ICE: A General and Validated Energy Complexity Model for Multithreaded Algorithms

Vi N. N. Tran, *Phuong H. Ha*

Department of Computer Science

UiT The Arctic University of Norway

ICPADS 2016, Wuhan, Dec. 13-16, 2016

Energy ?

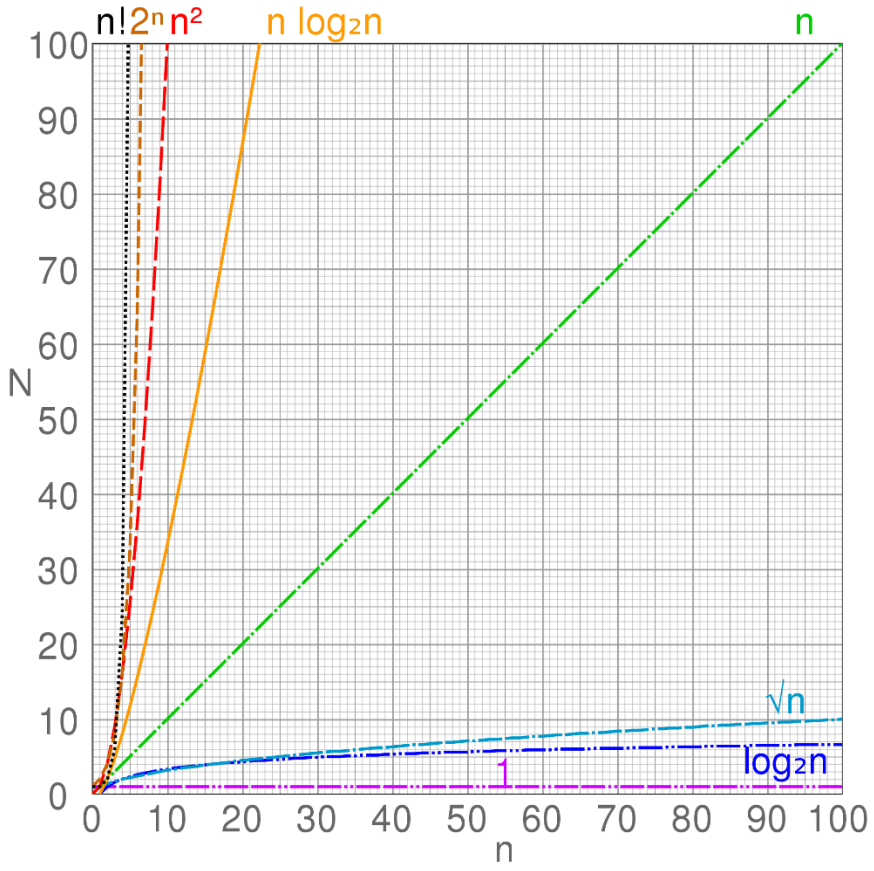

Energy has become the primary limiting factor in the development of all systems from data centers to mobile devices

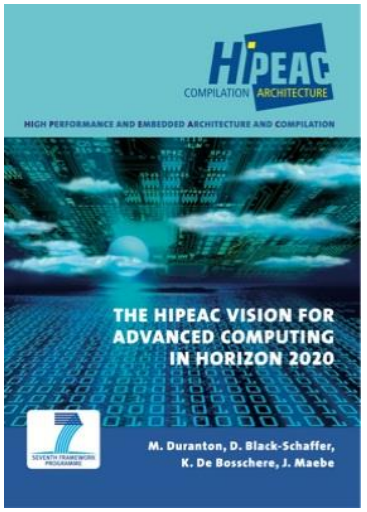

Given two parallel algorithms A and B, which algorithm consumes less energy analytically?

Study	Energy complexity for parallel algorithms	Algorithm generality	Validation
Roofline [IPDPS'13, IPDPS'14]	✗	✓	✓
Sequential energy complexity [ITCS'13]	✗	✓	✓
Koala [EuroSys'09]	✗	✓	✓
LEO [ASPLOS'15]	✗	✓	✓
POET [RTAS'15]	✗	✓	✓
Energy scalability [ICPP'09, SPAA'10]	✓	✓	✗
Dense matrix factorization [CCPE'14]	✓	✗	✓
Linear algebra [SusCom'15]	✓	✗	✓
ICE [This study]	✓	✓	✓

Challenges: parallel algorithm complexity + parallel platform complexity

This presentation focuses on a new general model ICE for analyzing the energy complexity of multithread algorithms

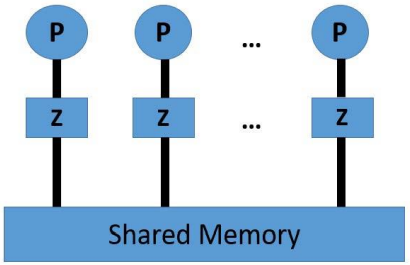

Platform abstractions

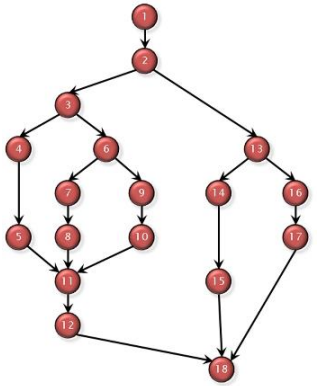

http://www.cprogramming.com/

Algorithm complexities

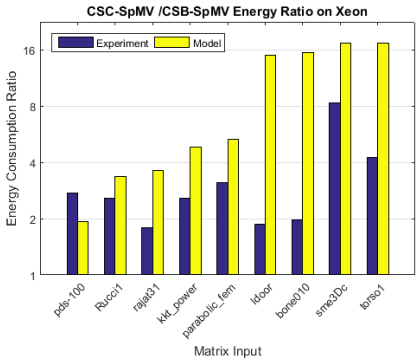

Validation

The ICE model abstracts away multicore platforms to keep itself simple and suitable for complexity analysis

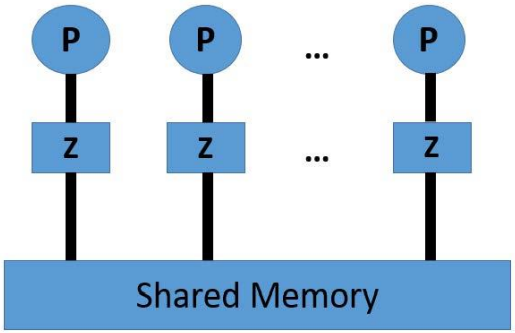

Shared memory with private caches

	Op	I/O
Static energy (max)	π_{op}	$\pi_{I/O}$
Dynamic energy (max)	ε_{op}	$\varepsilon_{I/O}$

Platform parameters, assuming the race-to-halt strategy

Platform	Processor	ϵ_{op}(nJ)	π_{op}(nJ)	$\epsilon_{I/O}$(nJ)	$\pi_{I/O}$(nJ)
Nehalem i7-950	Intel i7-950	0.670	2.455	50.88	408.80
Ivy Bridge i3-3217U	Intel i3-3217U	0.024	0.591	26.75	58.99
Bobcat CPU	AMD E2-1800	0.199	3.980	27.84	387.47
Fermi GTX 580	NVIDIA GF100	0.213	0.622	32.83	45.66
Kepler GTX 680	NVIDIA GK104	0.263	0.452	27.97	26.90
Kepler GTX Titan	NVIDIA GK110	0.094	0.077	17.09	32.94
XeonPhi KNC	Intel 5110P	0.012	0.178	8.70	63.65
Cortex-A9	TI OMAP 4460	0.302	1.152	51.84	174.00
Arndale Cortex-A15	Samsung Exynos 5	0.275	1.385	24.70	89.34
Xeon	2xIntel E5-2650l v3	0.263	0.108	8.86	23.29
Xeon-Phi	Intel 31S1P	0.006	0.078	25.02	64.40

Platform parameters of 11 platforms using the roofline method [IPDPS'14]

ICE energy complexity is derived from the algorithm complexities available in the literature

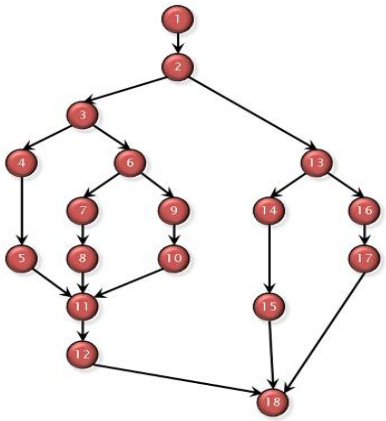

http://www.cprogramming.com/

Complexities	Description
Work	The number of ops
Span	Critical path length
I/O	# cache line transfers [FOCS'99]

$$E = E_{dynamic} + E_{static}$$
$$= \varepsilon_{op} * Work + \varepsilon_{I/O} * I/O + \mathbf{E_{static}}$$

$$E_{static} = \pi_{op} * Span$$

Compute-intensive algorithms

$$E_{static} = \pi_{I/O} * I/O \,/\, Parallelism$$
$$= \pi_{I/O} * I/O * Span/Work$$

Data-intensive algorithms

ICE energy complexity is confirmed by experimental results – Data-intensive SpMV

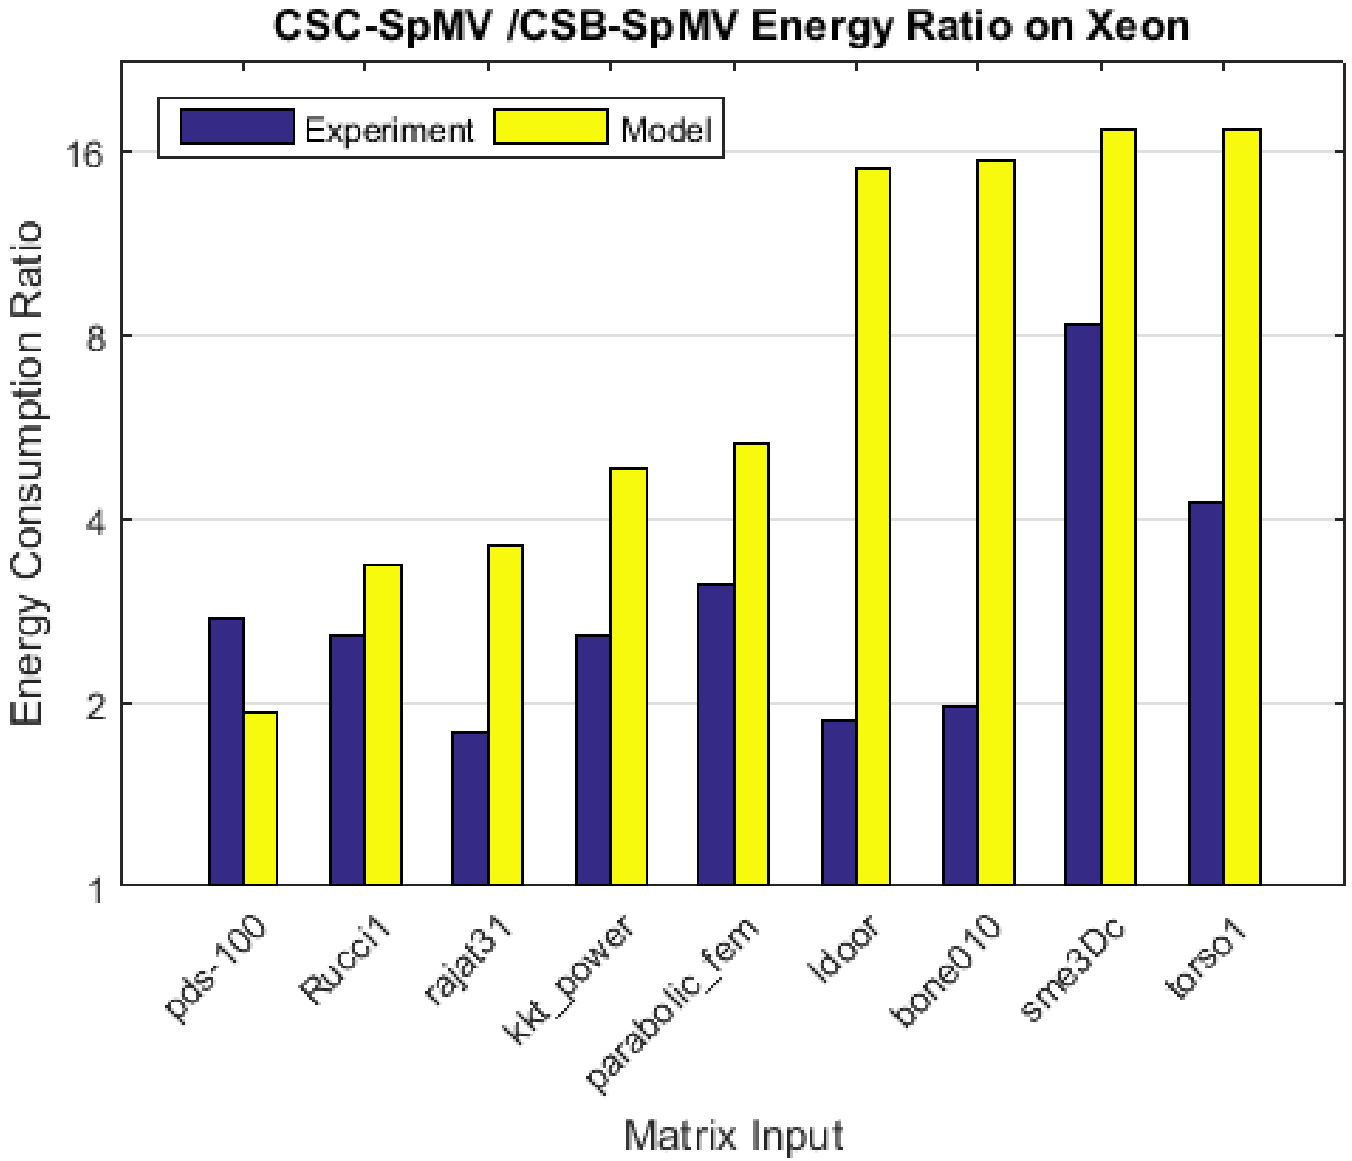

ICE energy complexity is confirmed by experimental results – Compute-intensive Matmul

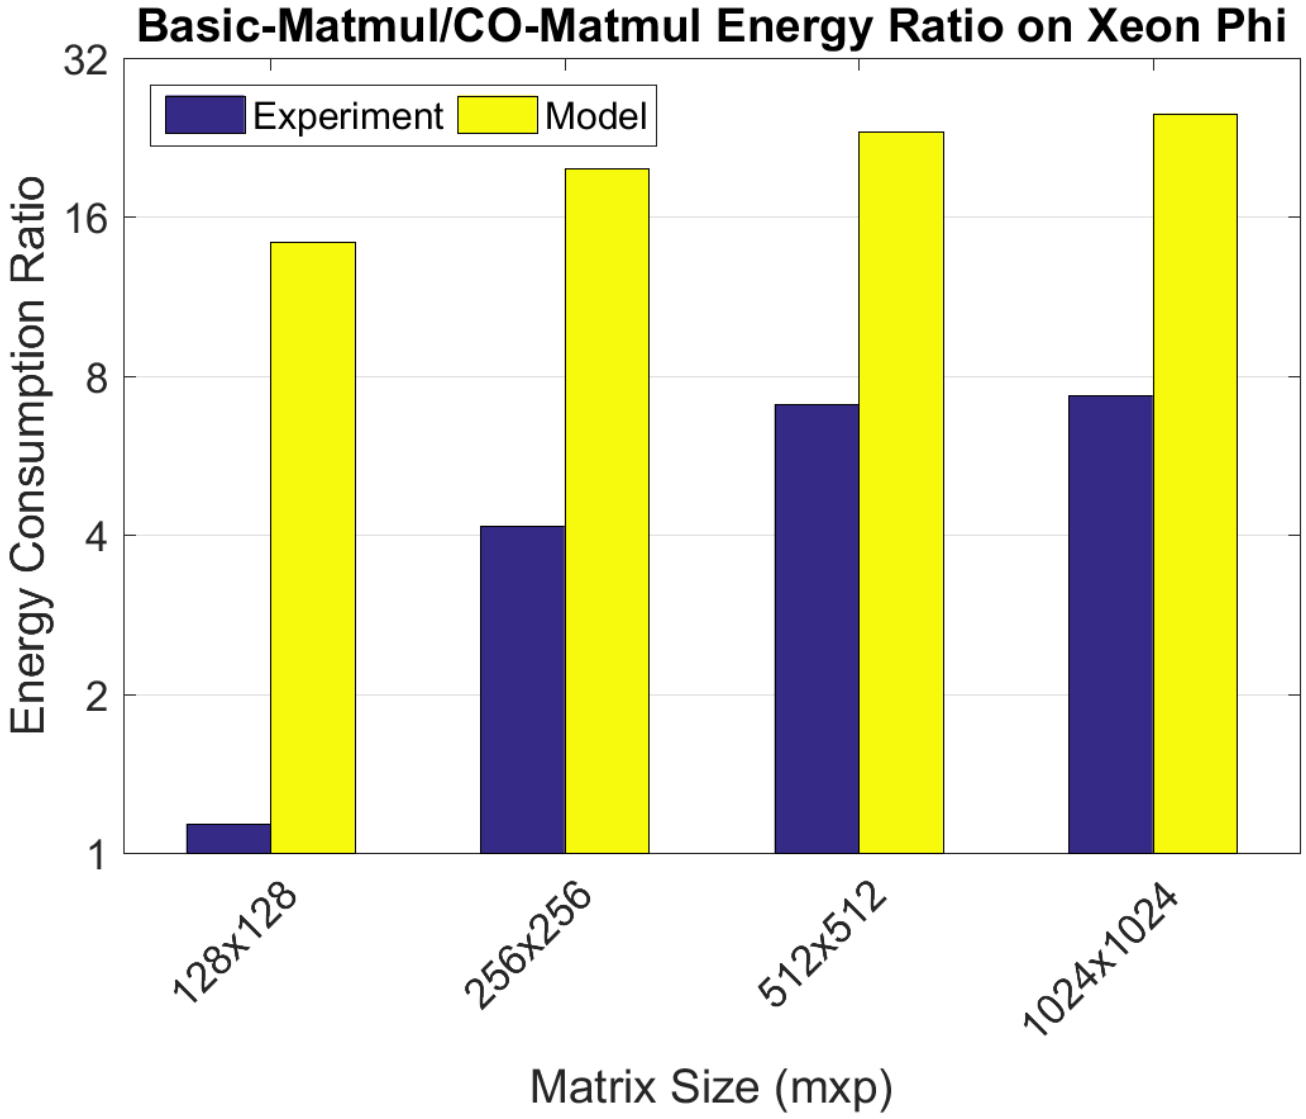

In summary, ICE is a general validated model for analyzing the energy complexity of multithreaded algorithms

It abstracts away parallel platforms by their static/dynamic energy

Energy complexity is derived from Work, Span and I/O complexities

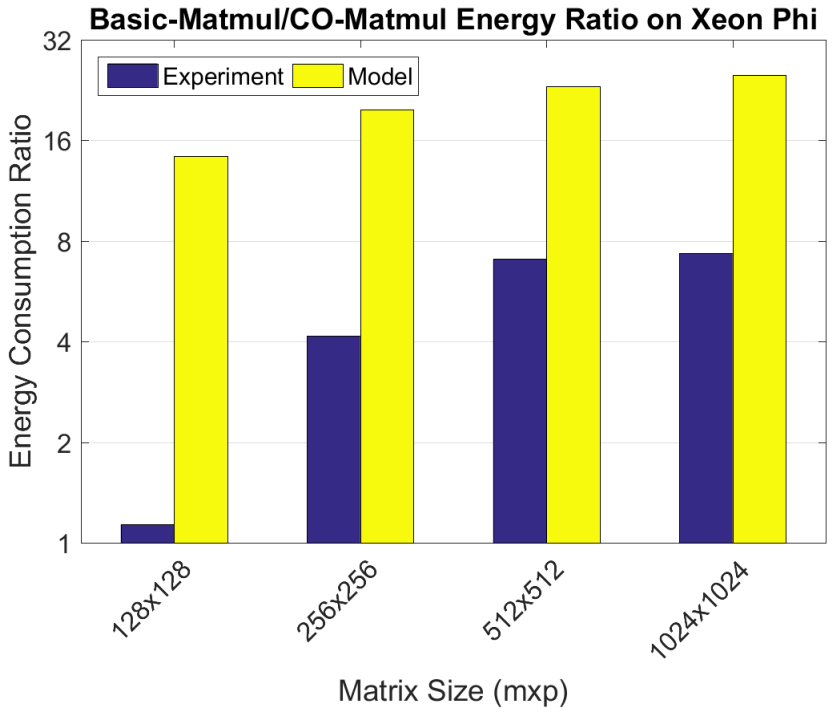

In summary, ICE is a general validated model for analyzing the energy complexity of multithreaded algorithms

It abstracts away parallel platforms by their static/dynamic energy

Energy complexity is derived from Work, Span and I/O complexities

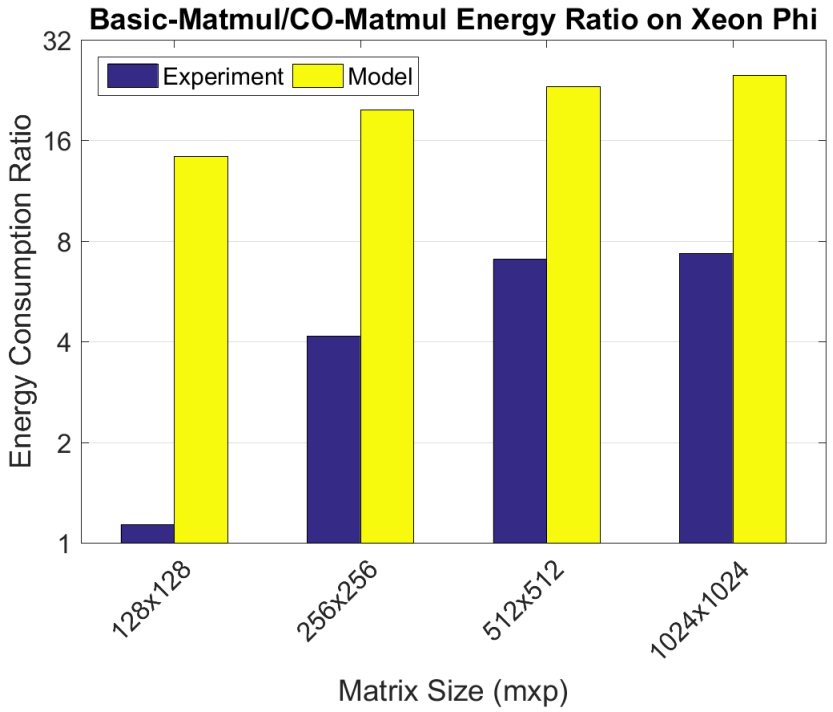

Thank you! Questions?

Contact information: Phuong Ha (phuong@cs.uit.no)