Methods and Tools for Energy Optimization

Piotr Łuszczek

PLASMA and MAGMA Software Scope

symmetric positive definite — Cholesky

symmetric, hermitian — LDL^t

non-symmetric — LU

overdetermined — QR

underdetermined

$$Ax = b$$

symmetric positive definite — symmetric eigen-solver

symmetric, hermitian

non-symmetric — general eigen-solver

overdetermined — SVD

underdetermined

$$Ax = b$$

generalized — Cholesky + symmetric eigen-solver

$$Ax = Bb$$

BSP (fork-join) vs. DAG (dynamically scheduled)

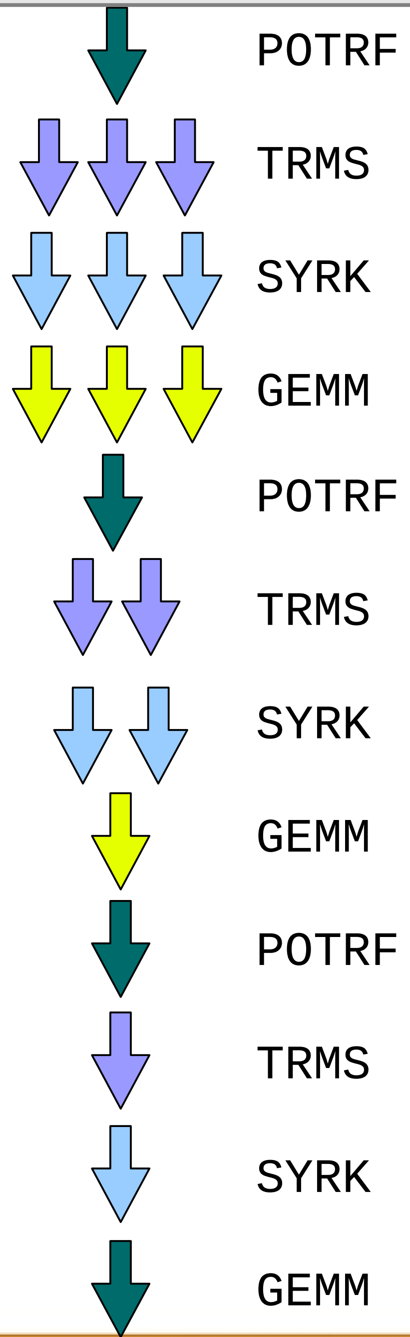

Cholesky factorization
Problem size: 4 by 4

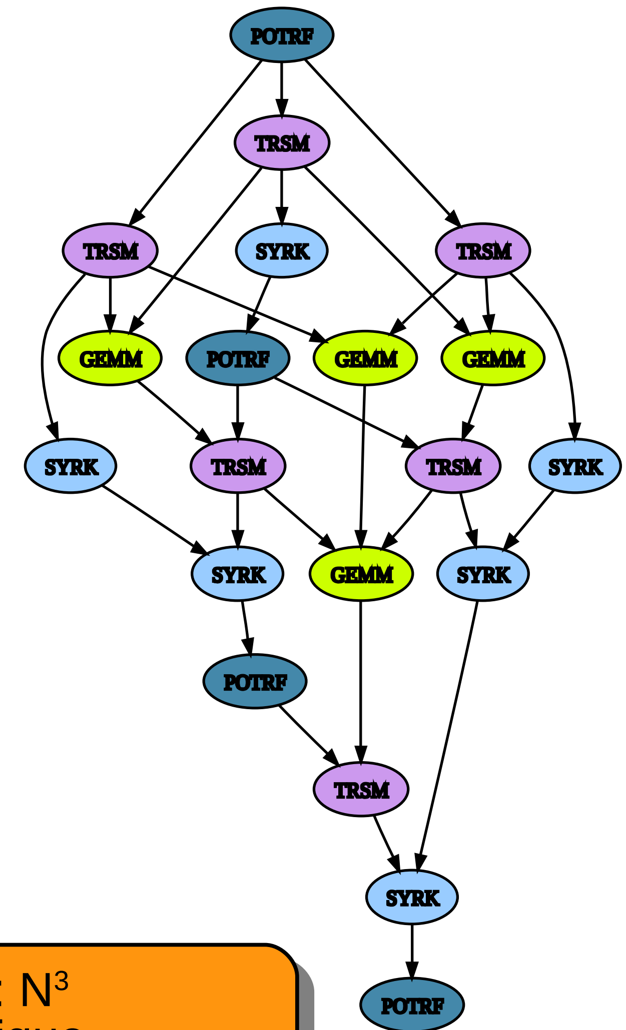

- DAGs get big quickly: N^3
- Sliding window technique
- DAGs are composed at runtime

Column-major (LAPACK) vs. Tile Storage (PLASMA)

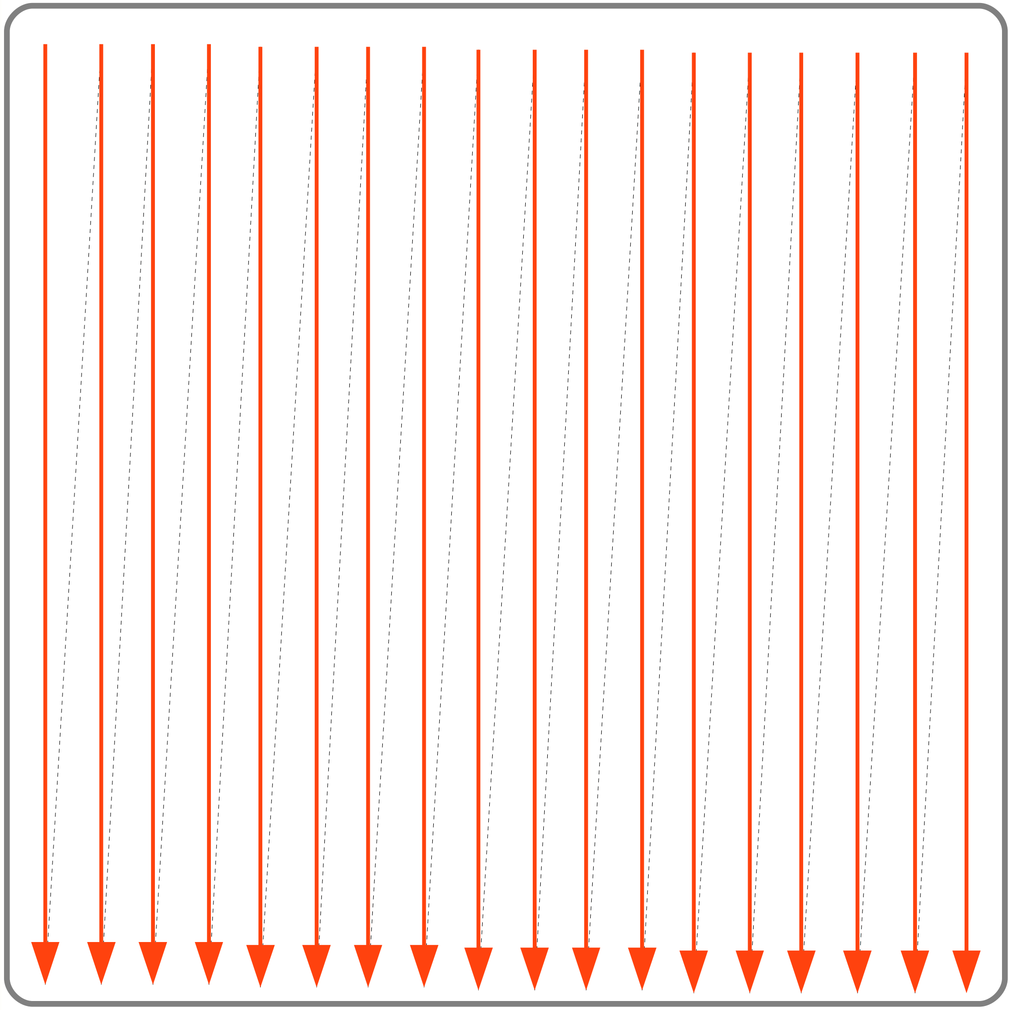

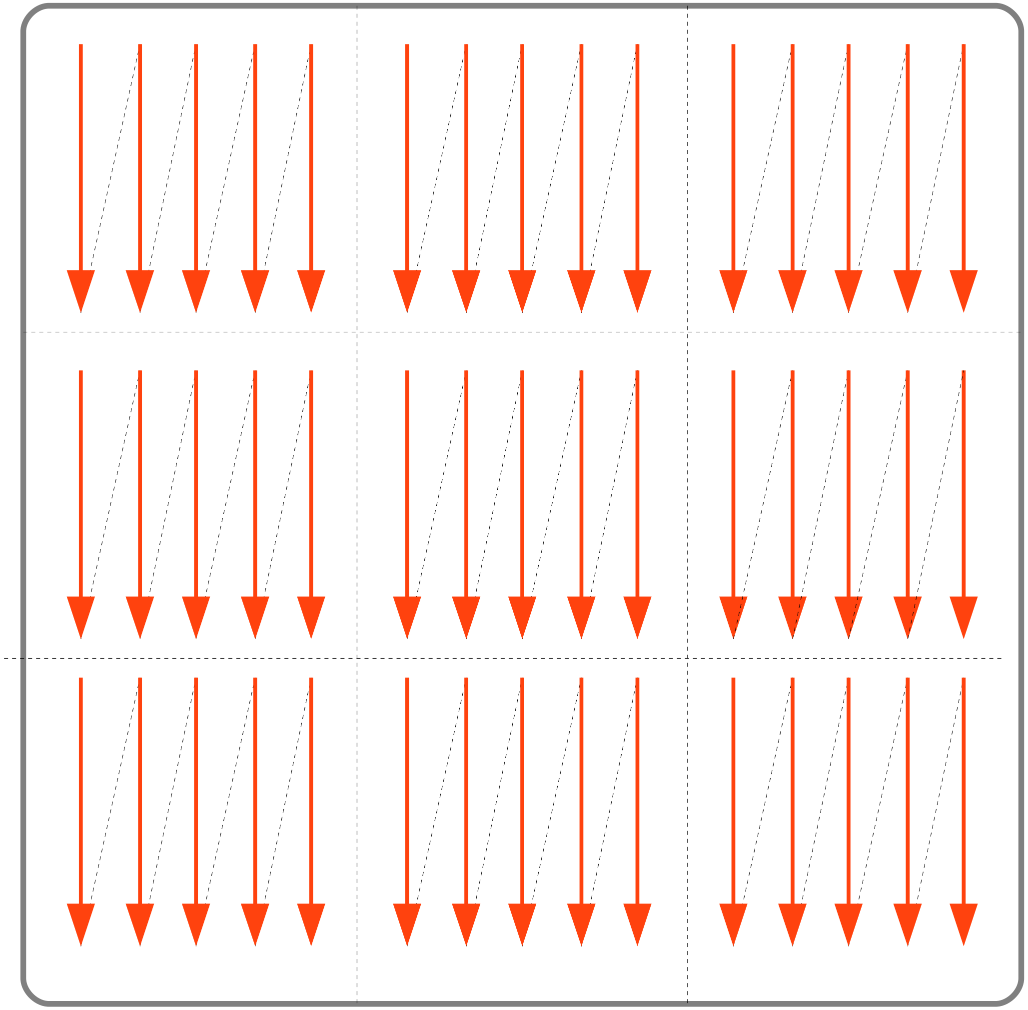

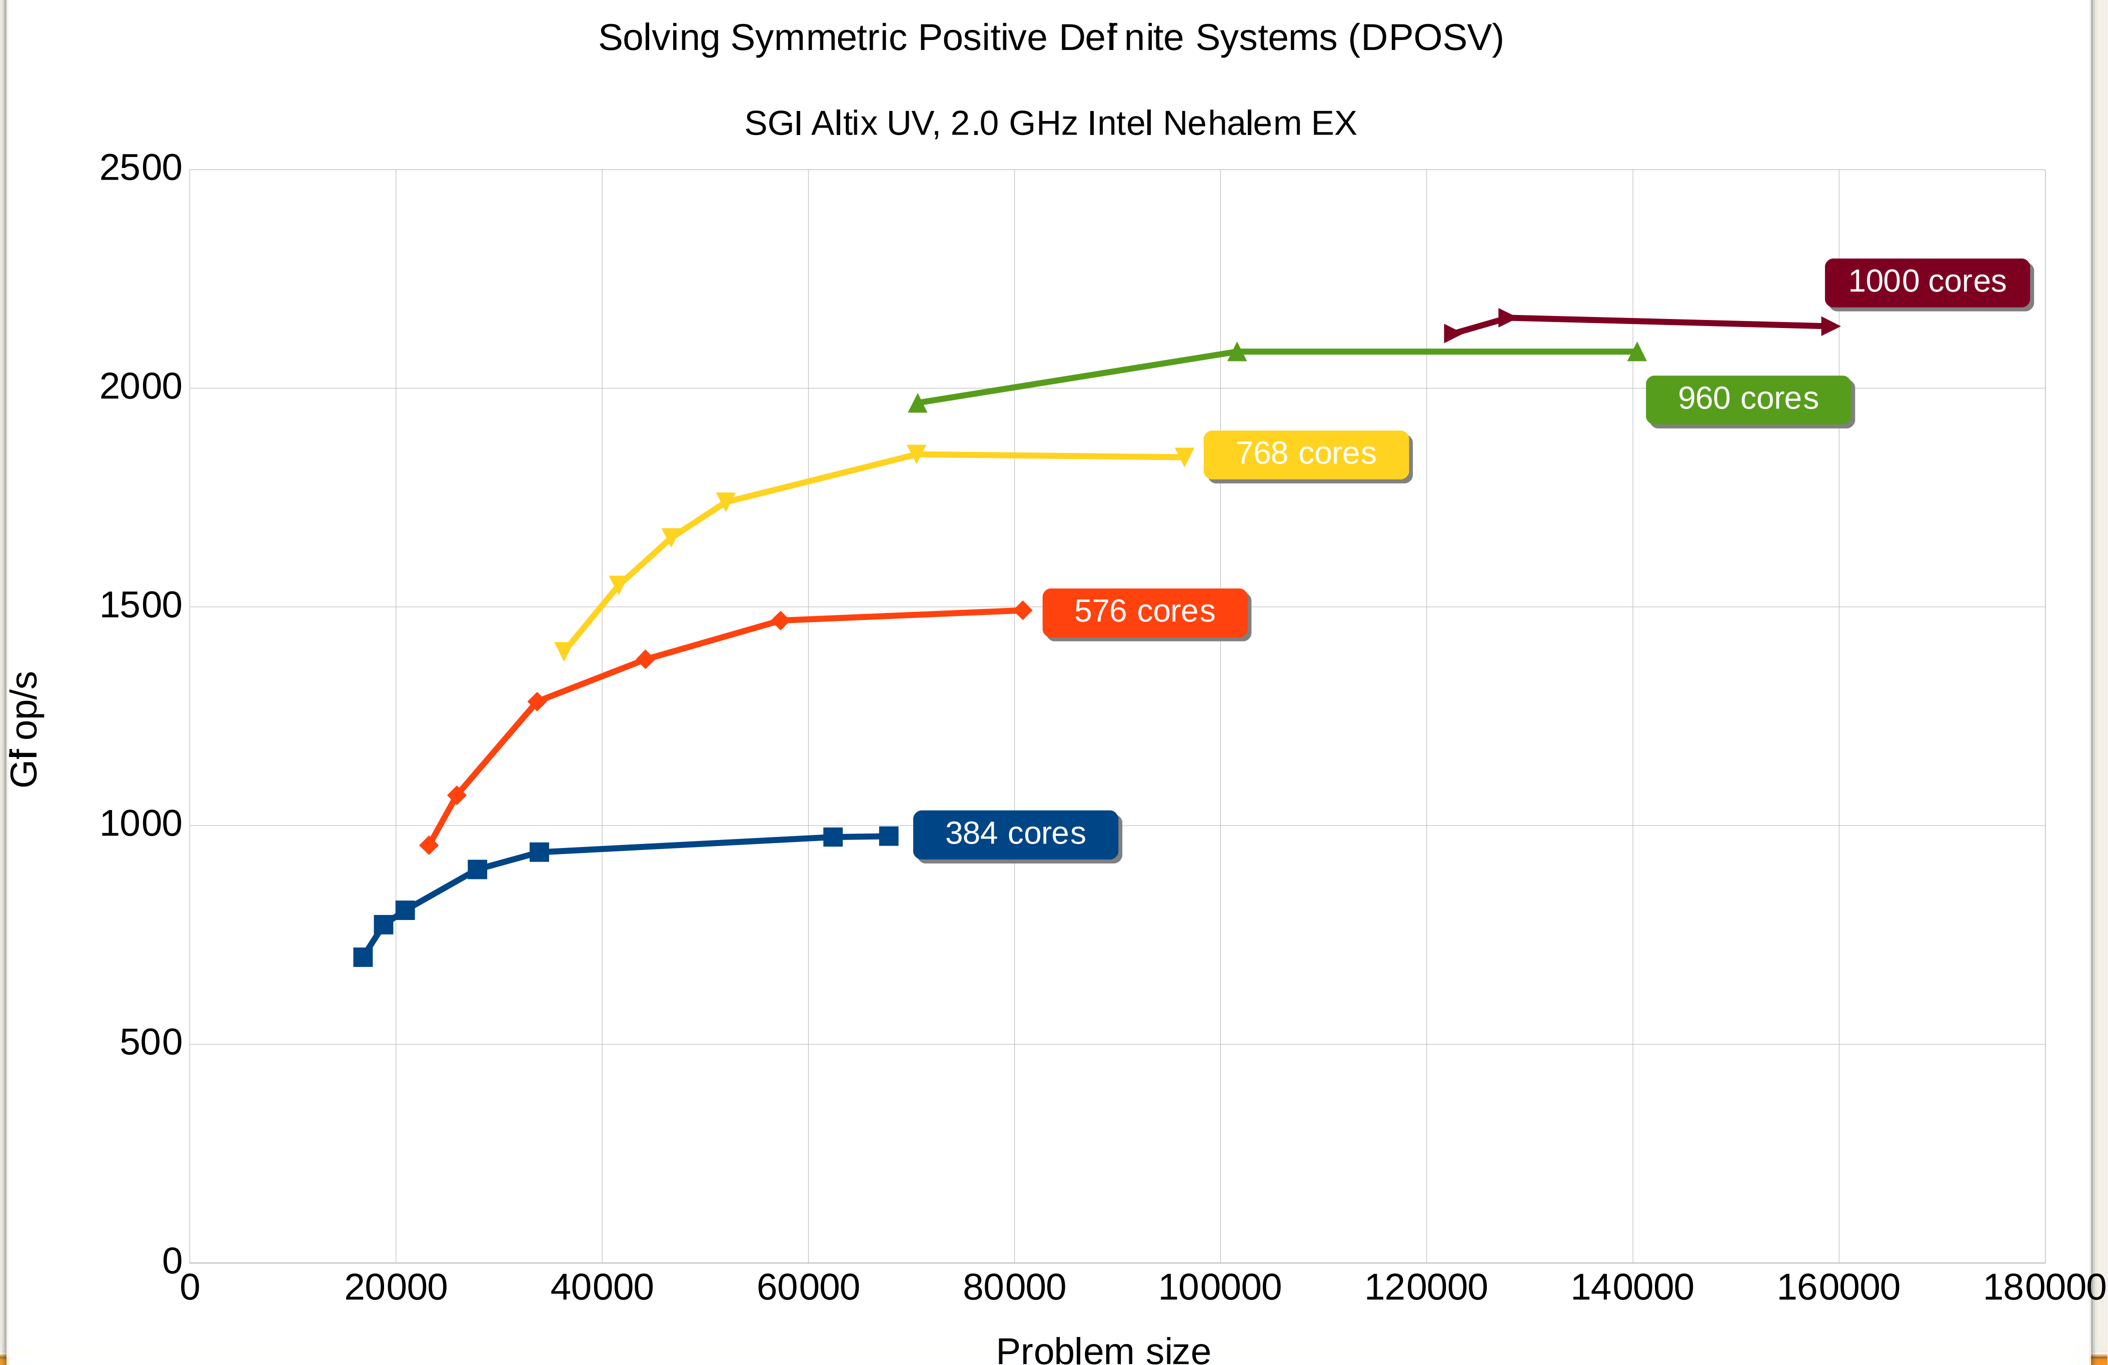

Solving Symmetric Positive Def nite Systems (DPOSV)
SGI Altix UV, 2.0 GHz Intel Nehalem EX
Gf op/s
2500
2000
1500
1000
500
0
0
20000
40000
60000
80000
100000
120000
140000
160000
180000
Problem size
1000 cores
960 cores
768 cores
576 cores
384 cores

Monitoring Power and Energy with PowerPack

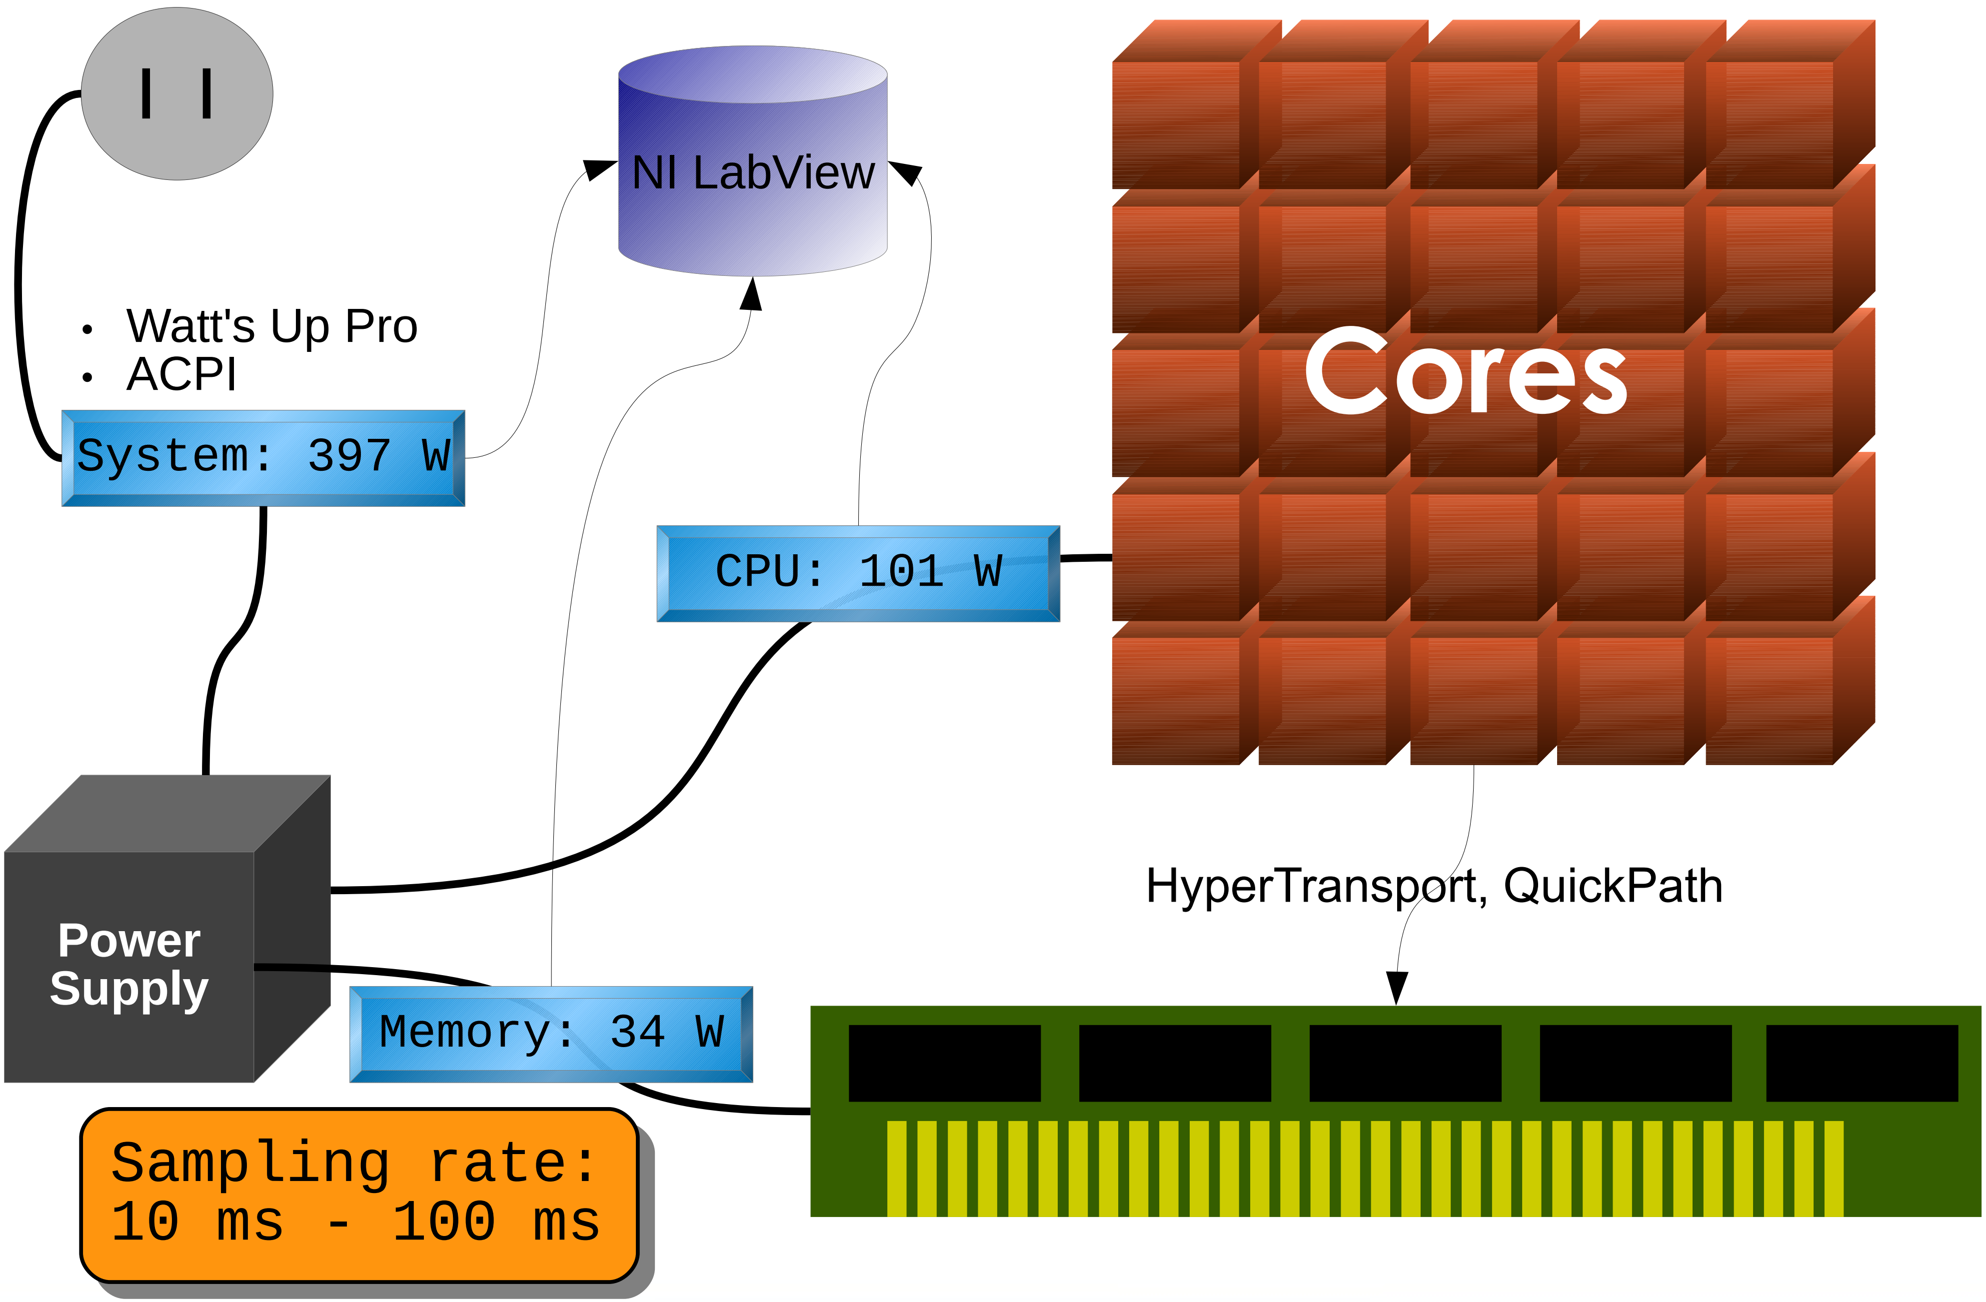

LU Factorization, 4 cores, Intel Core 2, 1.6 GHz

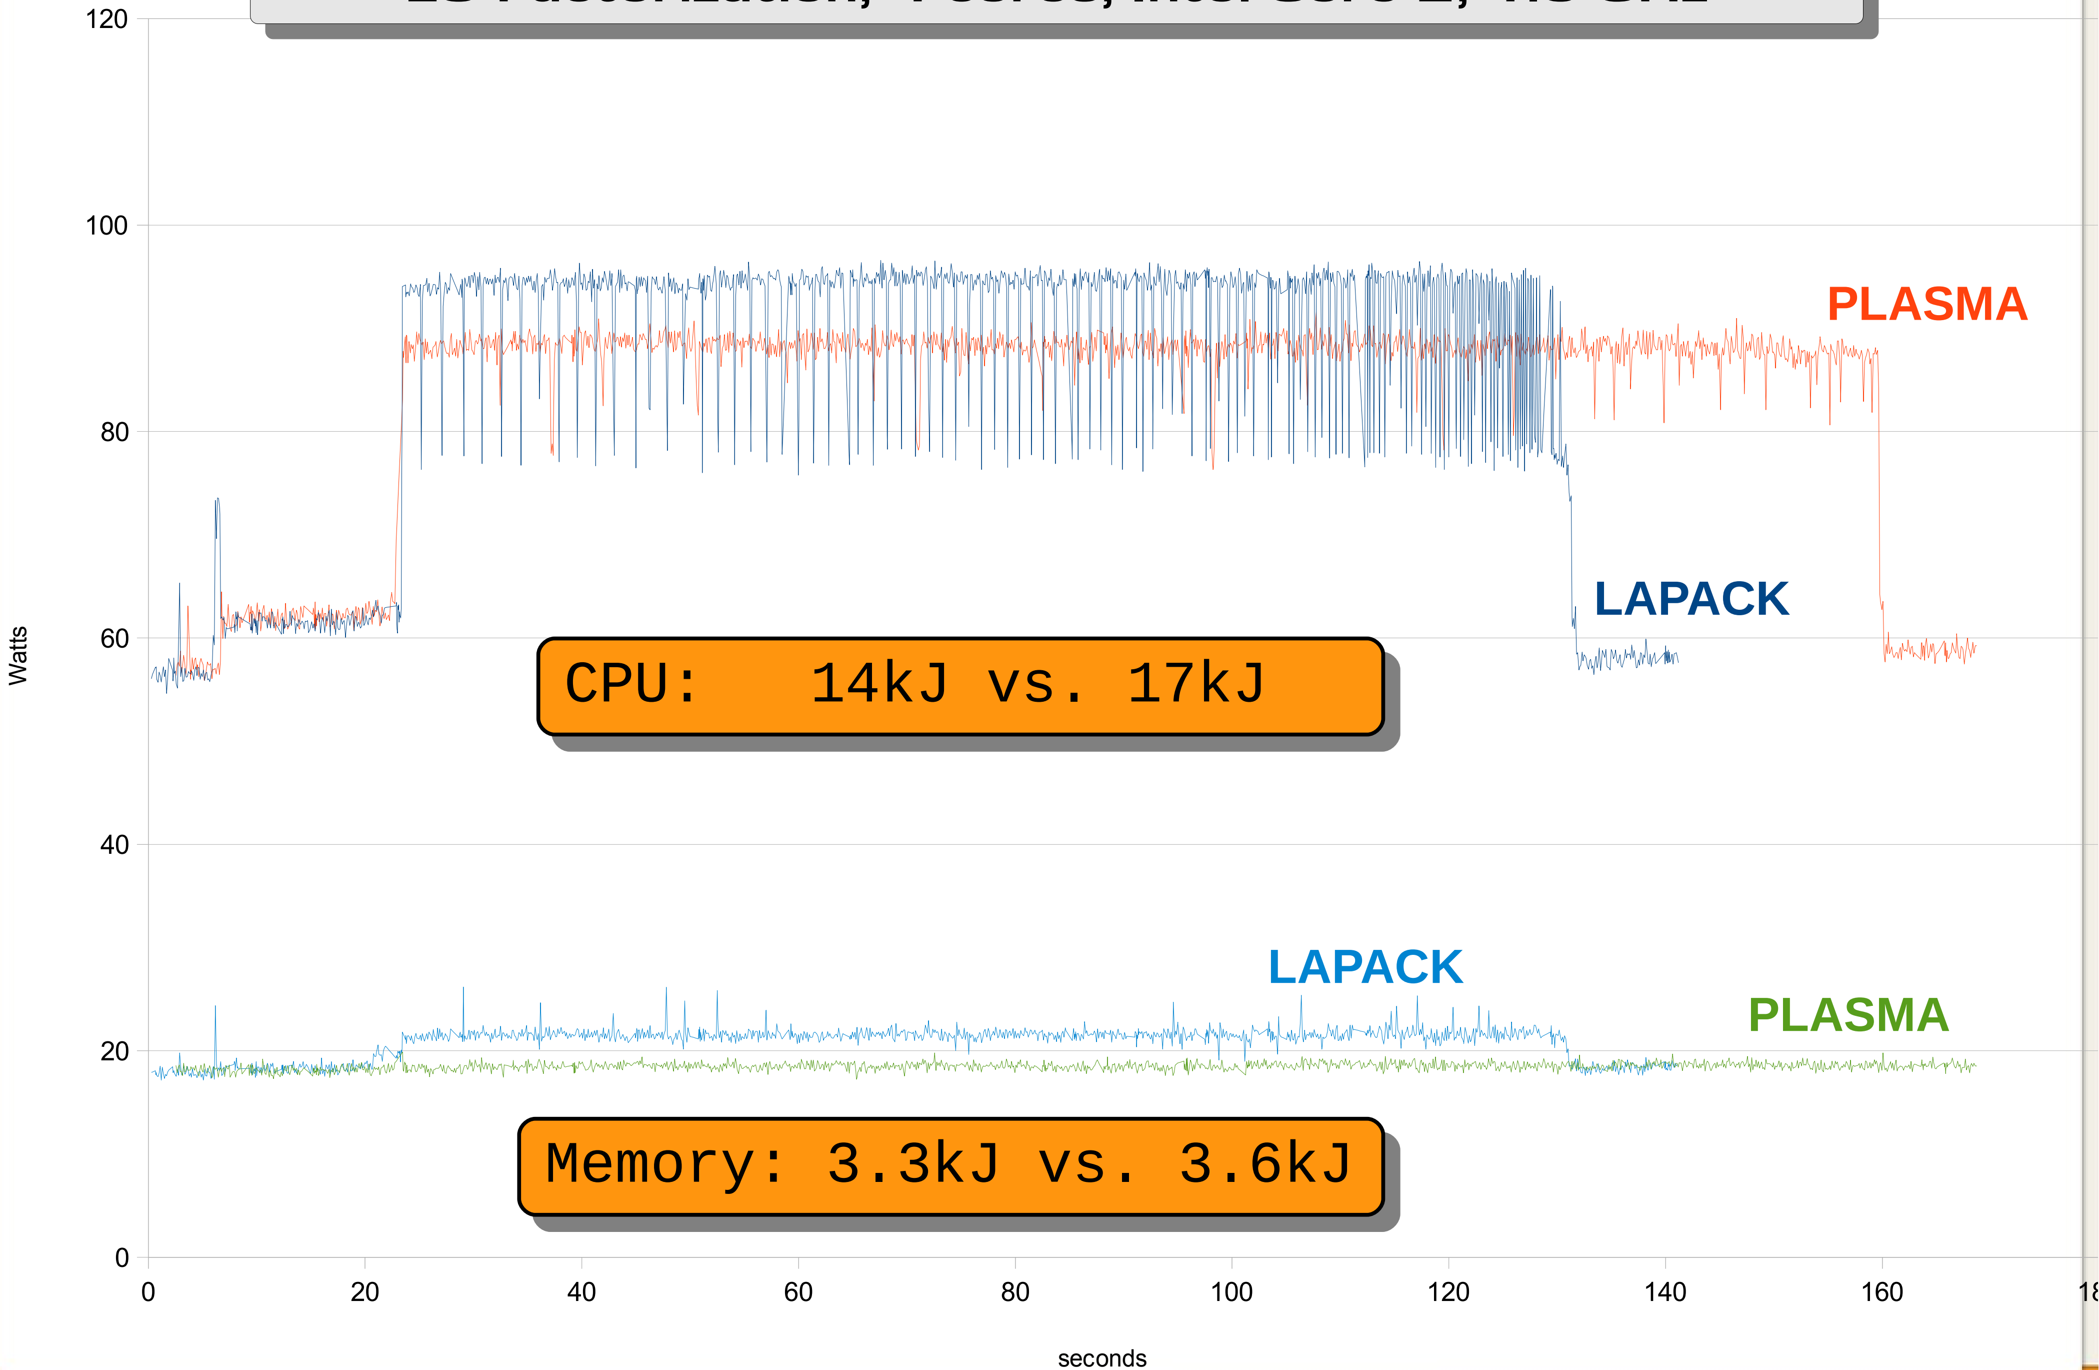

Power and Energy Metrics

- Hardware inspired metrics
  - Total power
  - Energy delay product
  - Power delay product
    - 1 / (Power*delay) = 1 / (Power/perf) = perf / Power [Gflop/s per Watt]
- Benchmarks

Rank	Name	TOP500	Green500	Power	Hours	M W h	Mil. Cores	MiB/core
1	Sequoia	16.3	2069	7890	23.1	182	1.57	780
2	K Comp	10.5	830	12660	29.5	373	0.71	1525
5	Tianhe-1	2.6	635	4040	3.4	14	0.19	531
6	Jaguar	1.9	377	5142	24.2	124	0.31	972
160	JAXA	0.1	119	1020	55.3	56	0.01	6942

LU-based Inversion
8-cores, 2.8 GHz

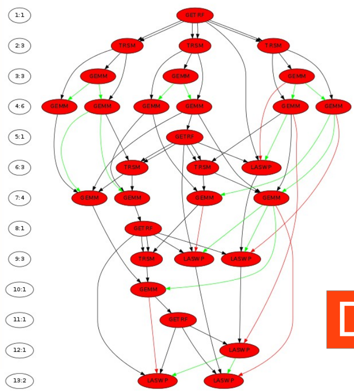

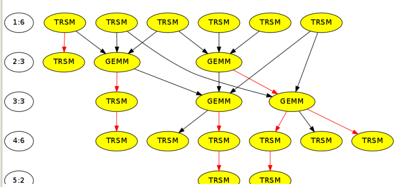

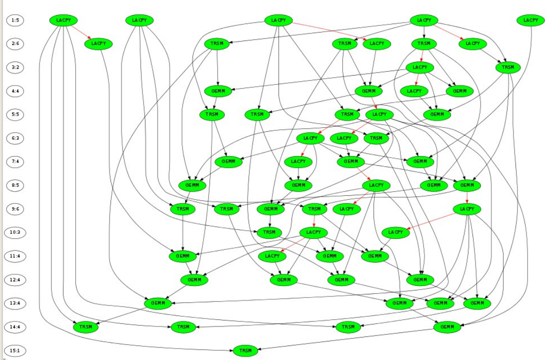

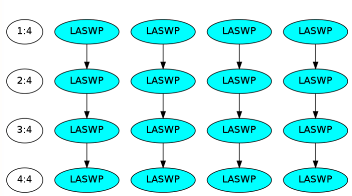

DAG Composition

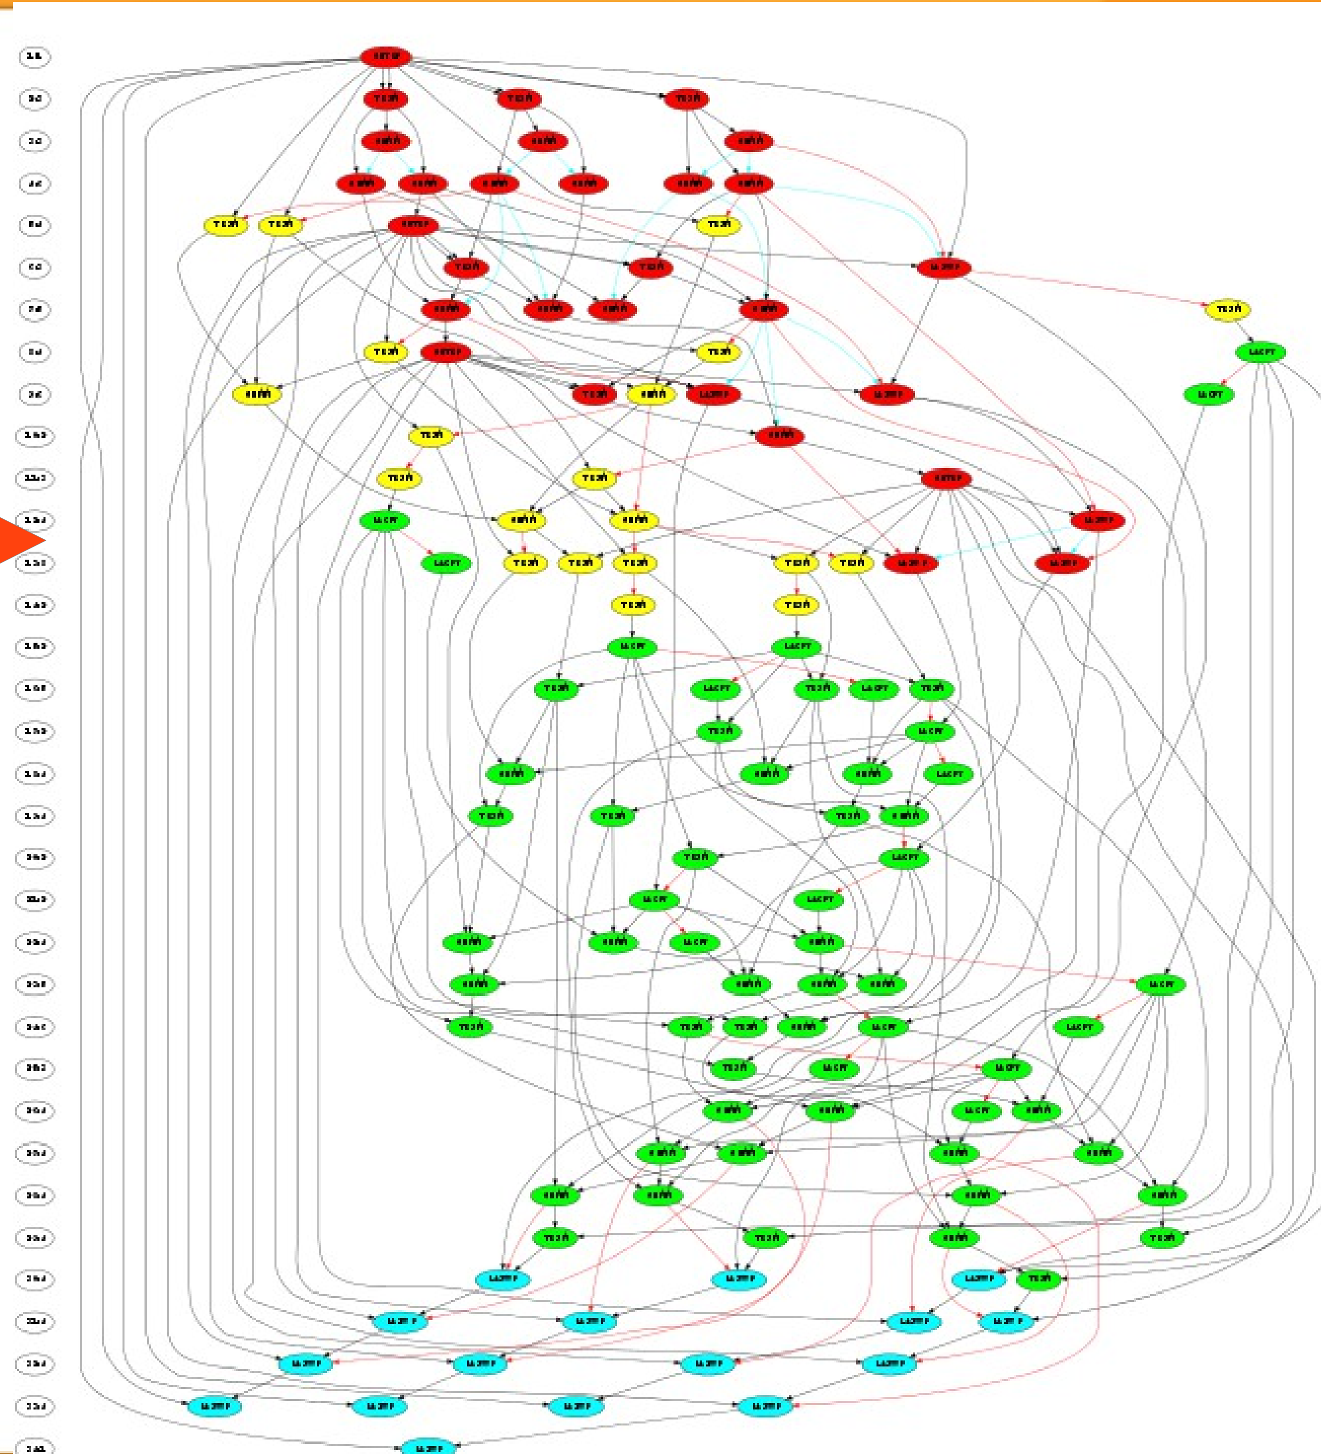

LU-based Inversion, Intel Core 2, 8-cores, 2.8 GHz

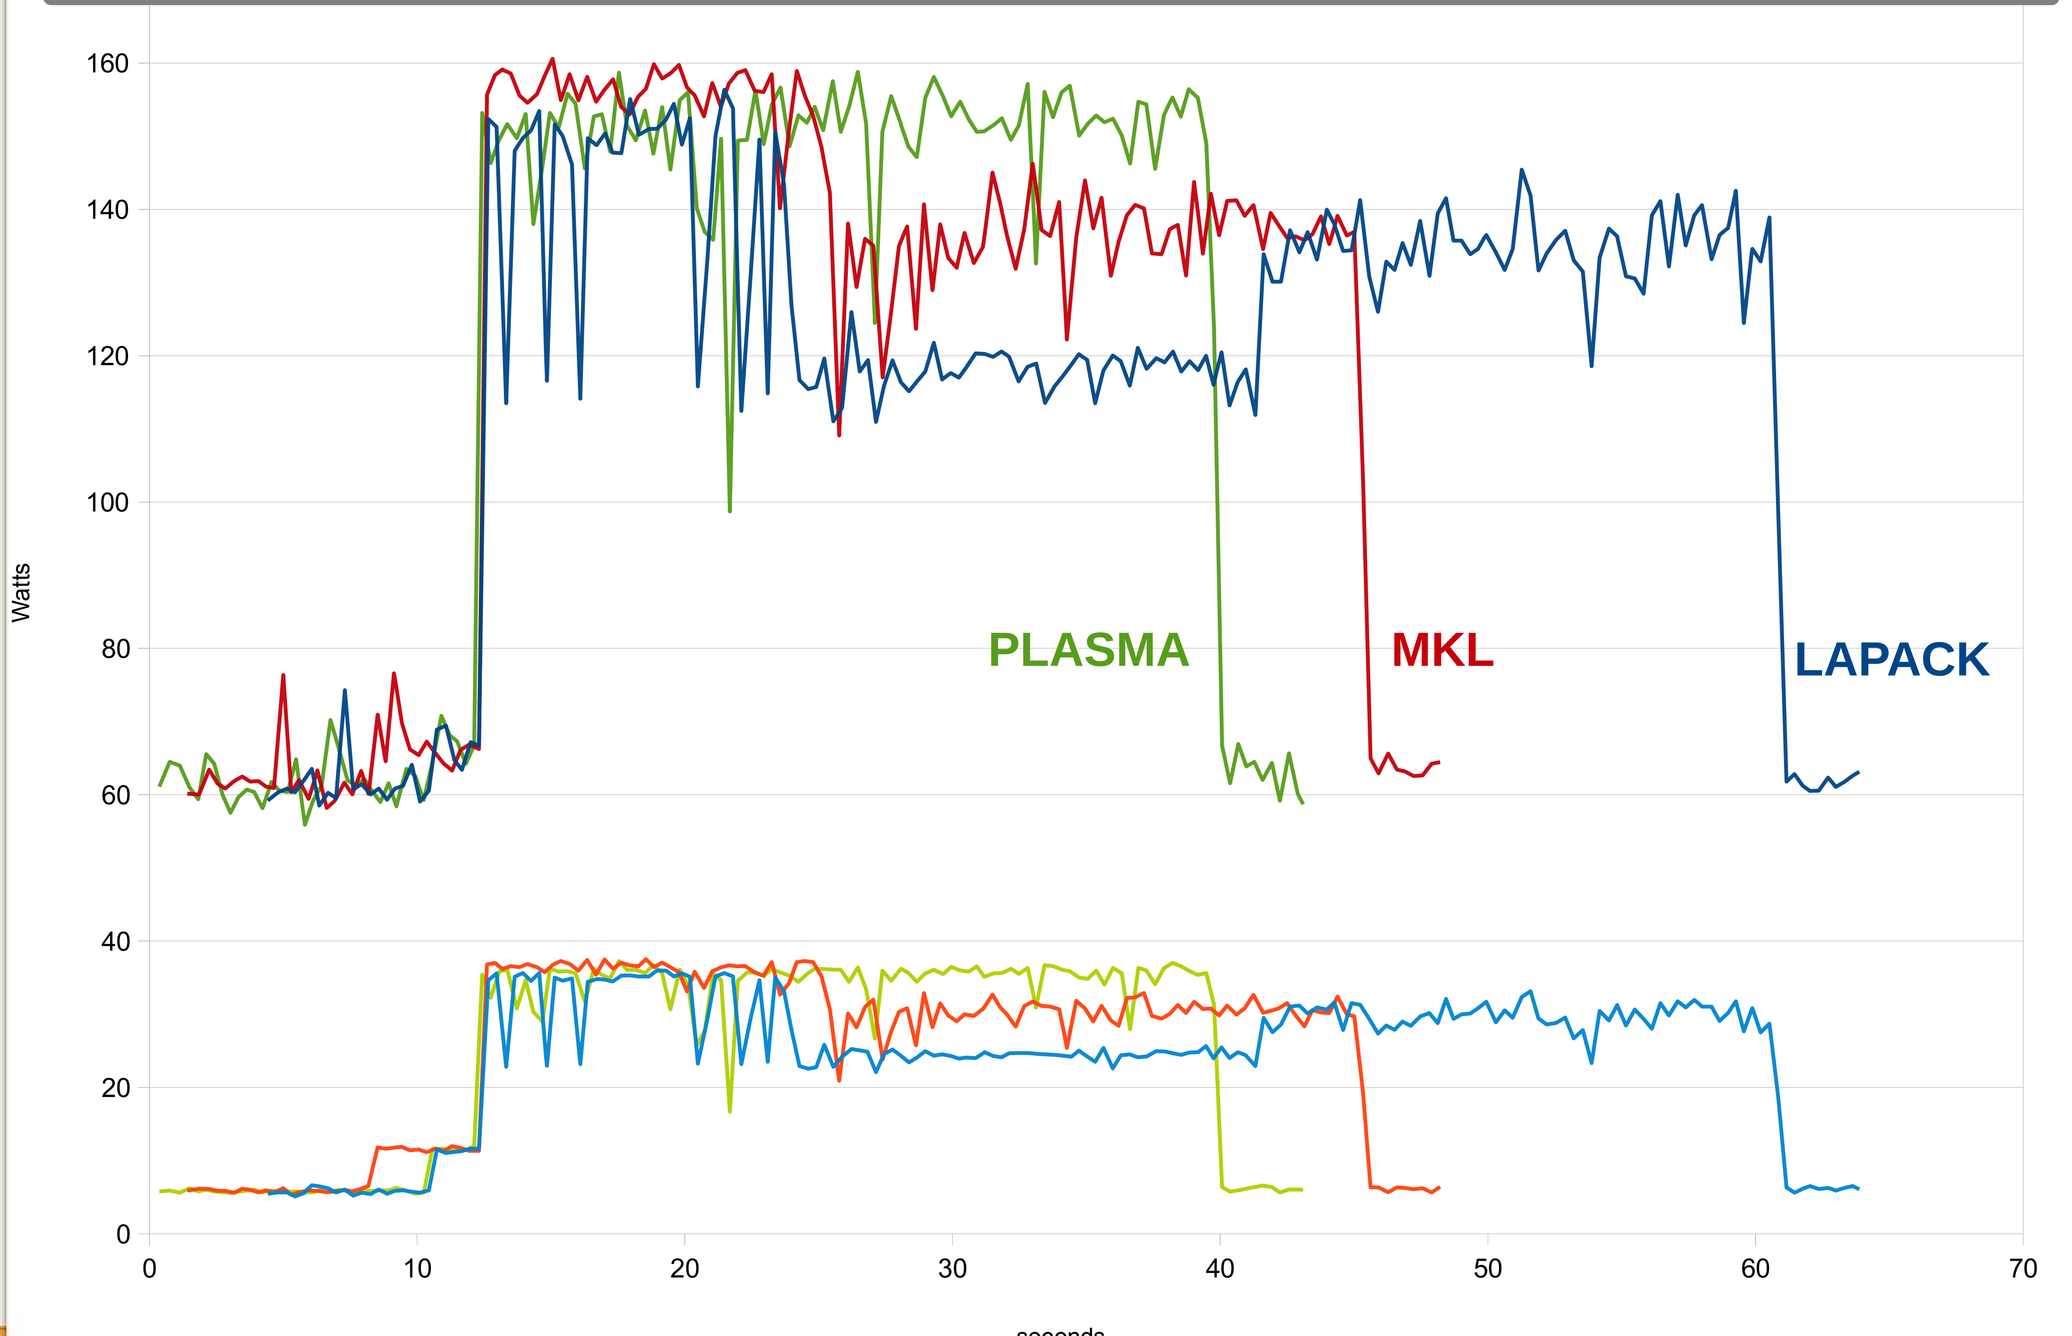

Tool Support for Power and Energy Measurement

Hardware-based solutions

PowerPack

PowerScope

Energy Endoscope

Counter-based solutions

PAPI-C

Intel RAPL

NVIDIA NVML

VMware

I/O

KVM

InifiniBand

LM

Cholesky Solve, Intel Sandy-Bridge E5-2690, 16-cores, 2.9 GHz

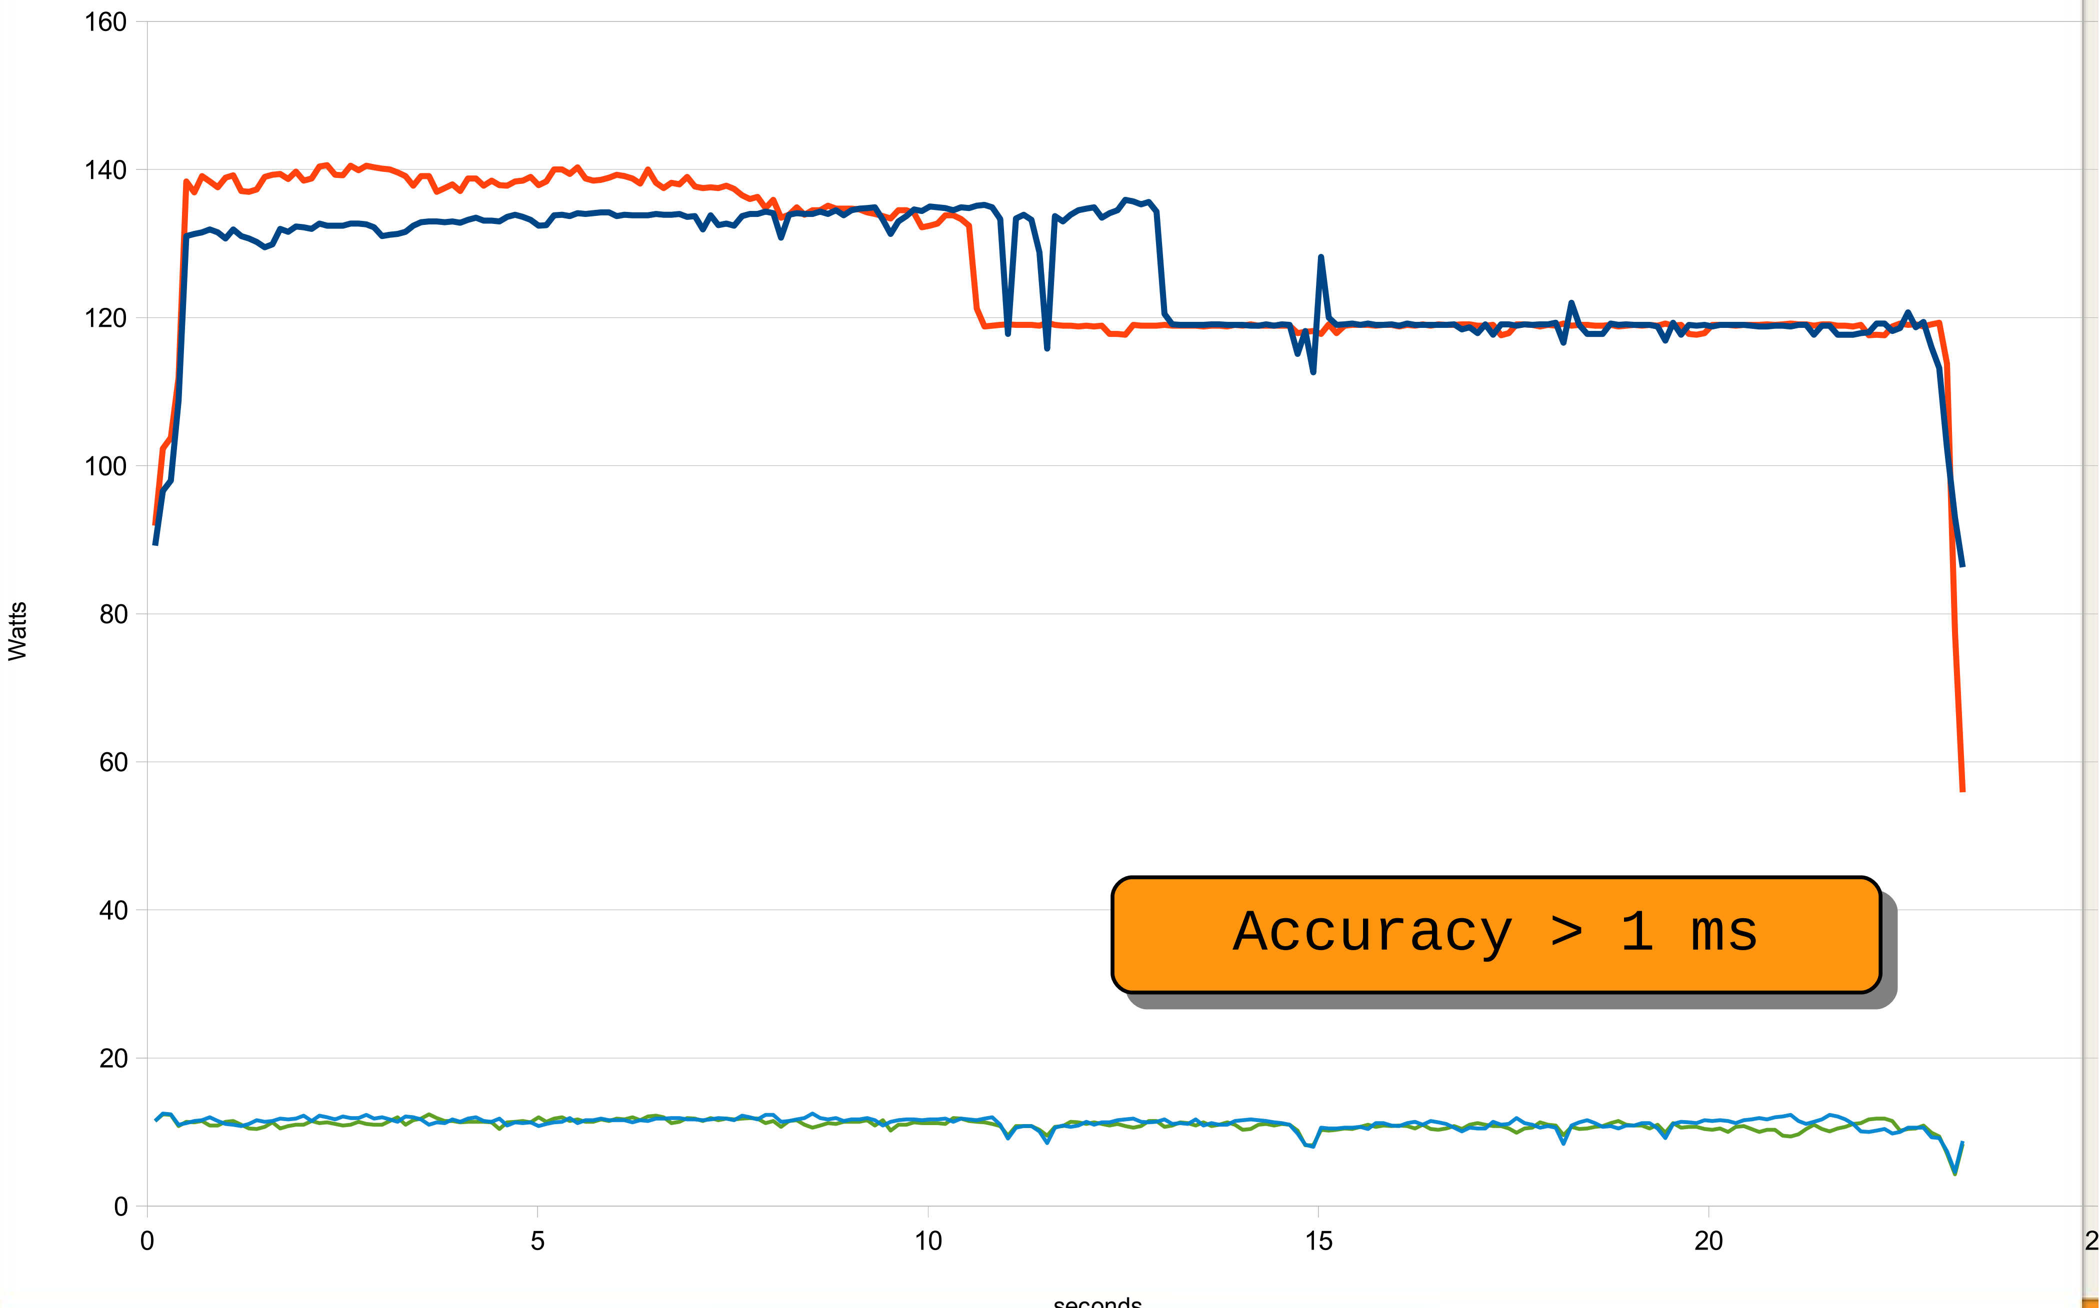

Accuracy > 1 ms

MAGMA Power Profile on Fermi C2075 with PAPI and NVML

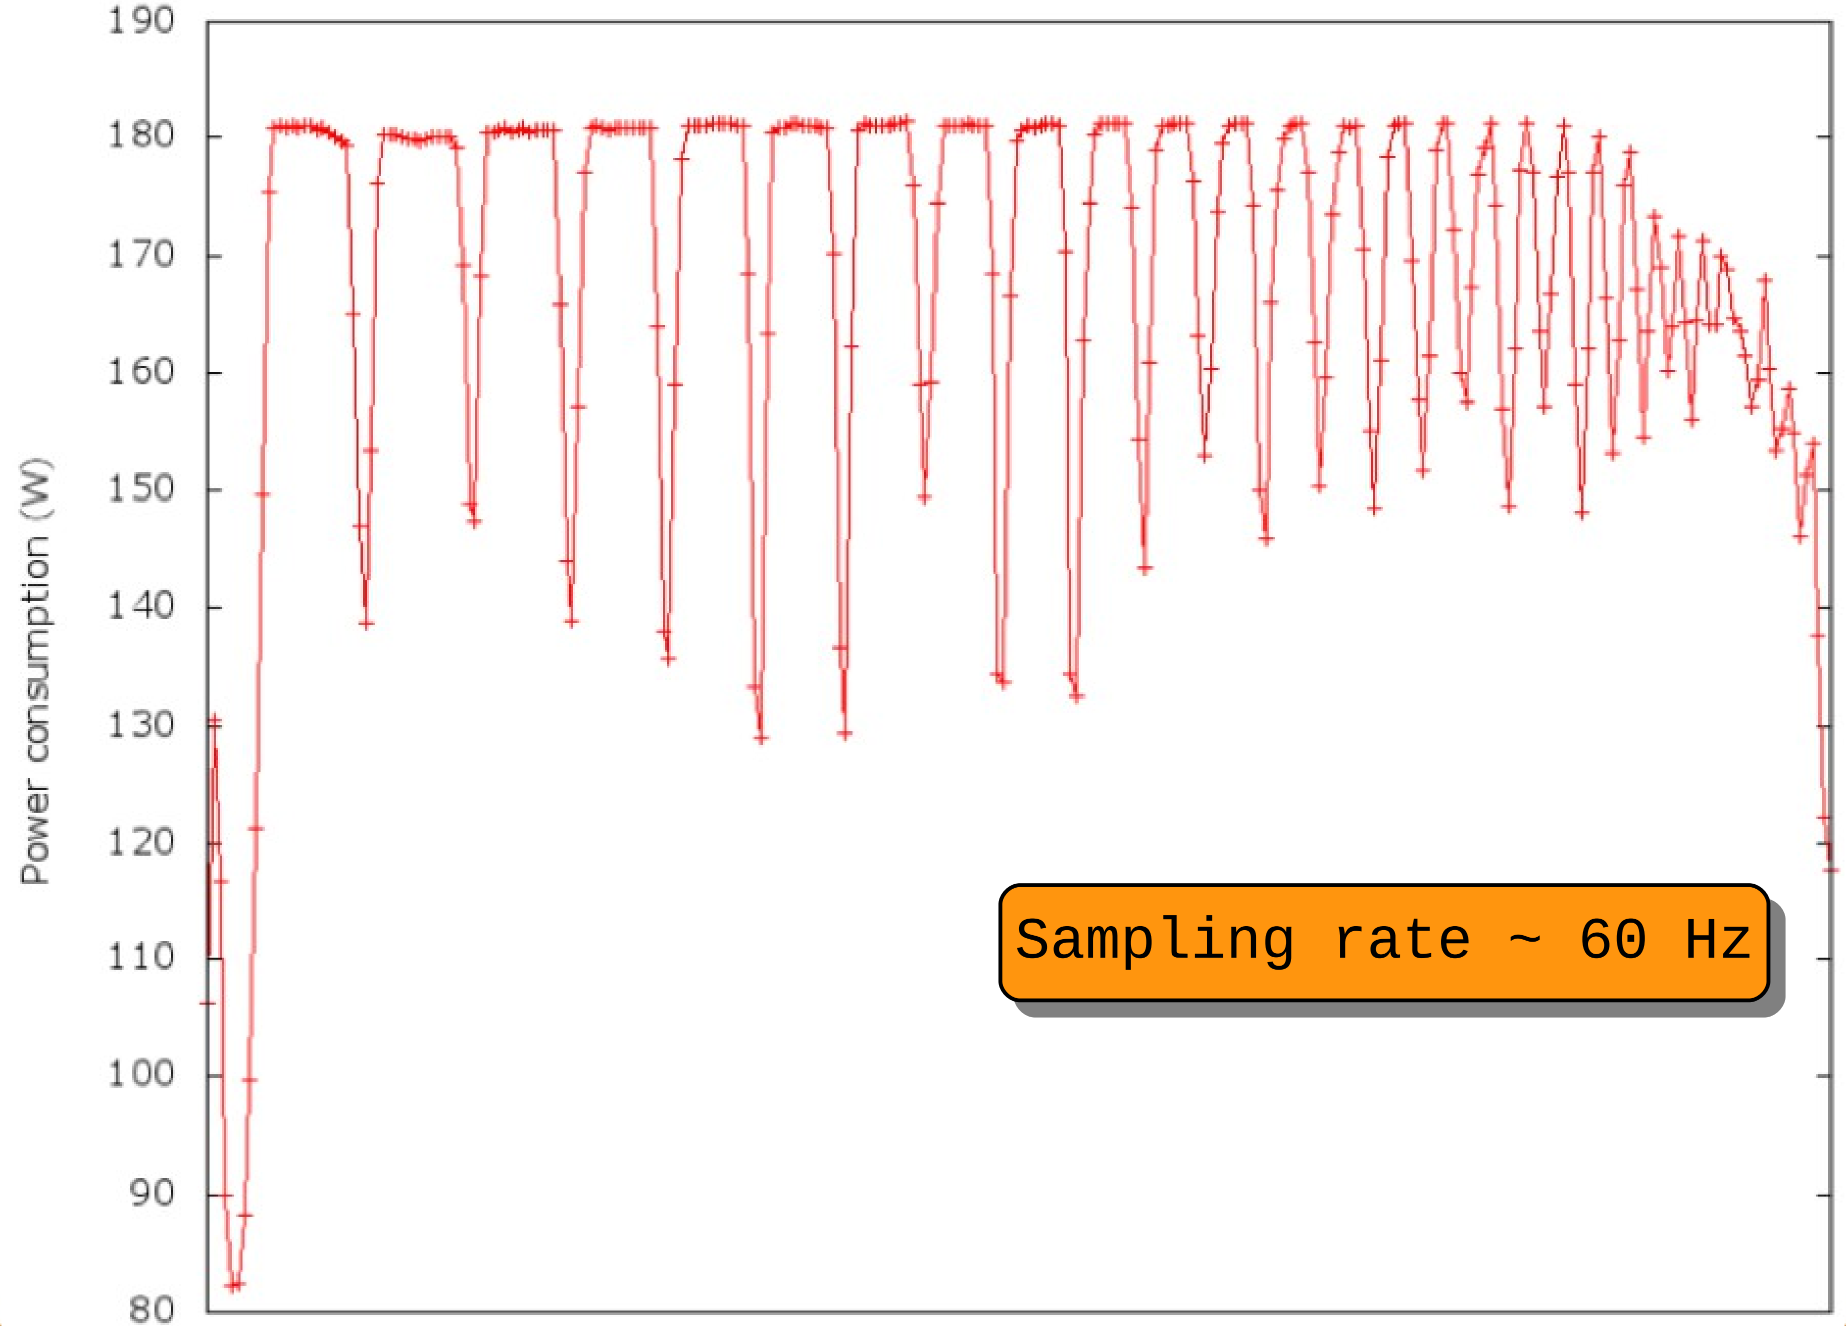

Collaborators

- Kirk Cameron, Virginia Tech
- Jack Dongarra, University of Tennessee
- Hatem Ltaief, KAUST
- Mathieu Faverge, ENSEIRB-Matmeca, INRIA, Bordeaux, France
- Jakub Kurzak, University of Tennessee
- Stanimire Tomov, University of Tennessee
- Vince Weaver, University of Maine