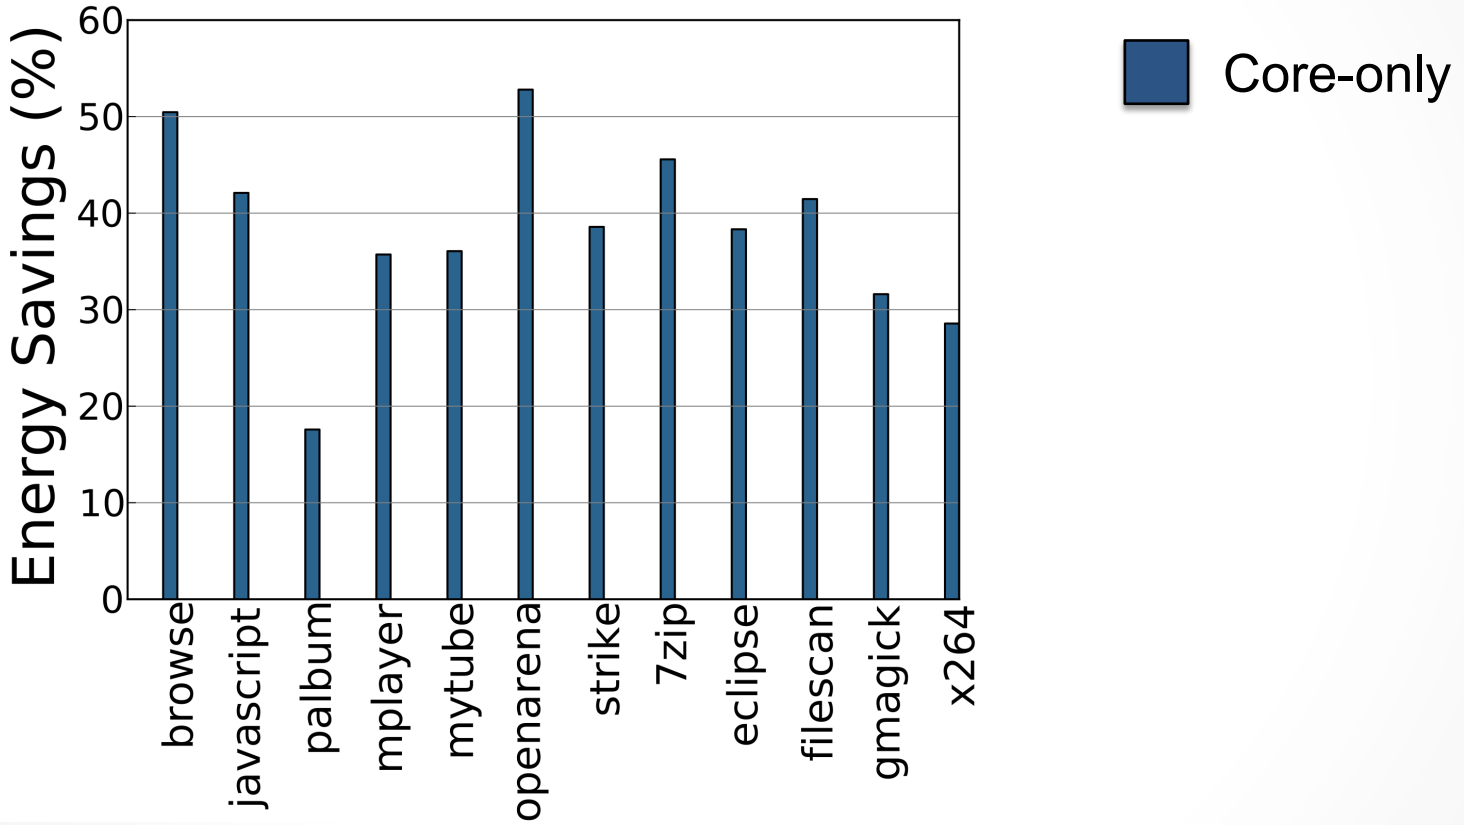

Significant core-only savings for all of the applications.

Energy Savings (Small vs. Big)

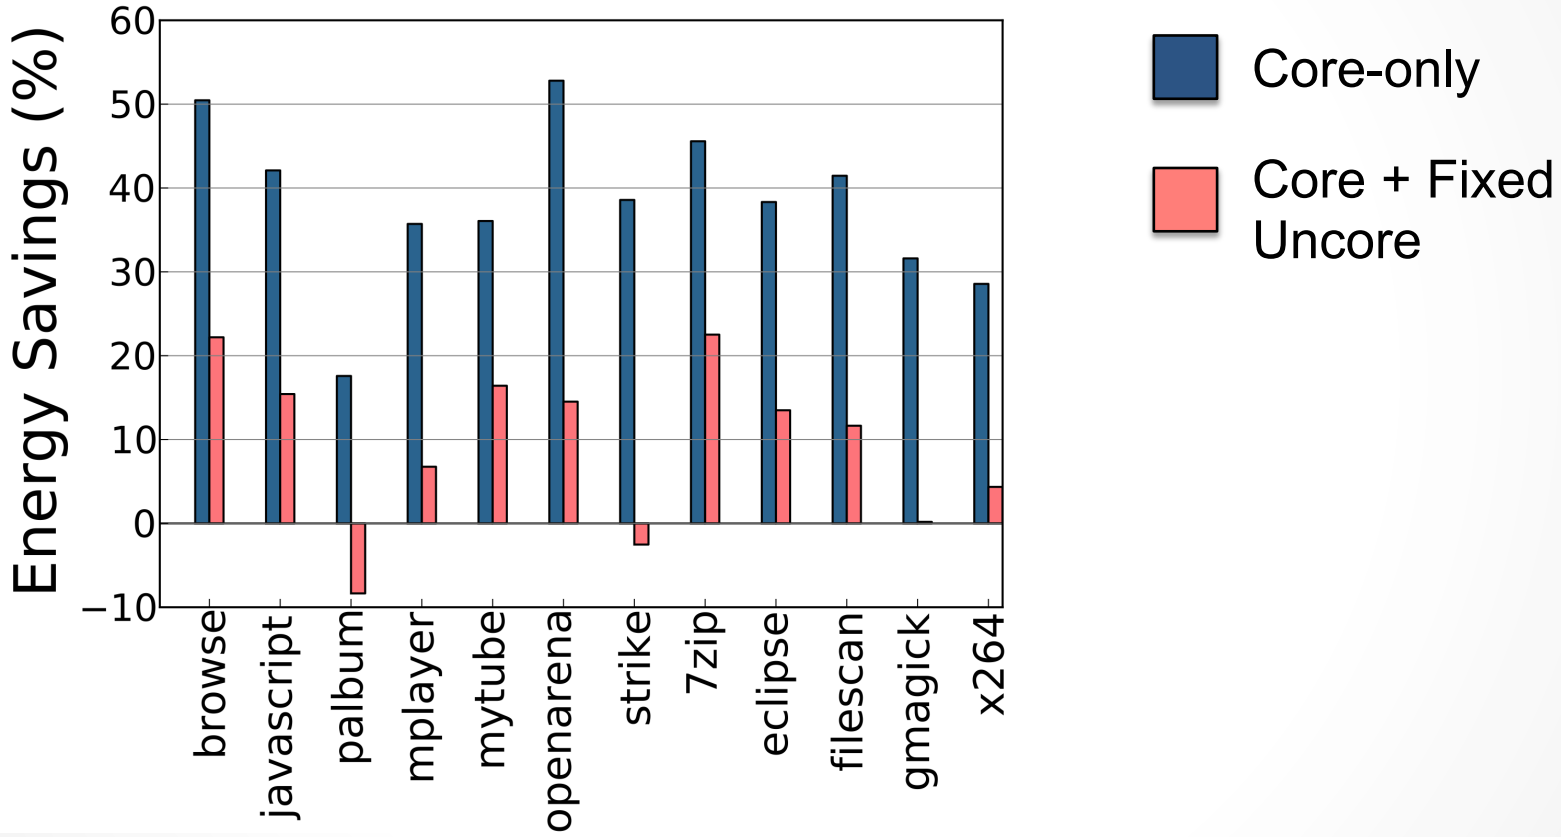

Savings strongly affected when uncore contribution is included in addition to core energy.

Energy Savings (Small vs. Big)

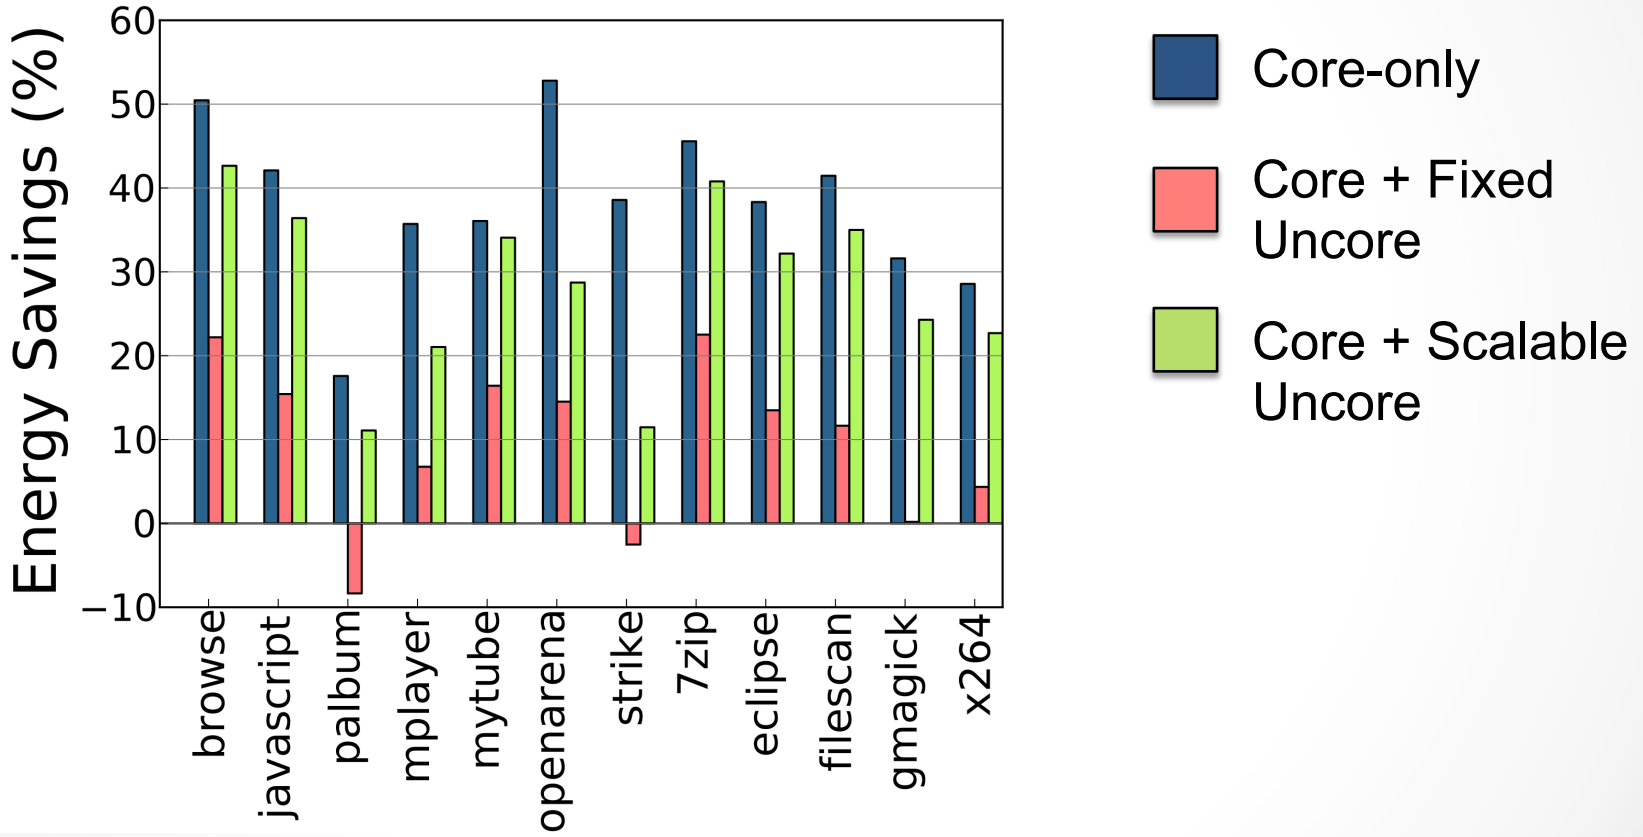

A scalable uncore configuration provides savings comparable to core-only savings

Energy Distribution

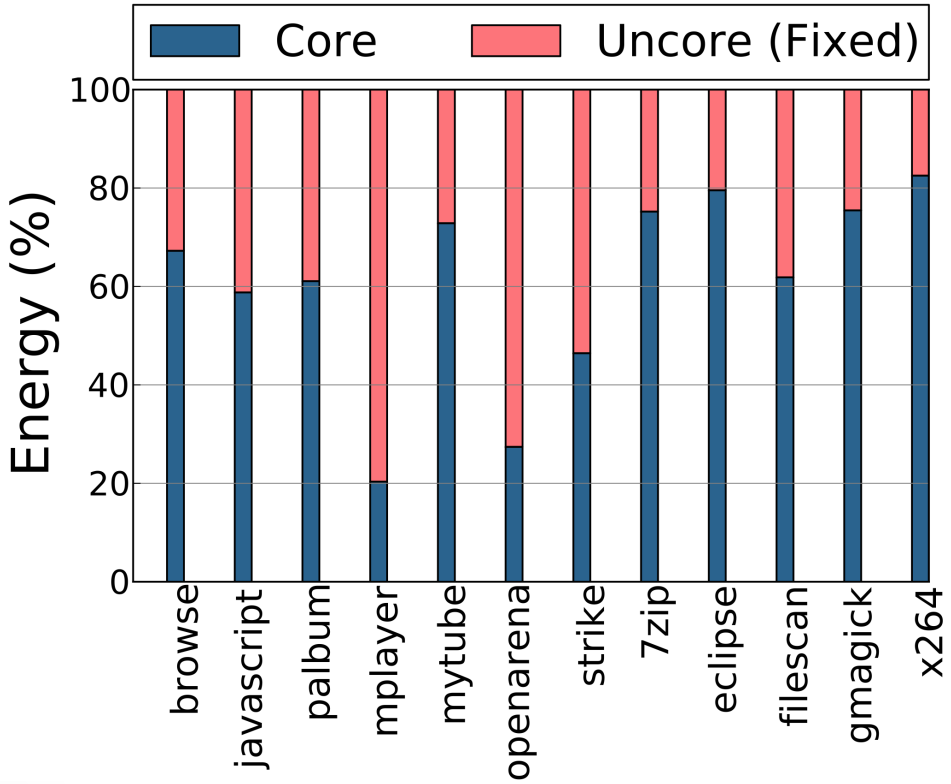

Varying dominant energy component across applications

Summary

- Heterogeneous cores for improving the energy-efficiency of mobile devices
- Evaluation using representative client workloads and experimental heterogeneous platform
- 'uncore-awareness' is a MUST. Need for a scalable uncore design to fully realize the intended gains

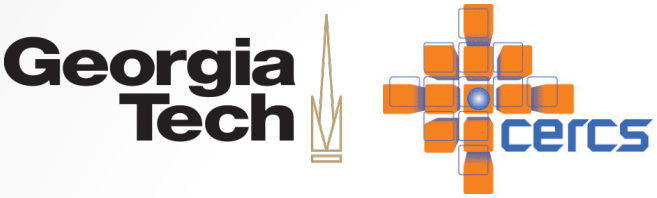

The Forgotten 'Uncore': On the Energy-Efficiency of Heterogeneous Cores

vishal@cc.gatech.edu

Thank you.