Evaluation of Emerging Energy-Efficient Computing Platforms for Biomolecular and Cellular Simulation Workloads

John E. Stone, Michael J. Hallock, James C. Phillips, Joseph R. Peterson, Zaida Luthey-Schulten, Klaus Schulten

NIH Center of Macromolecular Modeling and Bioinformatics

Beckman Institute for Advanced Science and Technology

University of Illinois at Urbana-Champaign

http://www.ks.uiuc.edu/

http://www.scs.illinois.edu/schulten/lm/

25th Heterogeneity in Computing Workshop

IEEE International Symposium on Parallel and Distributed Processing

Chicago, IL, May 23, 2016

Biomolecular and Cellular Simulation

- LM: Lattice Microbes
  - Lattice-based simulations of bacterial cells using reaction diffusion models
- NAMD: Molecular Dynamics
  - Classical mechanics simulation of proteins, biomolecular complexes, viruses, organelles
- VMD: Visual Molecular Dynamics
  - Visualization and analysis of molecular and lattice cell simulations

E.coli cell

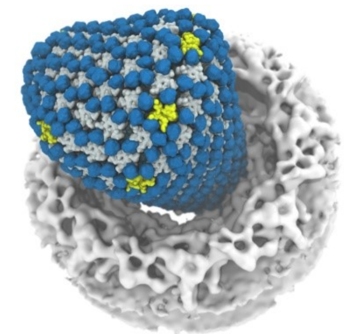

HIV-1 capsid and nuclear pore complex

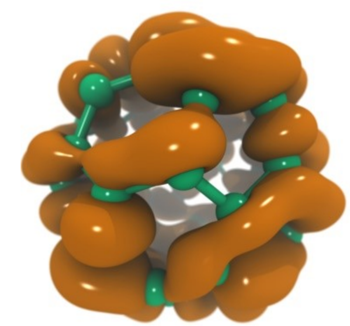

Molecular orbitals, vibrating C_{60}

Goal: A Computational Microscope

Study the molecular machines in living cells

Ribosome: target for antibiotics

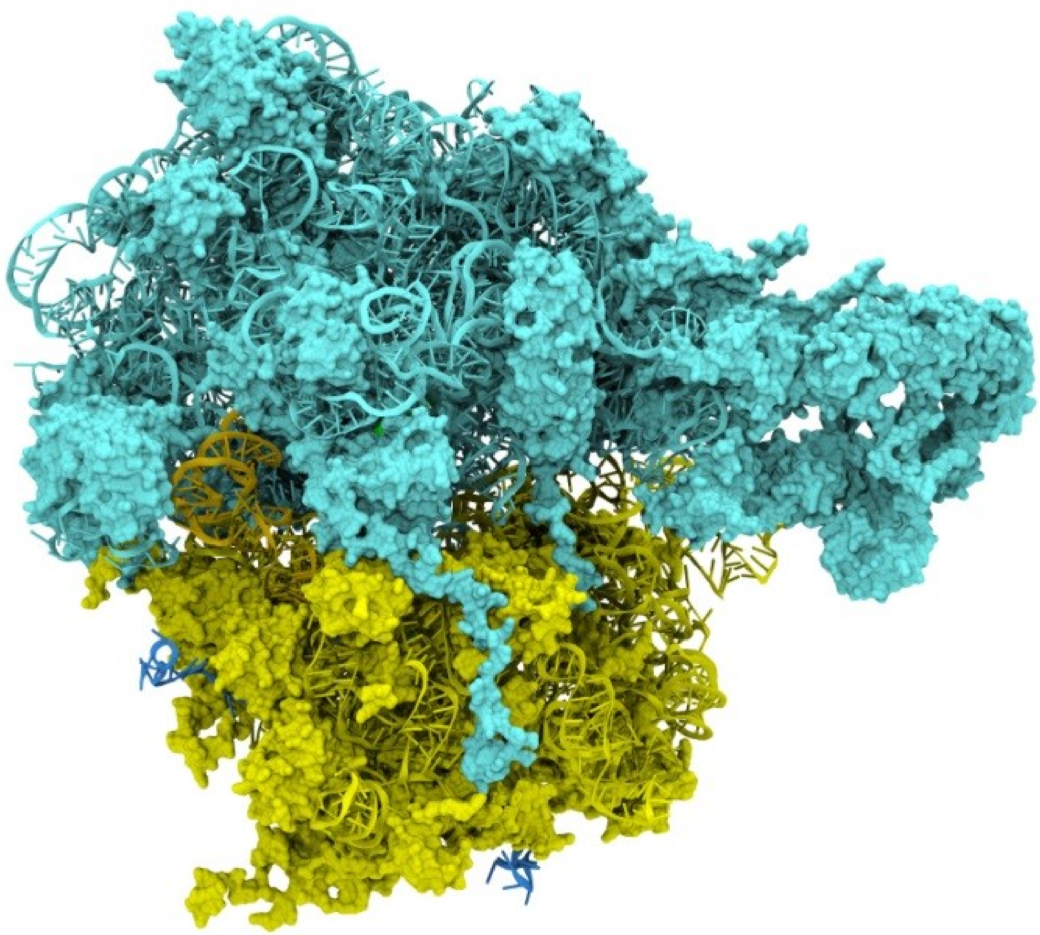

Poliovirus

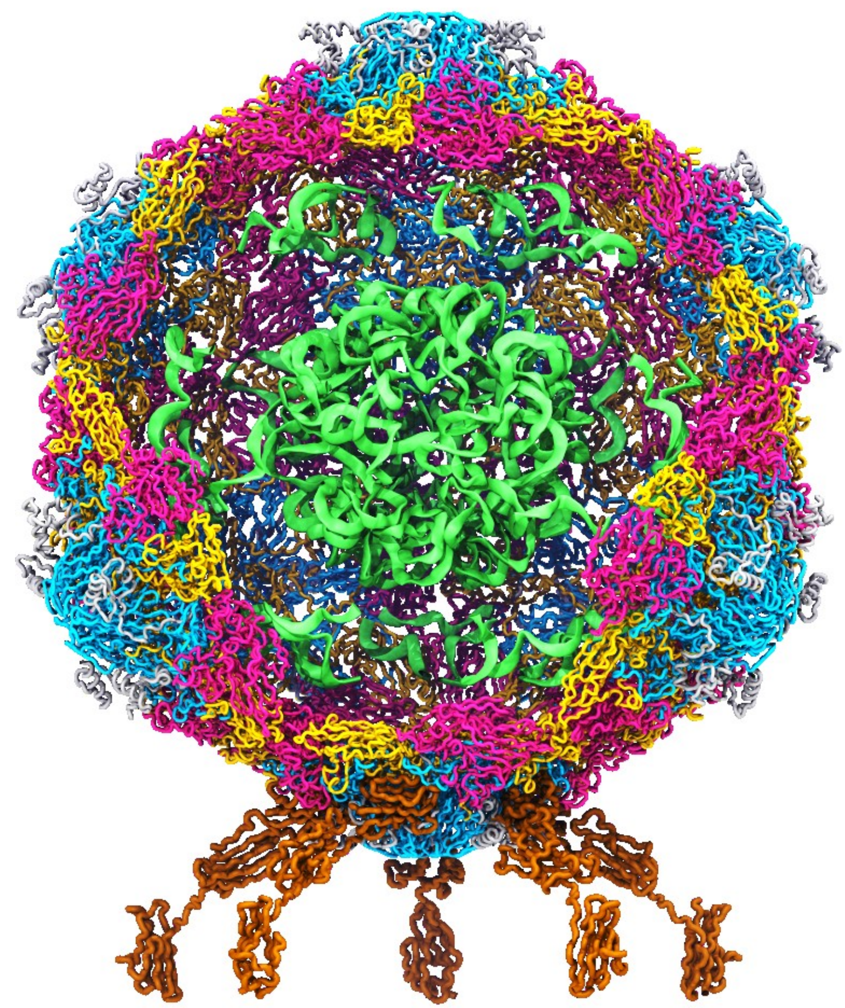

LM, NAMD, and VMD Use GPU Accelerators and Petascale Computing to Meet Computational Biology's Insatiable Demand for Processing Power

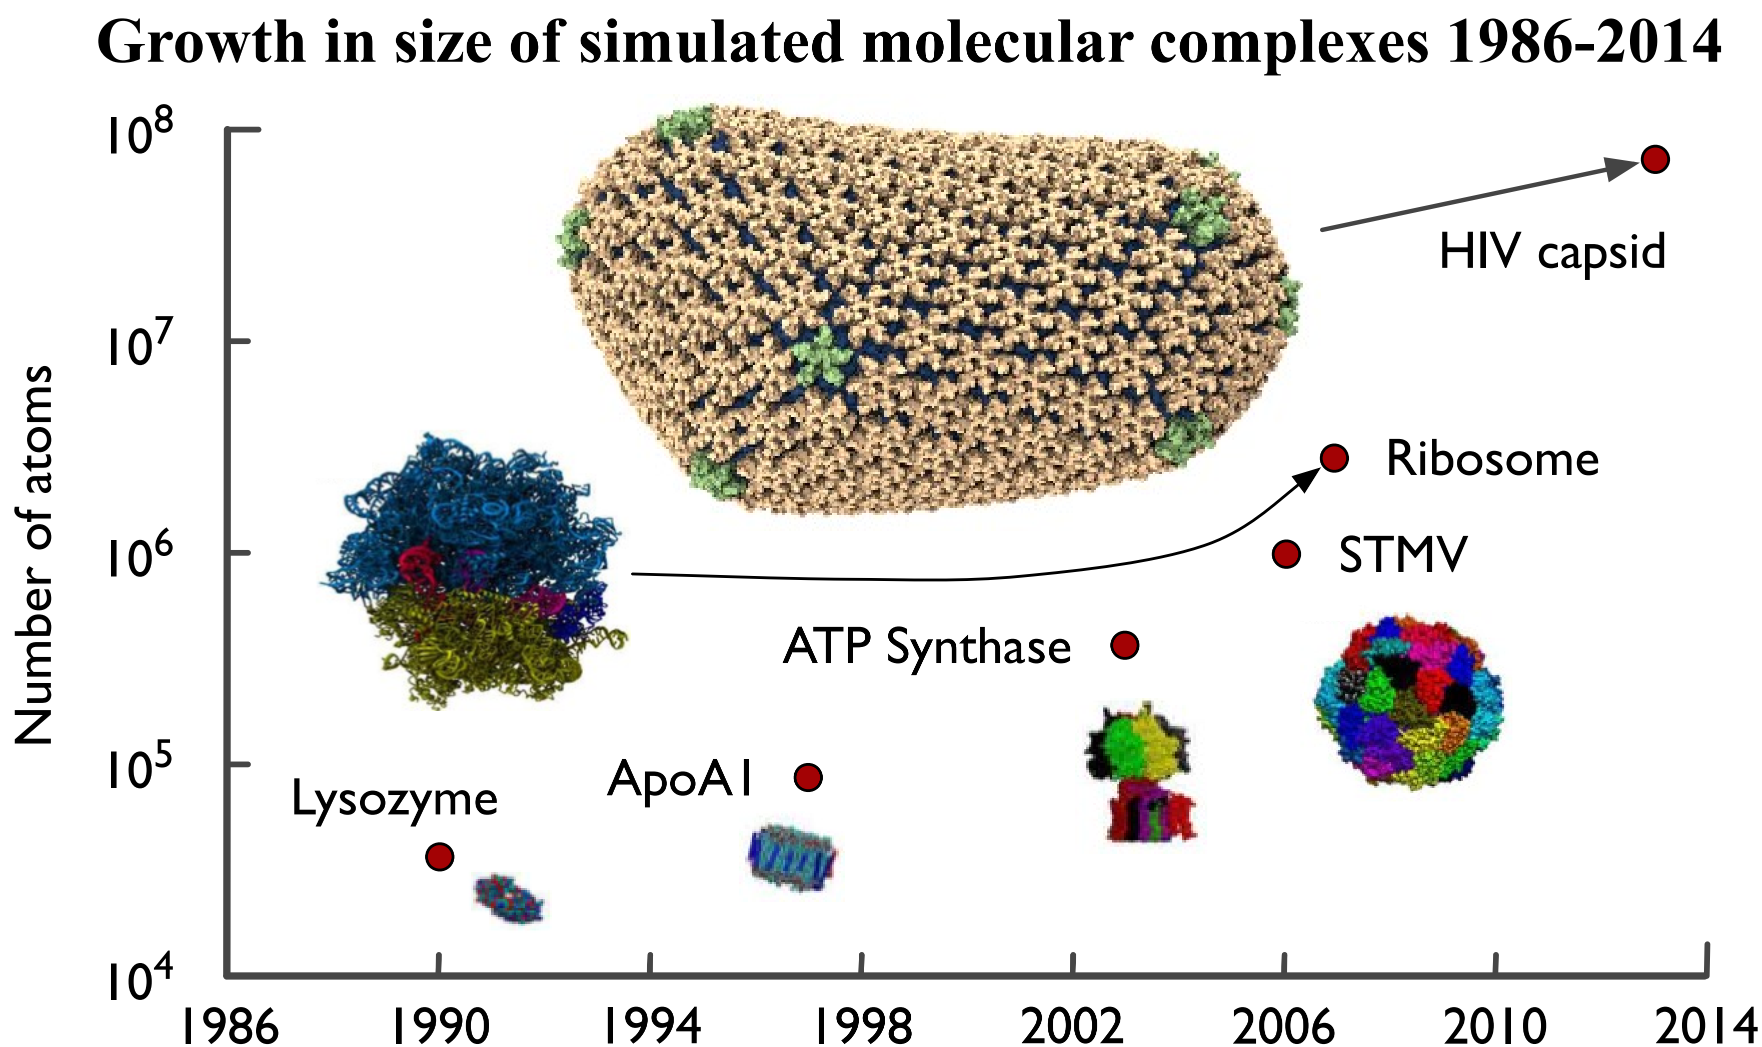

NAMD Titan XK7 Performance August 2013

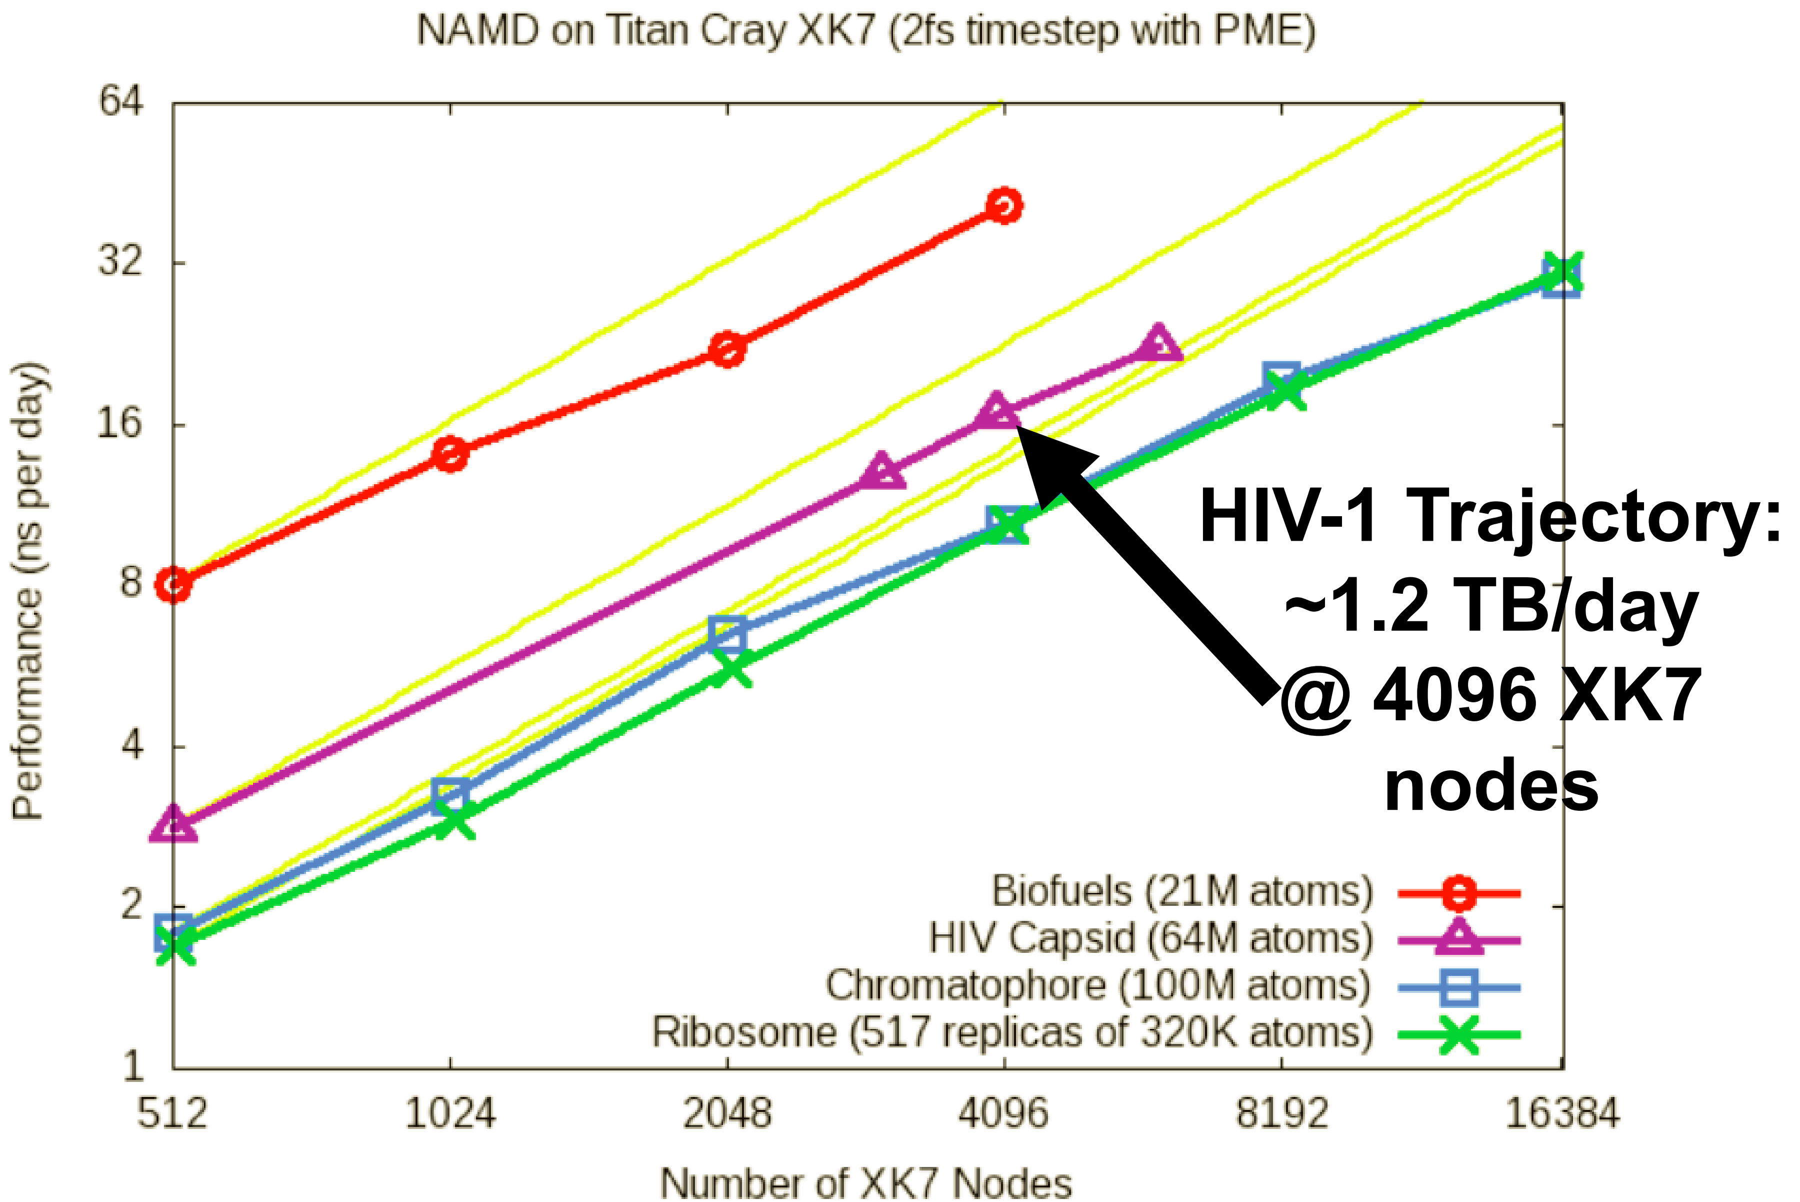

Continued Growth in Simulation Performance Requires Increased Energy Efficiency

- GPUs have revolutionized MD+cell sim., 2007-present
- Kernel perf. increases of 5x to 10x are common
- Amdahl's law pushes apps to leverage GPUs to an increasing degree for best performance
- Codes such as LM and HOOMD solely for GPU-accelerated platforms, CPU is just doing bookkeeping!
- **State-of-the-art GPUs are now often thermally limited**
- **Questions:**
  - **Are emerging ARM+GPU platforms competitive for MD+cell sim?**
  - **Are they more energy efficient than conventional x86?**
  - **Why, why not?**

NCSA AC Cluster GPU Performance and Power Efficiency Results (2010)

Application	GPU speedup	Host watts	Host+GPU watts	Perf/watt gain
NAMD	6	316	681	2.8
VMD	25	299	742	10.5
MILC	20	225	555	8.1
QMCPACK	61	314	853	22.6

Quantifying the Impact of GPUs on Performance and Energy Efficiency in HPC Clusters. J. Enos, C. Steffen, J. Fullop, M. Showerman, G. Shi, K. Esler, V. Kindratenko, J. Stone, J. Phillips. *The Work in Progress in Green Computing,* pp. 317-324, 2010.

Emerging ARM+GPU Platforms

A) CARMA: Tegra3 + Quadro 1000M GPU
B) KAYLA: Tegra3 + PCIe 2.0 x4 Discrete GPU
C) Jetson TK1: Tegra K1 (iGPU)
D) Jetson TX1: Tegra X1 (iGPU)
E) APM X-Gene: X-Gene + PCIe 2.0 x8 Tesla K20c GPU

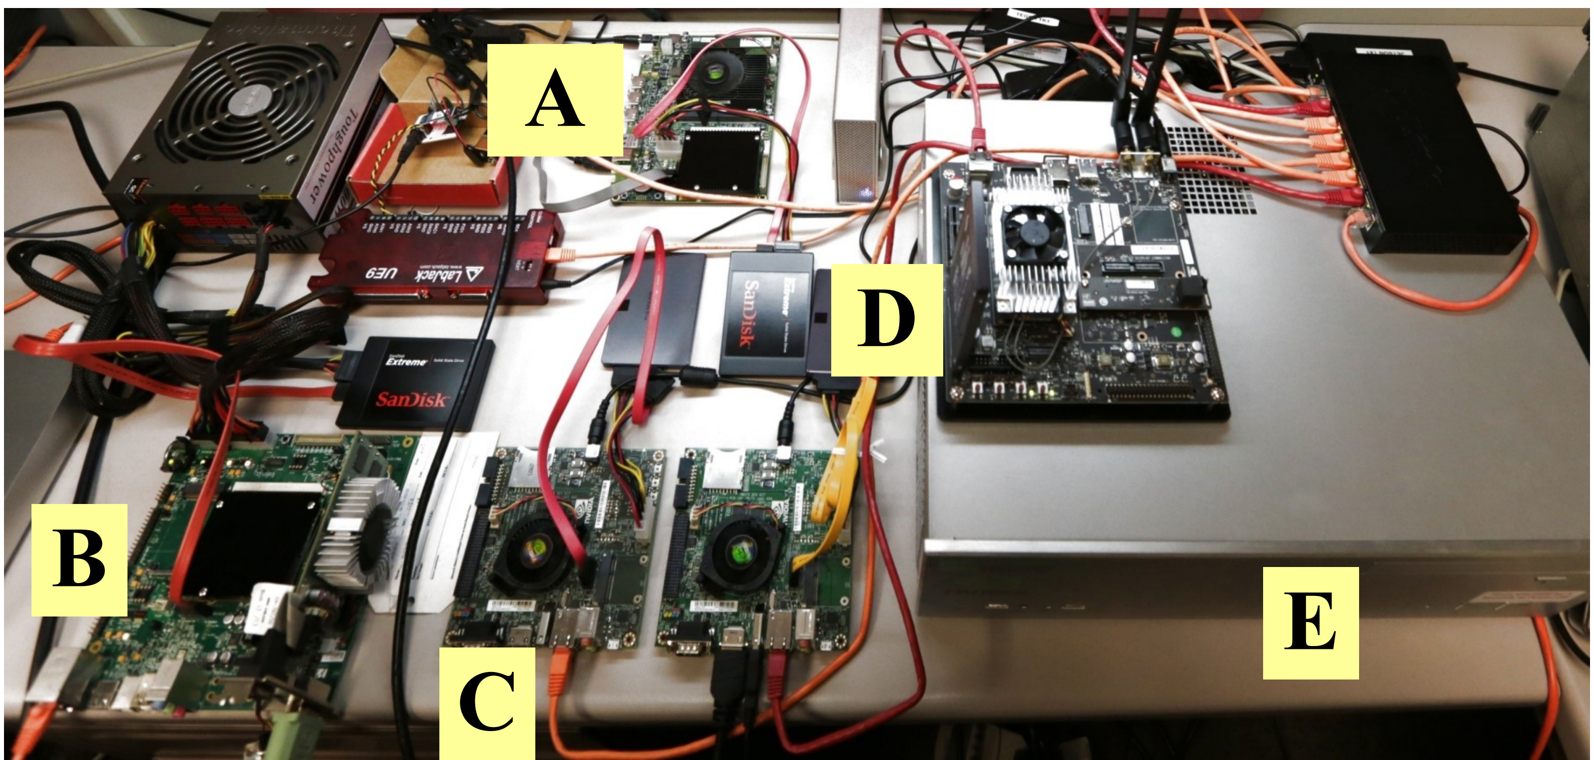

ARM Platform Porting + Eval Challenges

- ARM platform Linux differences:
  - ARM Linux much less standardized than x86
  - Kernel schduler **DVFS response time to load variation appears much longer than x86**
  - Kernel scheduler **powers off entire cores** and/or migrates processes between perf. and energy efficiency-optimized cores
  - **Dynamically varying number of available cores breaks conventional CPU work scheduler approaches** e.g. as in TBB, OpenMP – at startup they only see one available CPU core
  - **We modified apps to cope with varying core counts at runtime**
- ARM GPU drivers don't (yet) support DVFS on-par with x86, **neither platform supports app-controlled GPU DVFS in user-mode processes**
- ARM platforms tested lacked power monitoring APIs/hw – we used external instrumentation for all reported tests

VMD C_{60} Molecular Kernel Orbital Performance

Platform	C_{60} MO Kernel Execution Time (s)
CARMA Tegra 3 + Quadro 1000M	2.170 s
Jetson TK1 Tegra K1	2.020 s
Jetson TX1 Tegra X1 (Beta sw)	1.210 s
KAYLA Tegra 3 + GeForce 640	0.989 s
KAYLA Tegra 3 + GeForce Titan	0.396 s
APM X-Gene + Tesla K20c	0.243 s
i7-3960X + Tesla K20c	0.208 s
I7-3960X + GeForce Titan	0.157 s

For VMD MO kernels which are GPU-dominant, ARM platforms using comparable GPUs can approach conventional x86 platforms

Power Monitoring Instrumentation: Commercial Kill-a-Watt

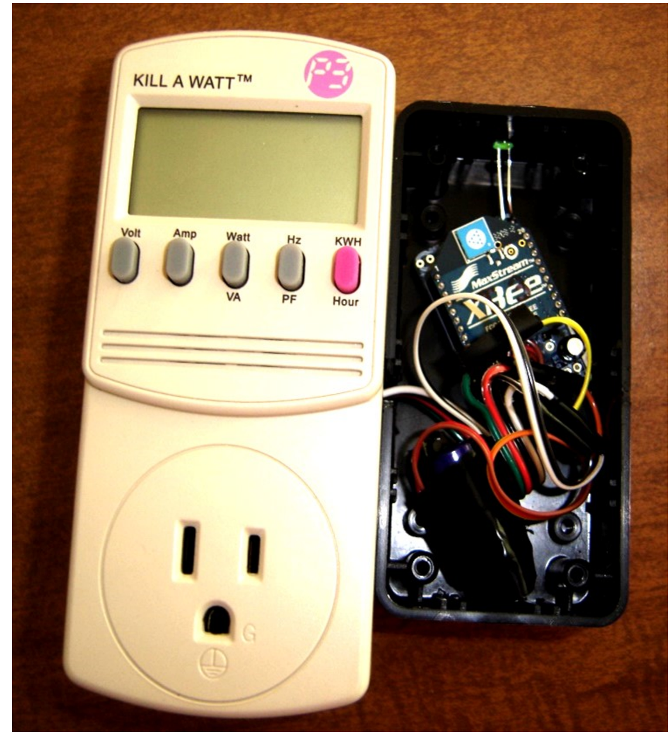

NCSA "AC" GPU cluster and Tweet-a-watt wireless power monitoring device. 0.2% accuracy w/ standard device

Quantifying the Impact of GPUs on Performance and Energy Efficiency in HPC Clusters. J. Enos, C. Steffen, J. Fullop, M. Showerman, G. Shi, K. Esler, V. Kindratenko, J. Stone, J. Phillips. *The Work in Progress in Green Computing,* pp. 317-324, 2010.

Power Monitoring Instrumentation: ACS712 Current Sensor + LabJack ADC

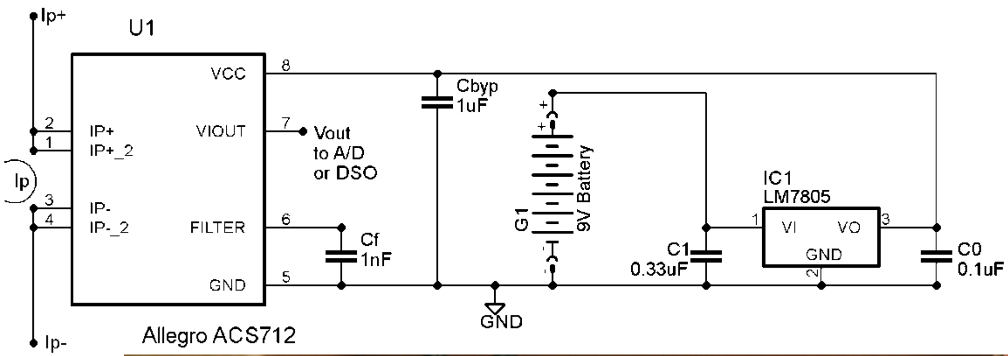

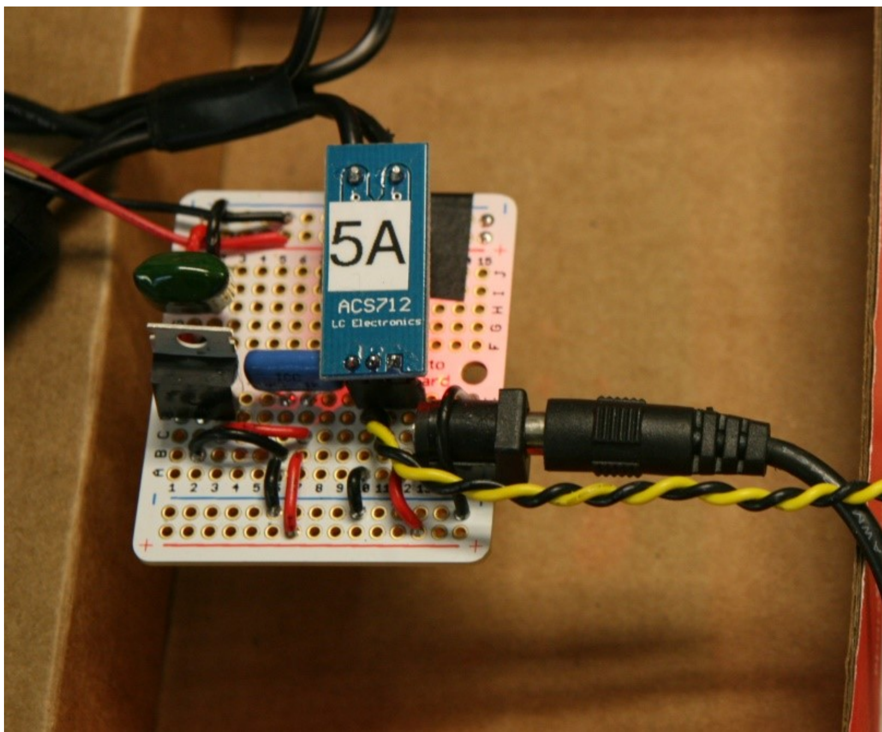

Power Monitoring Instrumentation: CARMA Attached to ACS712+DSO

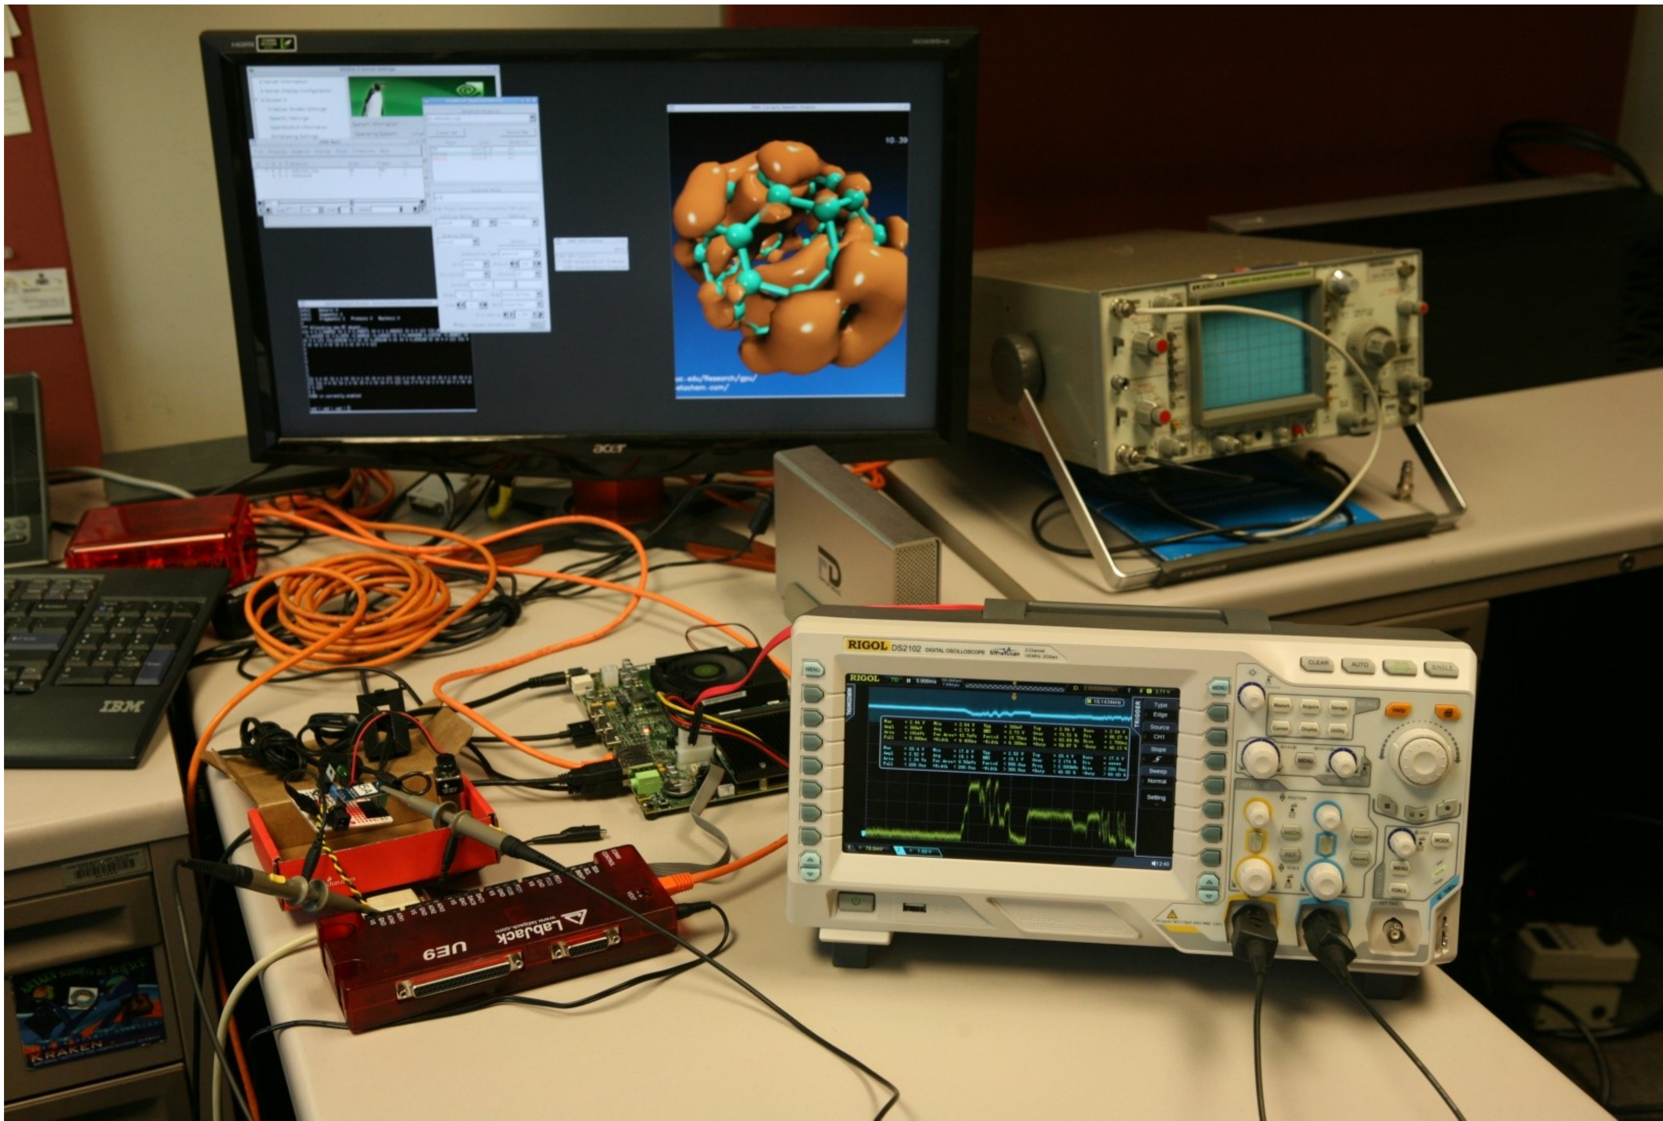

Power Monitoring Instrumentation: CARMA Attached to ACS712+ADC

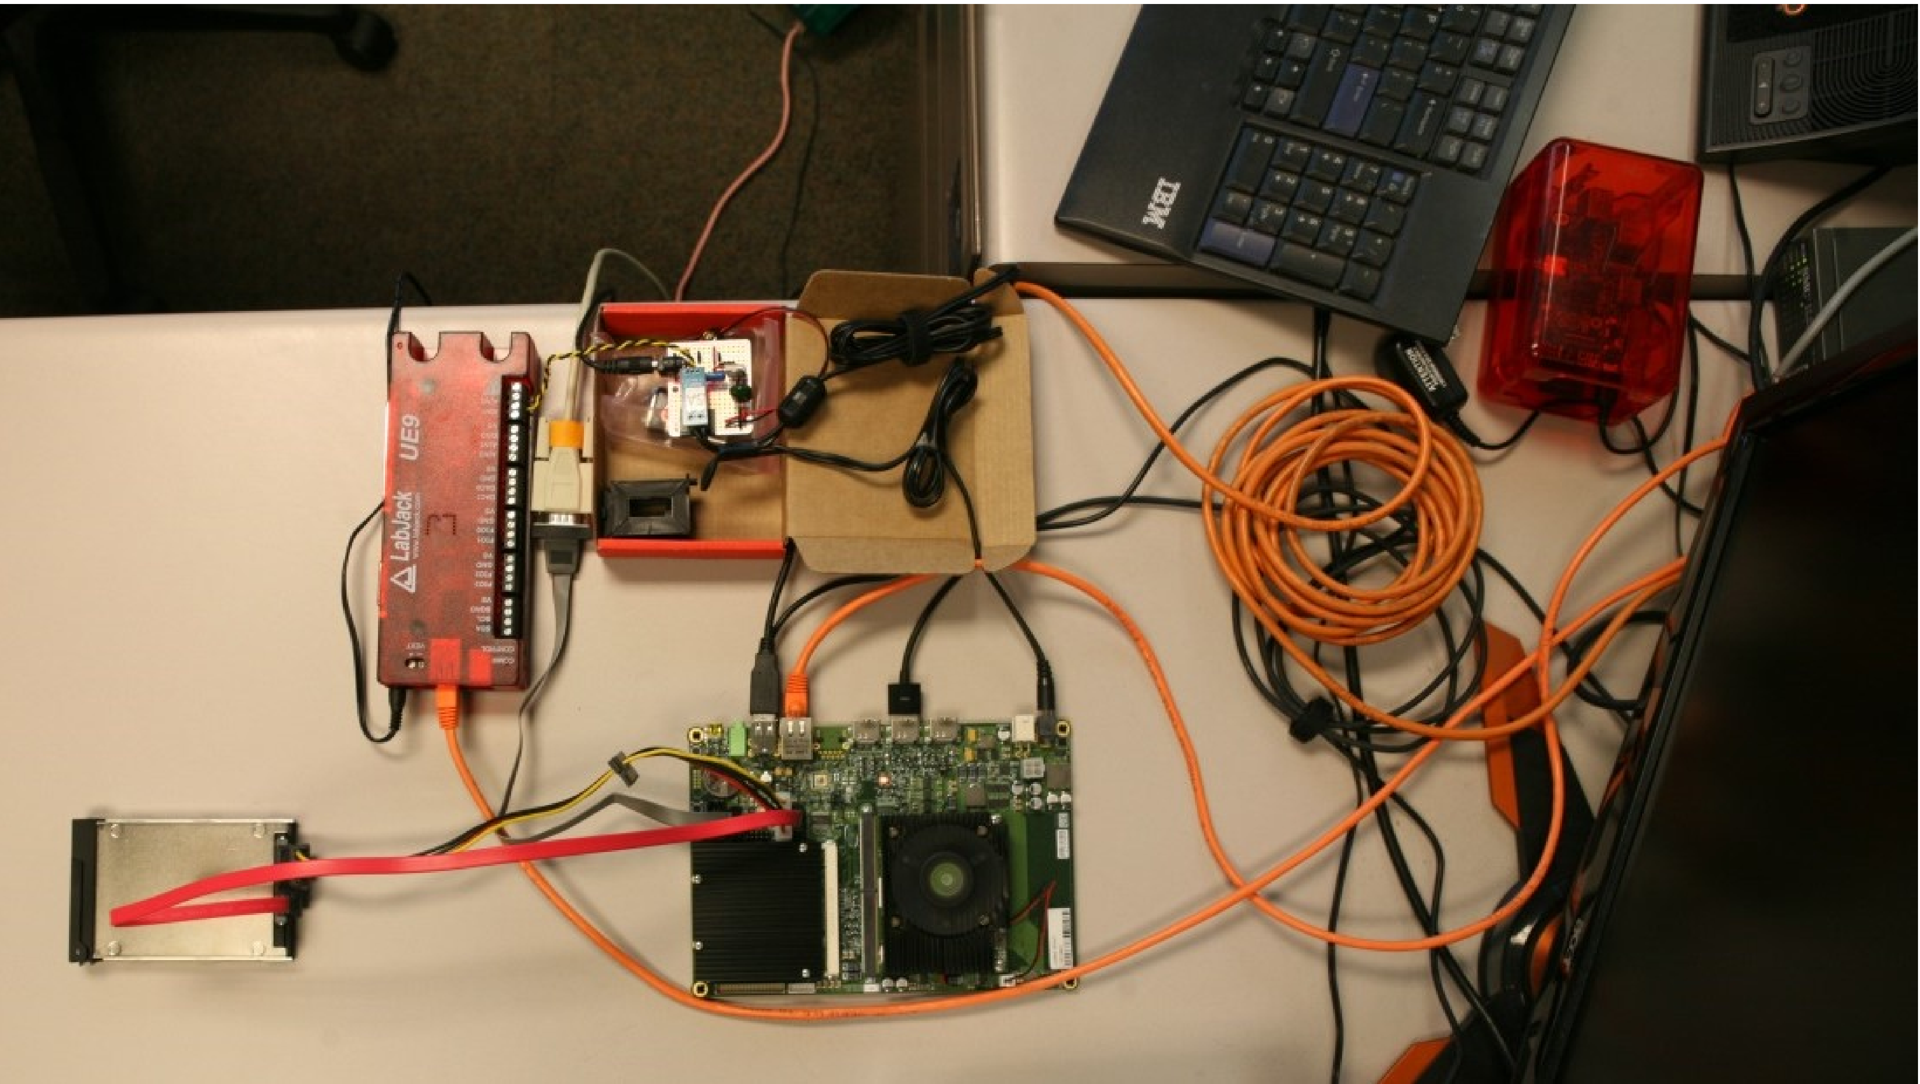

NAMD Simulation Performance

Platform	GPU	Time/step	Power	Steps/kJ	GPU Speedup
CARMA	Quadro 1000M	0.350 s	34 W	84	4.3x
KAYLA	GeForce Titan	0.283 s	93 W	38	5.9x
i7-3960X	GeForce Titan	0.0185 s	444 W	122	5.8x

- NAMD perf. and efficiency on ARM GPU platforms far below that of x86
- ARM+GPU system software lacks host-mapped GPU memory, preventing GPU from "streaming" output to host
- NAMD sensitive to CPU-GPU small-transfer overheads
- ARM+GPU platform CPU-GPU transfer perf. Is low…

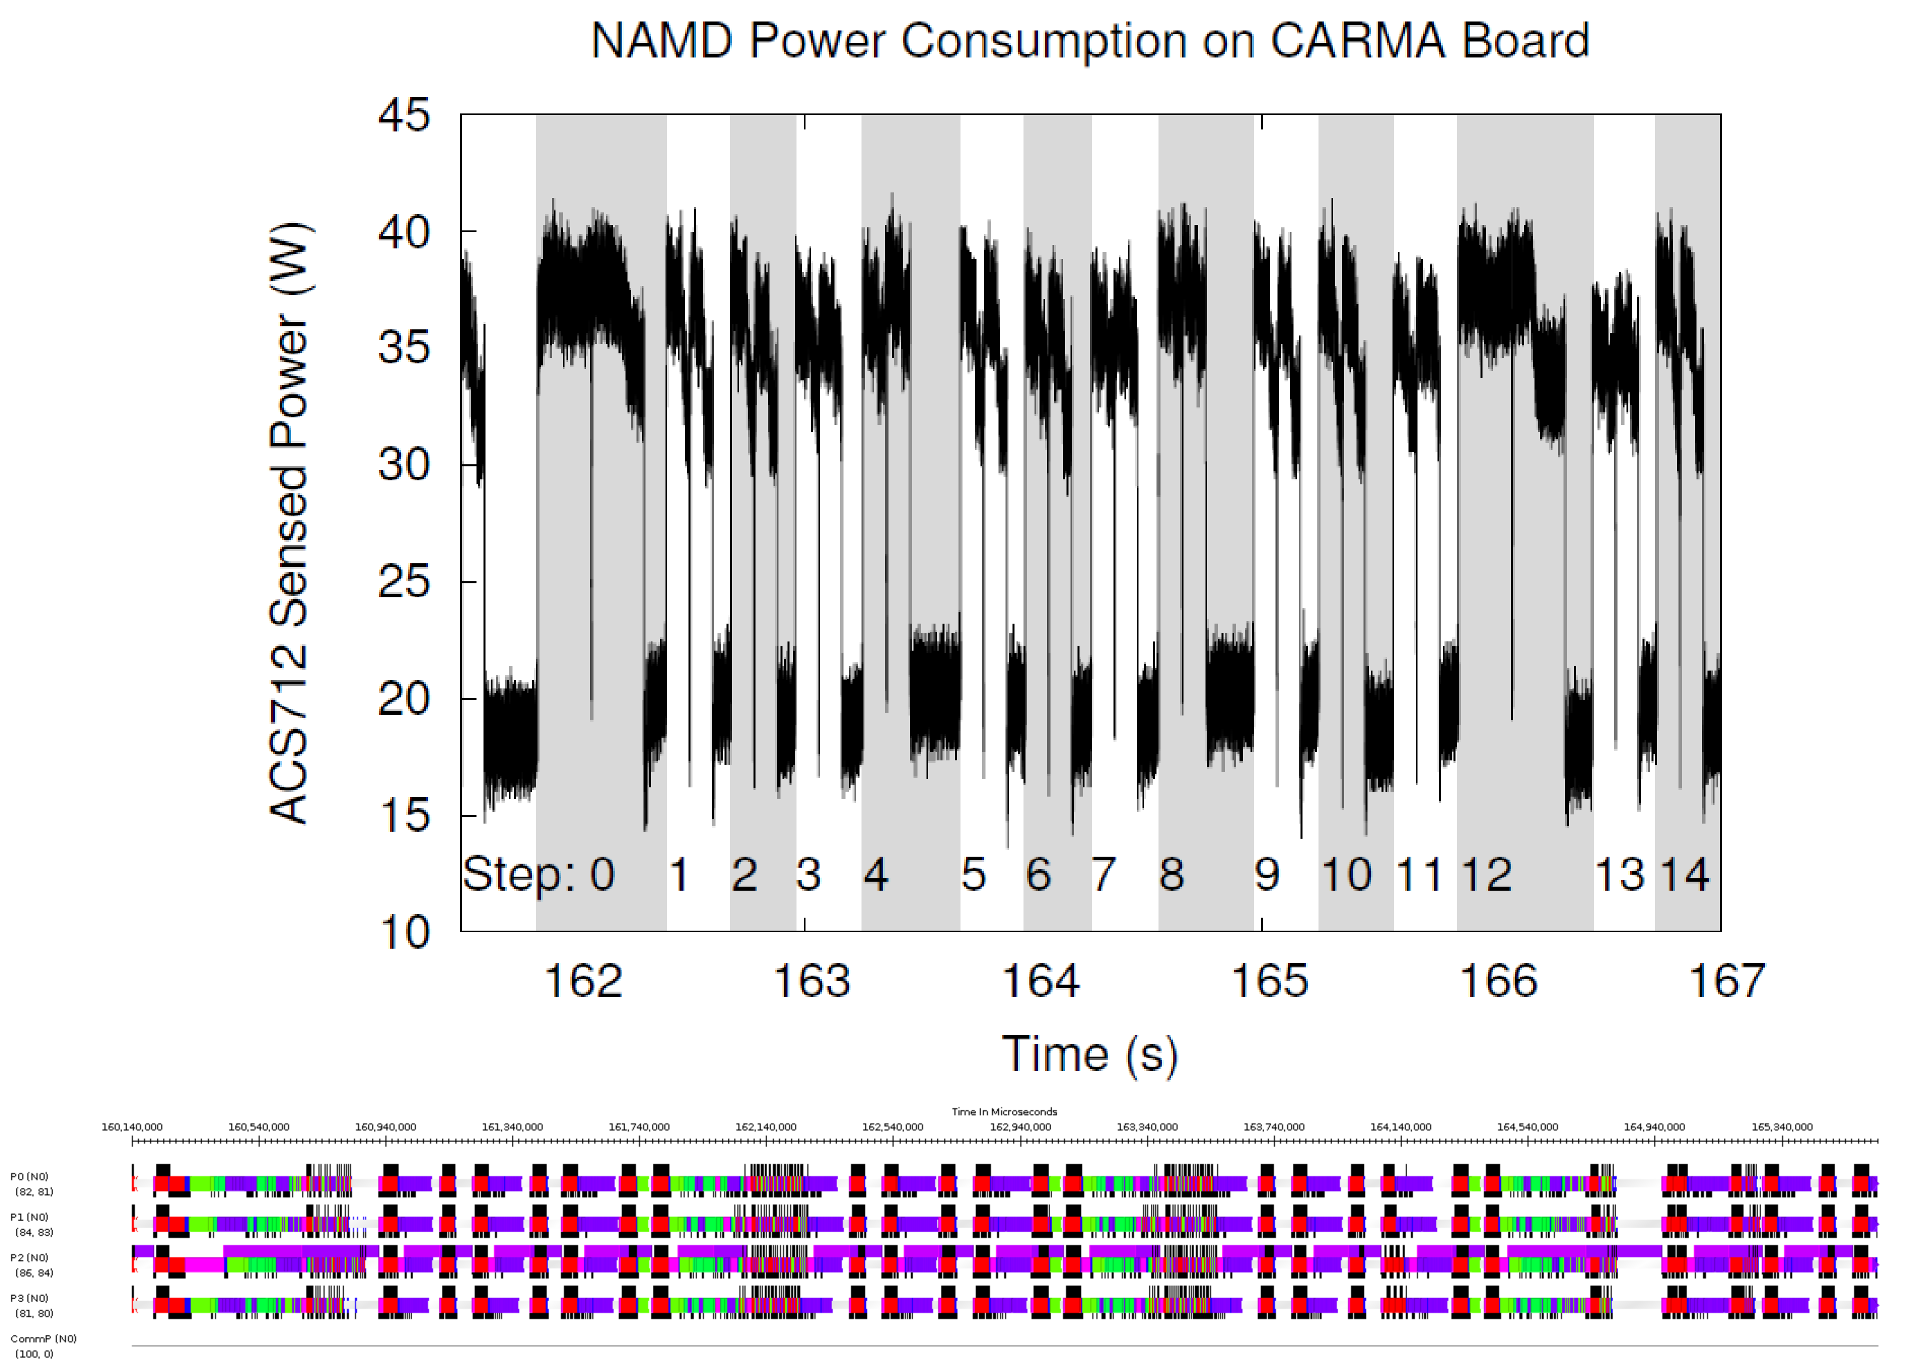
NAMD Power Consumption on CARMA Board
ACS712 Sensed Power (W)
45
40
35
30
25
20
15
10
Step: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14
162
163
164
165
166
167
Time (s)
Time In Microseconds
160,140,000
160,540,000
160,940,000
161,340,000
161,740,000
162,140,000
162,540,000
162,940,000
163,340,000
163,740,000
164,140,000
164,540,000
164,940,000
165,340,000
P0 (N0)
(82, 81)
P1 (N0)
(84, 83)
P2 (N0)
(86, 84)
P3 (N0)
(81, 80)
CommP (N0)
(100, 0)

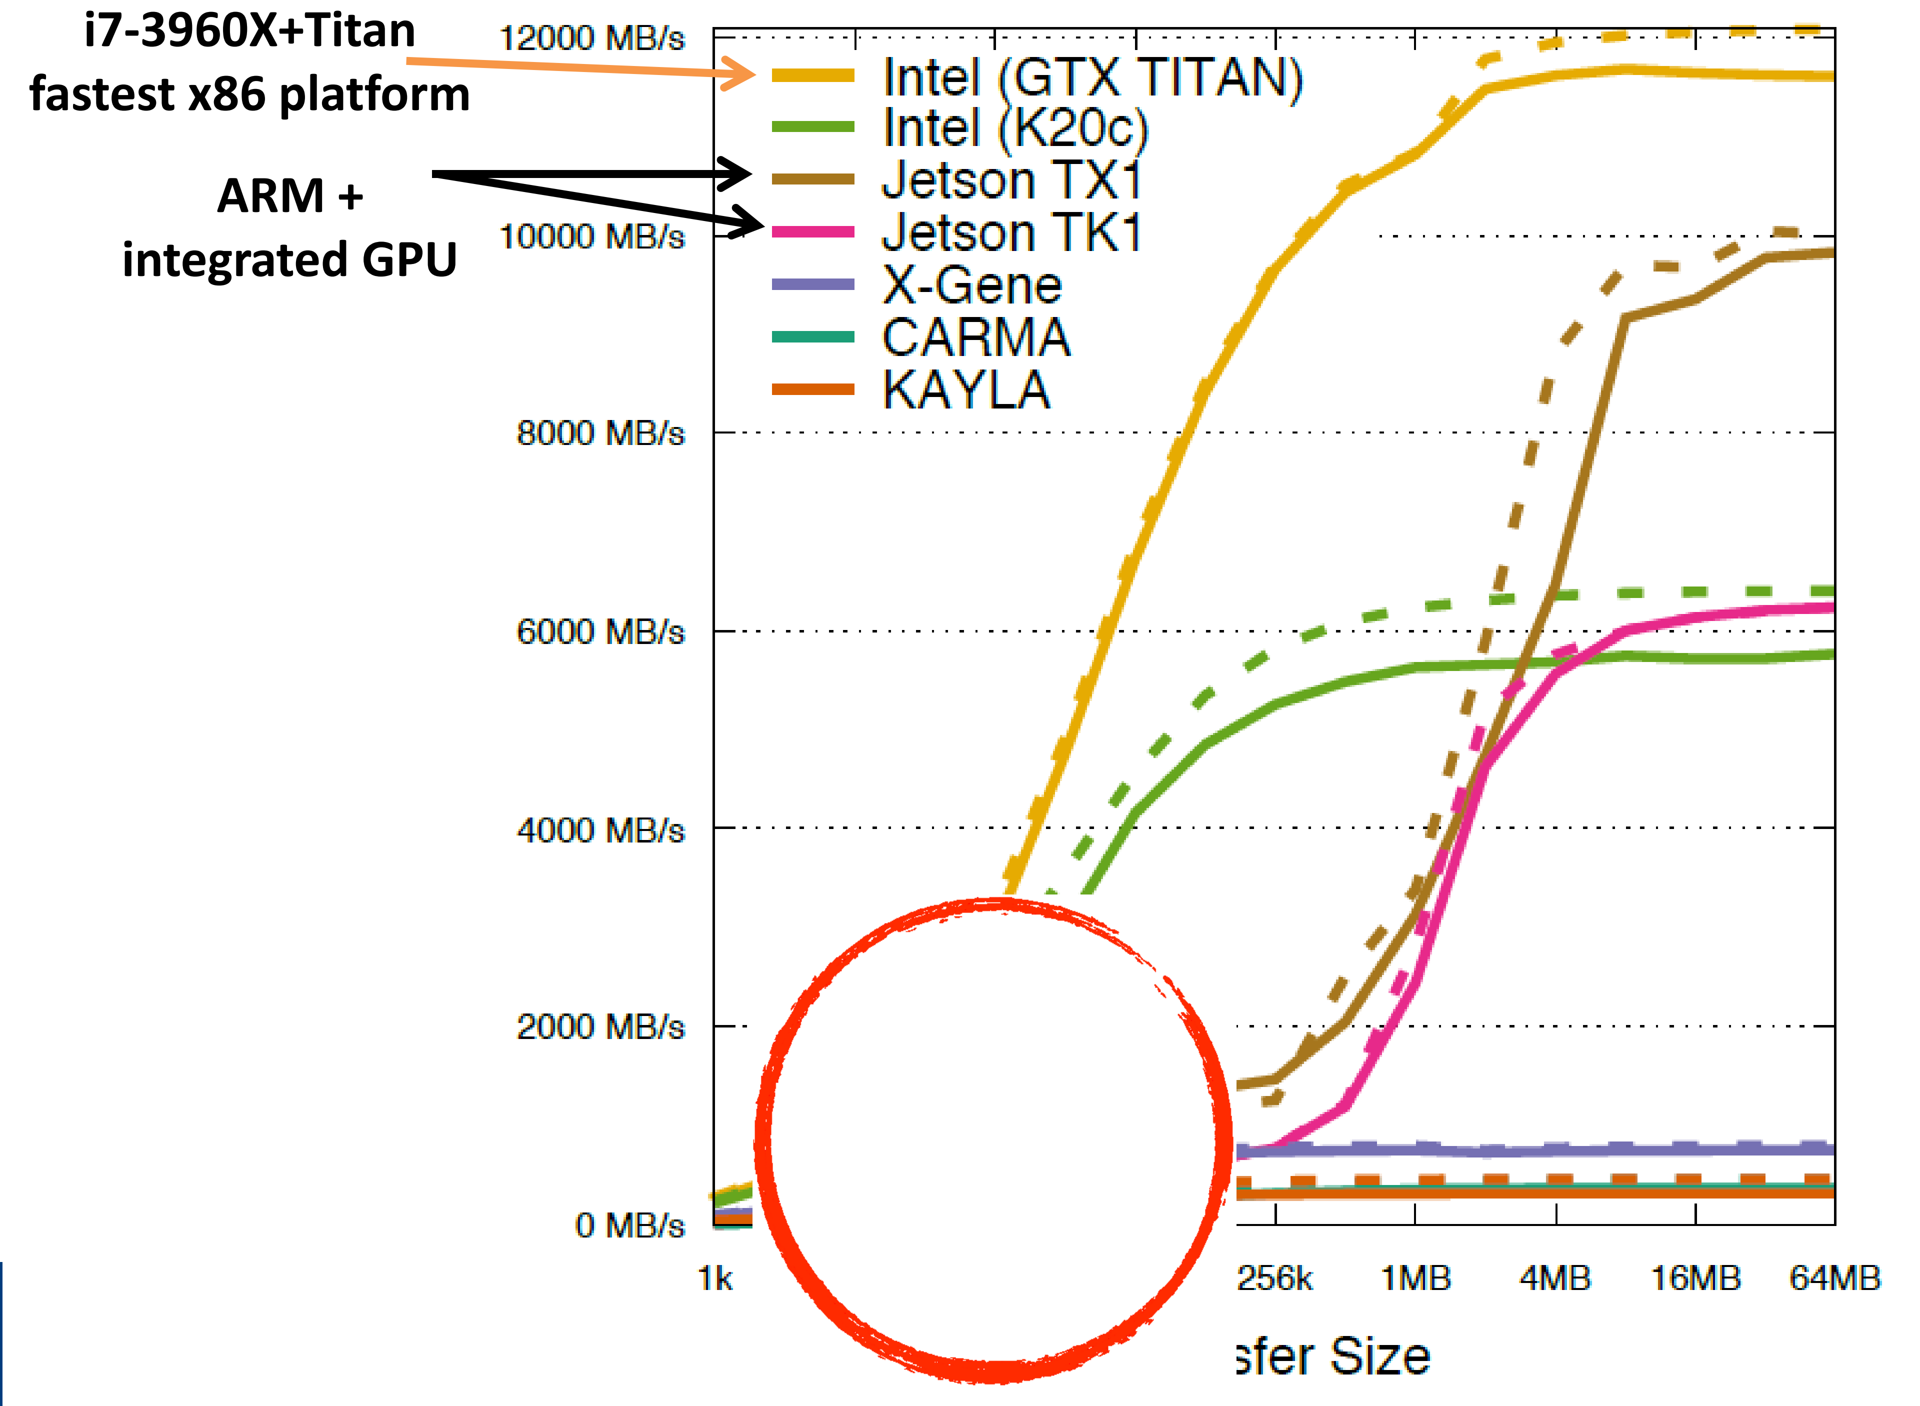

i7-3960X+Titan
fastest x86 platform
ARM +
integrated GPU
12000 MB/s
10000 MB/s
8000 MB/s
6000 MB/s
4000 MB/s
2000 MB/s
0 MB/s
Intel (GTX TITAN)
Intel (K20c)
Jetson TX1
Jetson TK1
X-Gene
CARMA
KAYLA
1k
256k
1MB
4MB
16MB
64MB
sfer Size

LM Simulation Performance

Lattice Size	Particles	Carma Steps/s	W	Efficiency	Kayla+ GeForce Titan Steps/s	W	Efficiency
32^3	2K	726	31W	23.4 steps/J	1304	137W	9.5 steps/J
128^3	256K	21.7	34W	0.64 steps/J	253	203W	1.2 steps/J
256^3	512K	3.0	35W	0.09 steps/J	40.4	212W	0.19 steps/J

X-Gene: Efficiency winner in all but one LM test case, perf is competitive with x86

Lattice Size	Particles	Intel i7-3960X+Tesla K20c Steps/s	W	Efficiency	APM X-Gene+Tesla K20c Steps/s	W	Efficiency
32^3	2K	5463	226W	24.2 steps/J	4638	142W	32.6 steps/J
128^3	256K	305	266W	1.2 steps/J	300	189W	1.6 steps/J
256^3	512K	48.3	270W	0.18 steps/J	47.7	195W	0.24 steps/J

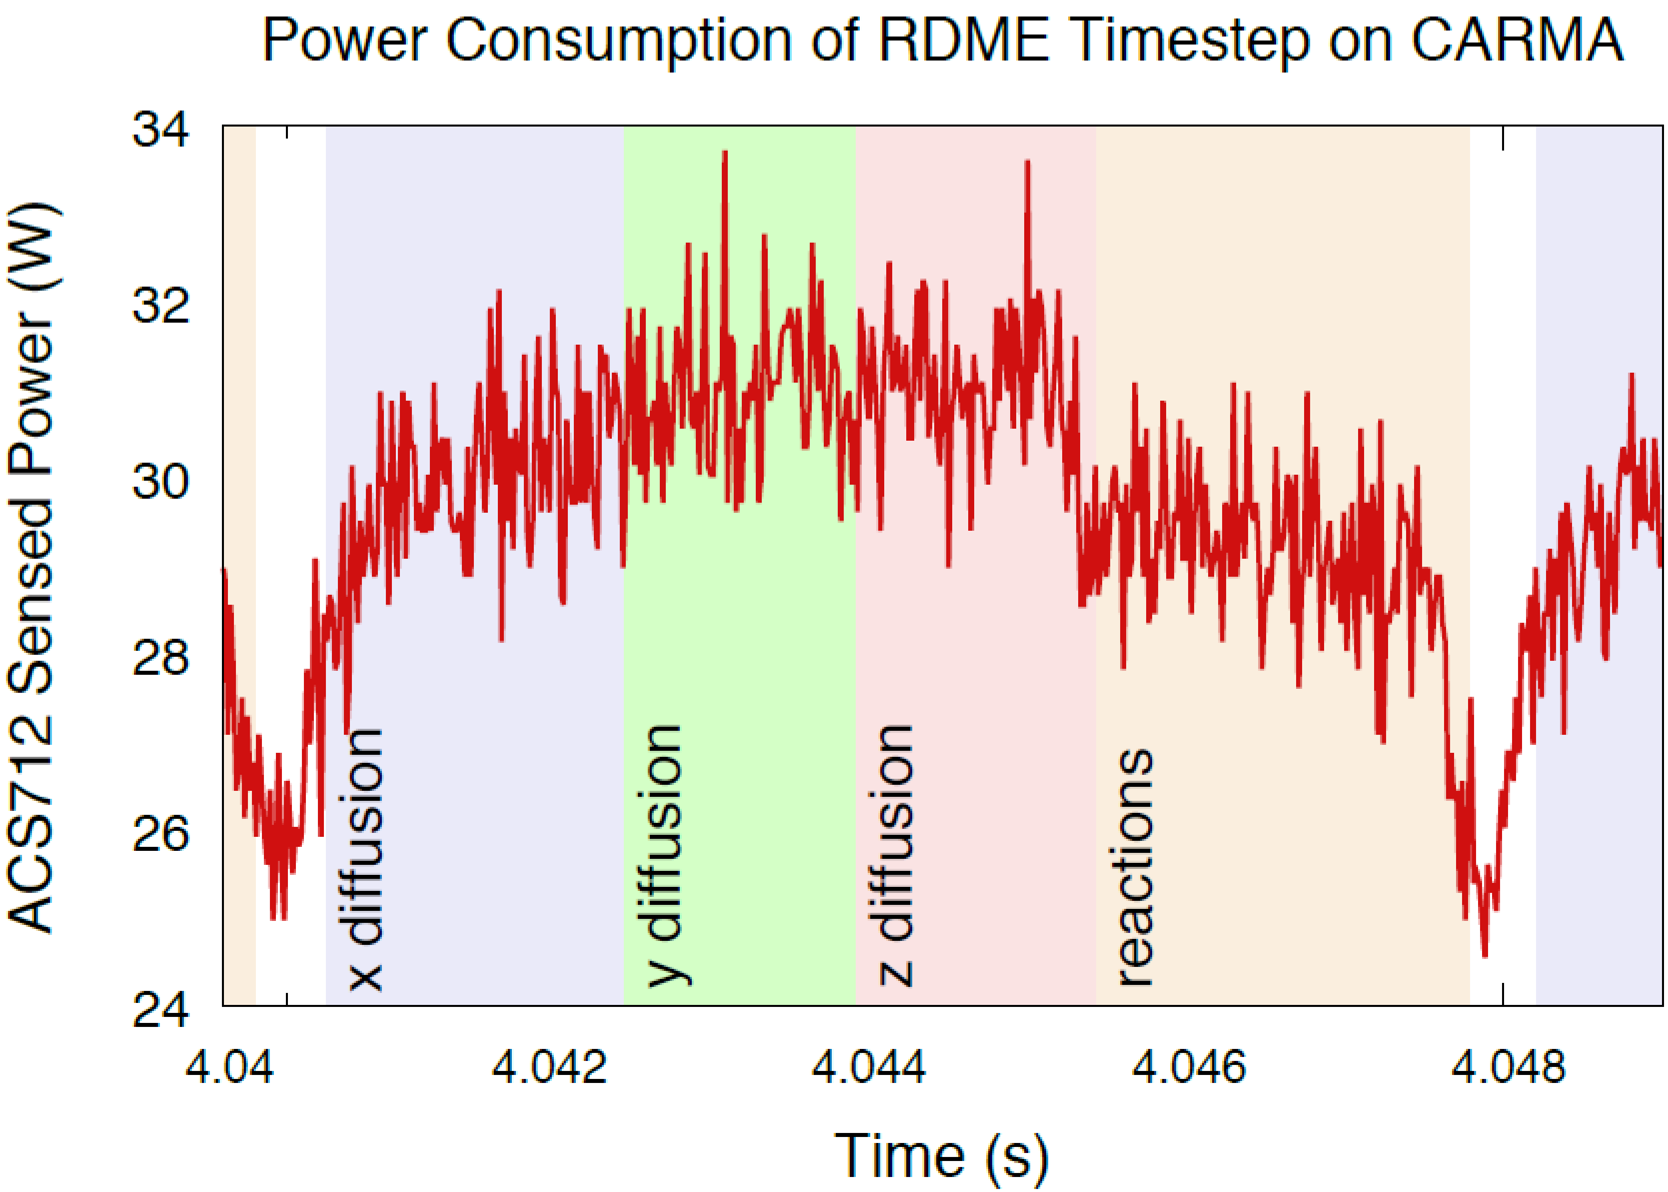
Power Consumption of RDME Timestep on CARMA
ACS712 Sensed Power (W)
34
32
30
28
26
24
x diffusion
y diffusion
z diffusion
reactions
4.04
4.042
4.044
4.046
4.048
Time (s)

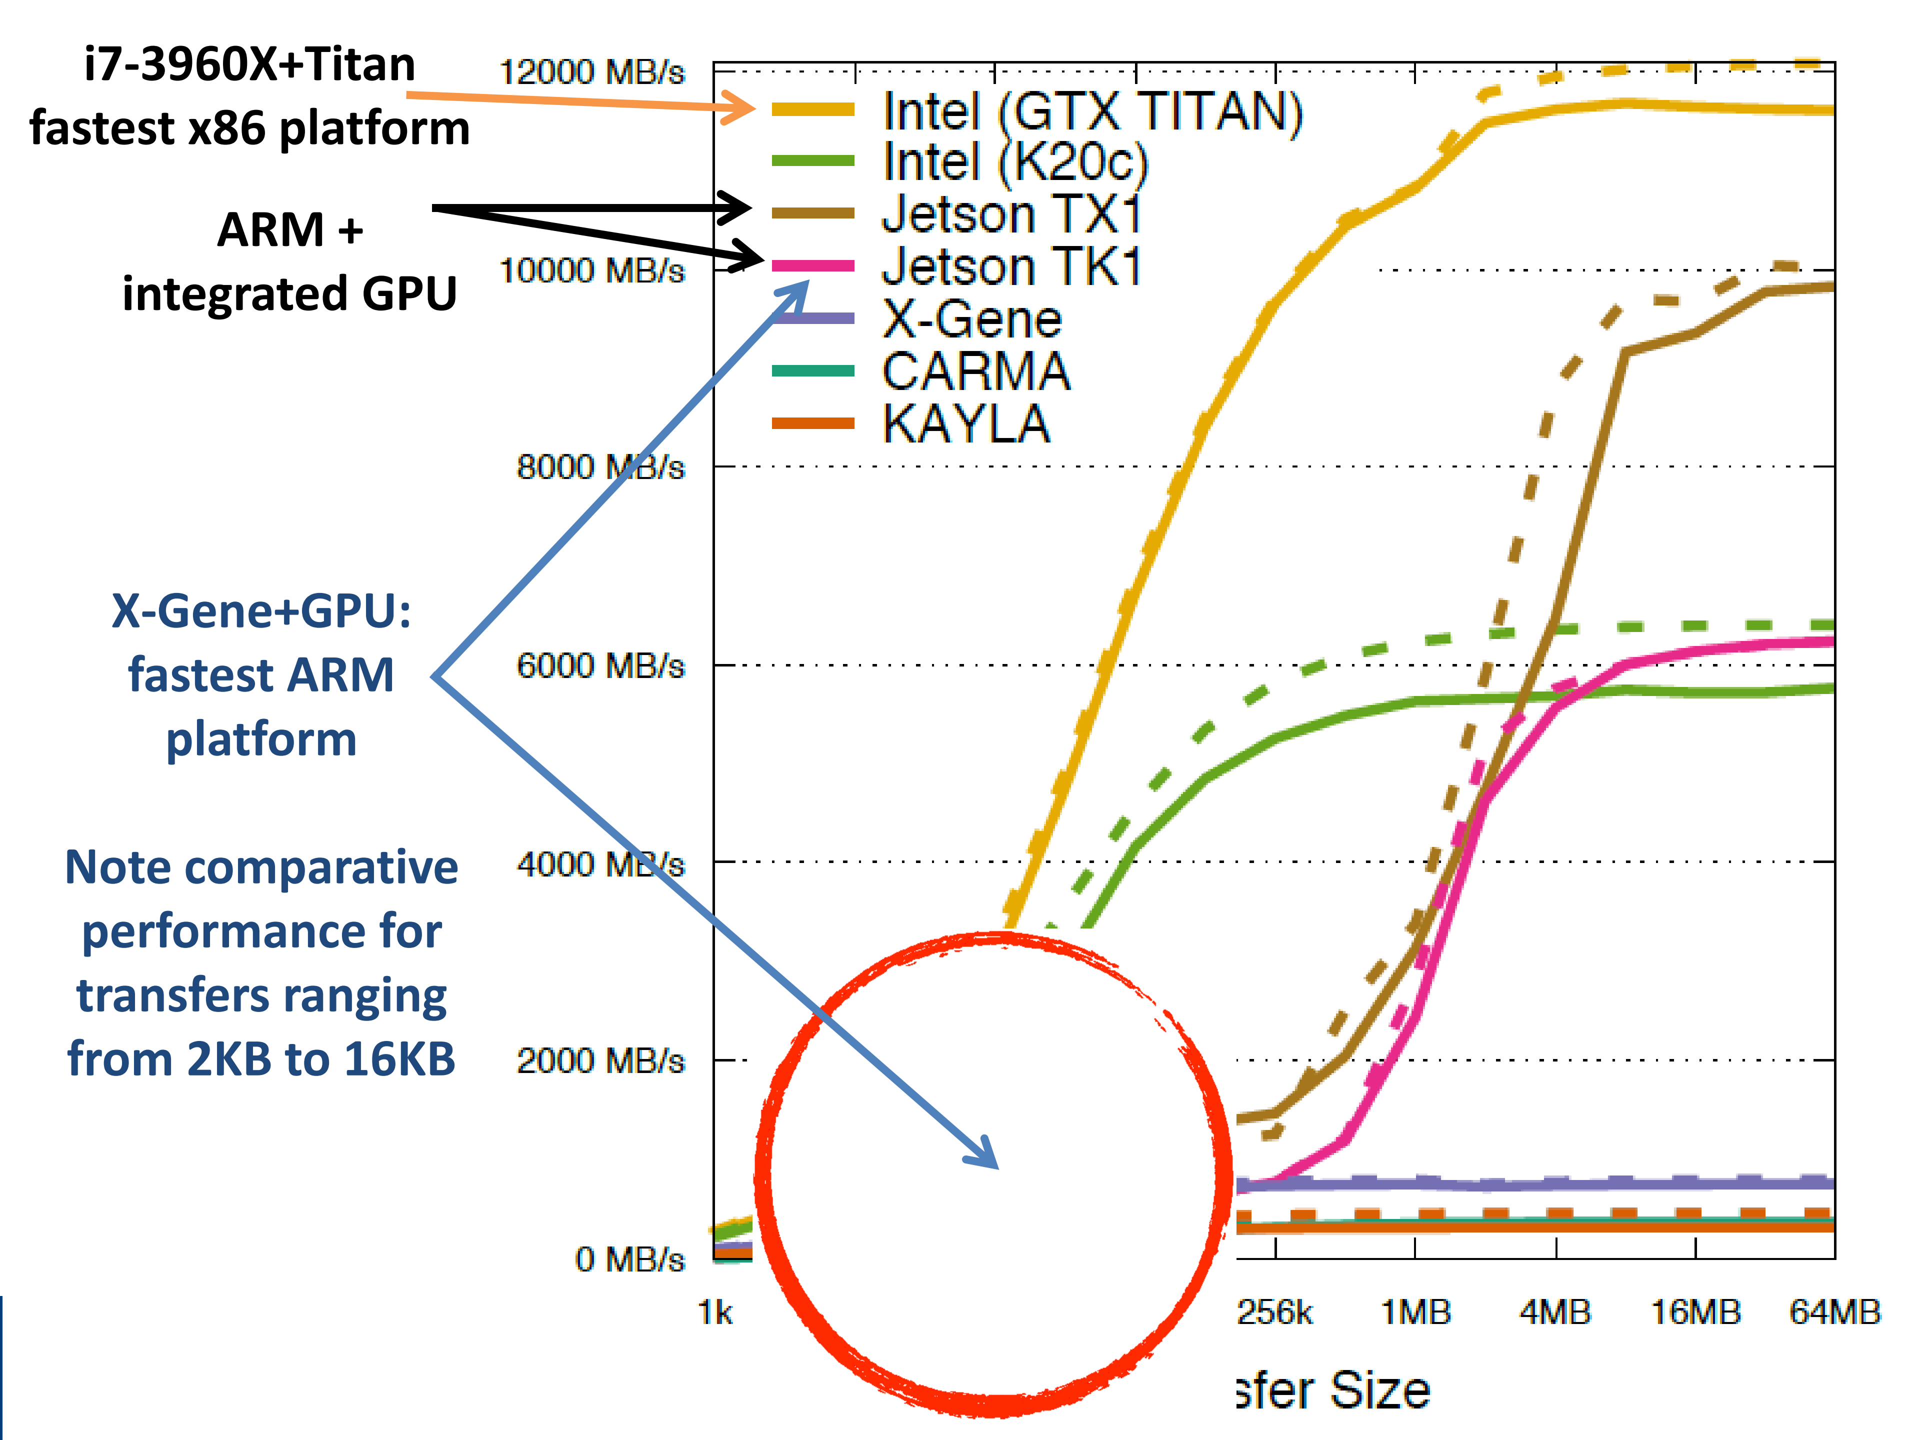

i7-3960X+Titan
fastest x86 platform
ARM +
integrated GPU
X-Gene+GPU:
fastest ARM
platform
Note comparative
performance for
transfers ranging
from 2KB to 16KB
12000 MB/s
10000 MB/s
8000 MB/s
6000 MB/s
4000 MB/s
2000 MB/s
0 MB/s
Intel (GTX TITAN)
Intel (K20c)
Jetson TX1
Jetson TK1
X-Gene
CARMA
KAYLA
1k
256k
1MB
4MB
16MB
64MB
sfer Size

Digging Deeper Into CPU-GPU Transfer Performance Issues

- ARM consistently underperformed vs. x86 comparison cases
- ARM PCIe interfaces run at lower rates: x4 or x8 vs. x16 on x86
- ARM perf. for small-sized transfers is very low, even w/ integrated GPUs:
  – ARM cores lack sophistication in single-thread code paths, no out-of-order, etc
  – ARM arch has complex procedure for VM ops, e.g., **TLB shootdown**; tree-like sync of all CPU cores
  – Lower clock rates
  – Driver stack less mature than x86

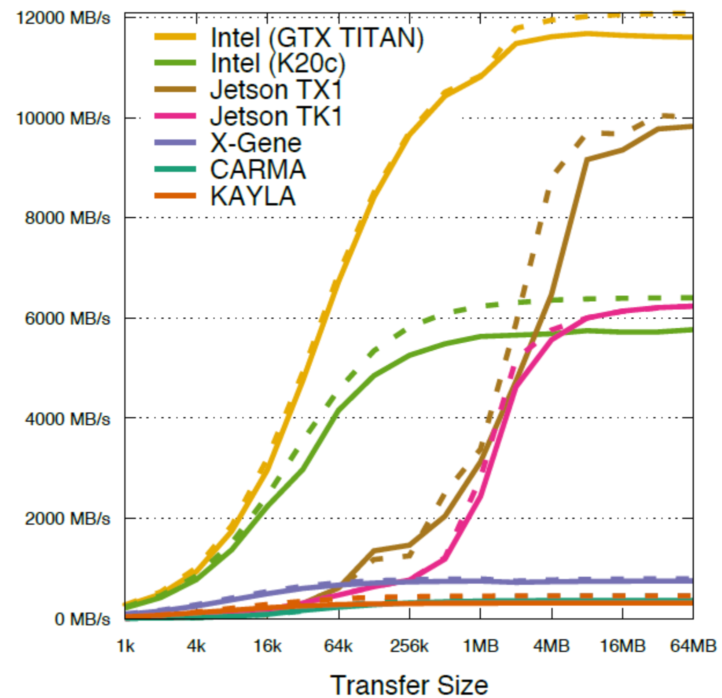

Improving VMD C_{60} MO Perf. on KAYLA ARM with Optimized CPU-GPU Transfers

C_{60} MO Algorithm	Perf. (FPS)	Power (W)	Energy Efficiency (frames/kJ)
Original	2.12 FPS	88 W	24 frames/kJ
New Transfer-optimized	3.89 FPS	89 W	49 frames/kJ

- Eliminate copy of MO wavefunction densities from GPU to CPU
- Perform intermediate marching cubes step in-place on GPU
- **1.8x performance increase on ARM**
- On x86 strategy is usually but not universally better, closer to break-even point, multi-GPU, etc.

Future Work

- Ongoing study of sources of PCIe CPU-GPU transfer overheads, schemes to mitigate performance loss on ARM or other platforms
- Direct link high-freq. power monitoring instrumentation to conventional performance instrumentation tools
- Develop new GPU algorithms that are tolerant of low-perf. CPU-GPU transfers: new kernels, even simple ones, that eliminate small transfers when possible
- Re-test platforms reported on here with new and improved compilers, kernels, drivers, an other system software expected to become available later this year
- Develop new low-cost instrumentation schemes that separate GPU and CPU power on commodity x86
- Compare w/ POWER8, other x86 platforms, Xeon Phi, etc.

Acknowledgements

- Theoretical and Computational Biophysics Group, University of Illinois at Urbana-Champaign
- NVIDIA GPU Center of Excellence, University of Illinois at Urbana-Champaign
- NVIDIA CUDA team
- NCSA
- Funding:
  - DOE INCITE, ORNL Titan: DE-AC05-00OR22725
  - NSF Blue Waters: NSF OCI 07-25070, PRAC “The Computational Microscope”, ACI-1238993, ACI-1440026
  - NIH support: 9P41GM104601, 5R01GM098243-02, PHS 5 T32 GM008276

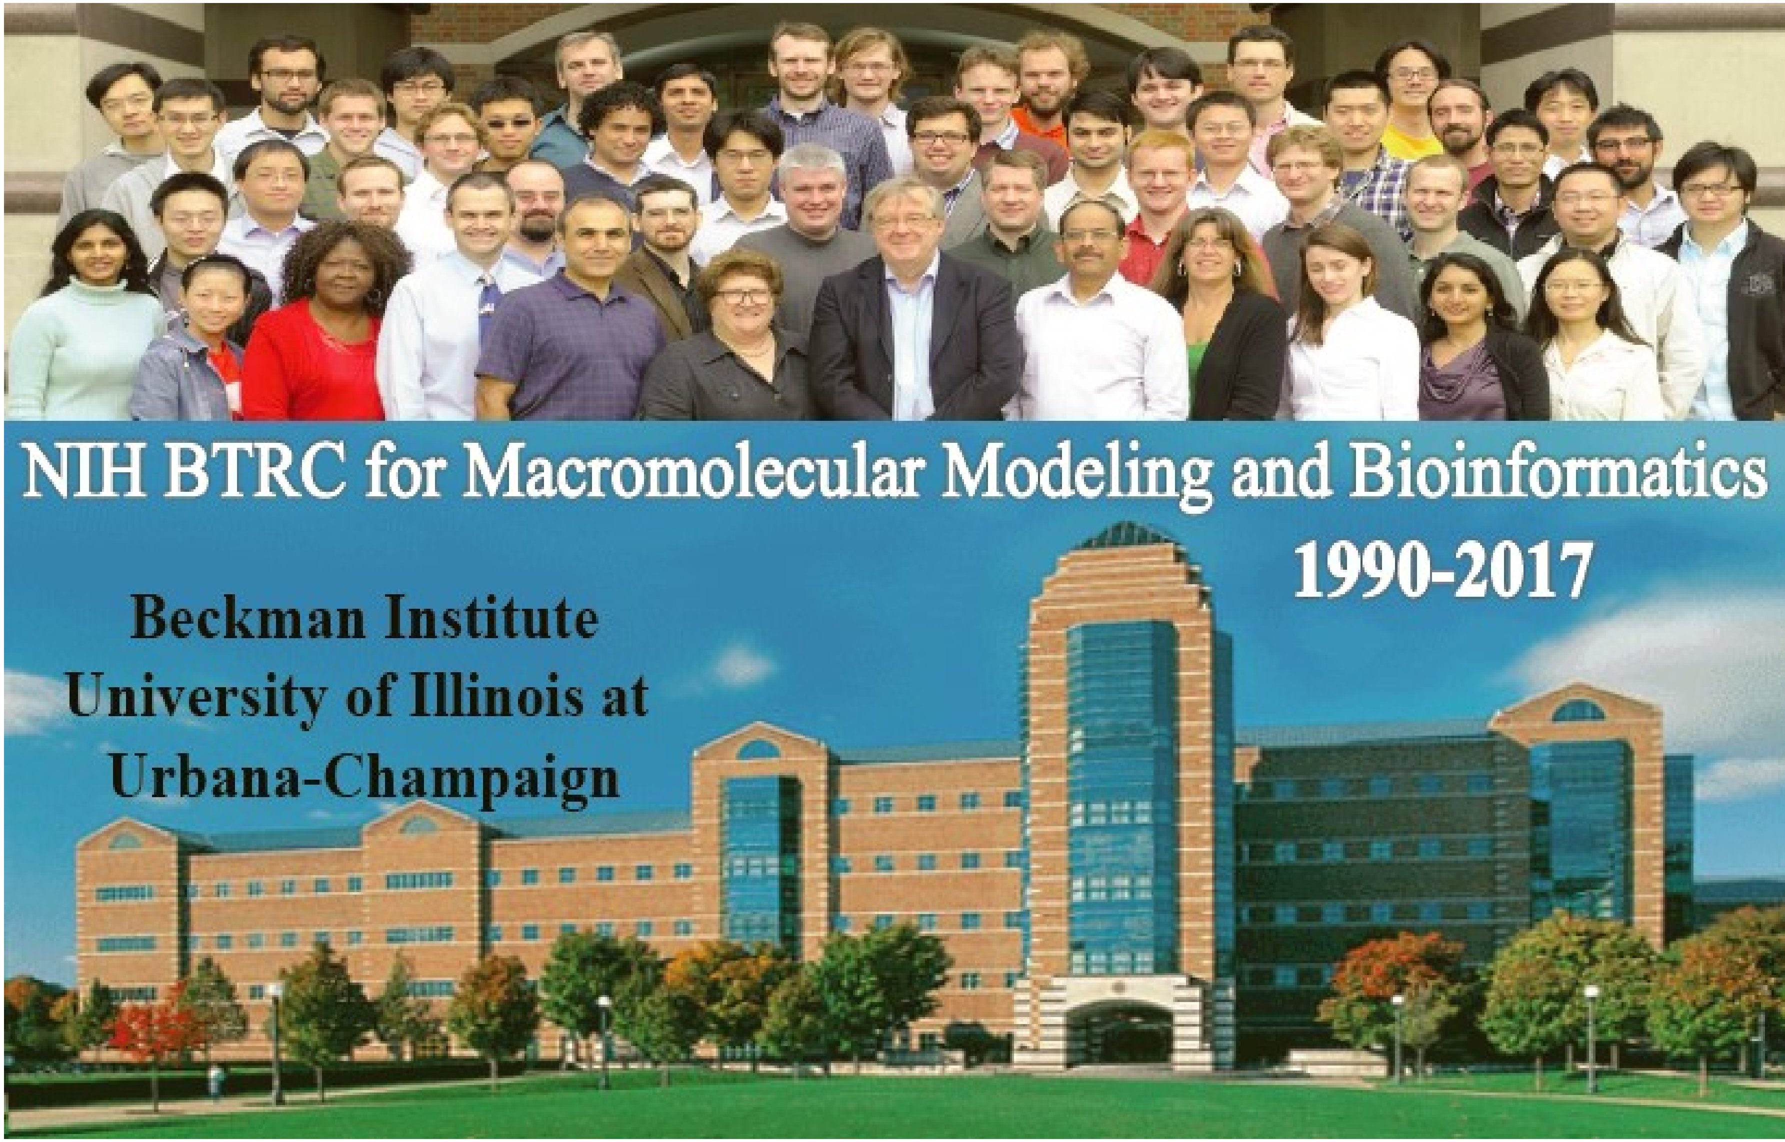
NIH BTRC for Macromolecular Modeling and Bioinformatics
1990-2017
Beckman Institute
University of Illinois at
Urbana-Champaign

Related Publications

http://www.ks.uiuc.edu/Research/gpu/

- **Immersive Molecular Visualization with Omnidirectional Stereoscopic Ray Tracing and Remote Rendering.** John E. Stone, William R. Sherman, and Klaus Schulten.High Performance Data Analysis and Visualization Workshop, IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW), 2016. **(In-press)**
- **High Performance Molecular Visualization: In-Situ and Parallel Rendering with EGL.** John E. Stone, Peter Messmer, Robert Sisneros, and Klaus Schulten.High Performance Data Analysis and Visualization Workshop, IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW), 2016. **(In-press)**
- **Evaluation of Emerging Energy-Efficient Heterogeneous Computing Platforms for Biomolecular and Cellular Simulation Workloads.** John E. Stone, Michael J. Hallock, James C. Phillips, Joseph R. Peterson, Zaida Luthey-Schulten, and Klaus Schulten.25th International Heterogeneity in Computing Workshop, IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW), 2016. **(In-press)**
- **Atomic Detail Visualization of Photosynthetic Membranes with GPU-Accelerated Ray Tracing.** J. E. Stone, M. Sener, K. L. Vandivort, A. Barragan, A. Singharoy, I. Teo, J. V. Ribeiro, B. Isralewitz, B. Liu, B.-C. Goh, J. C. Phillips, C. MacGregor-Chatwin, M. P. Johnson, L. F. Kourkoutis, C. Neil Hunter, and K. Schulten. J. Parallel Computing, 2016. **(In-press)**
- **Chemical Visualization of Human Pathogens: the Retroviral Capsids.** Juan R. Perilla, Boon Chong Goh, John E. Stone, and Klaus SchultenSC'15 Visualization and Data Analytics Showcase, 2015.

Related Publications
http://www.ks.uiuc.edu/Research/gpu/

- **Visualization of Energy Conversion Processes in a Light Harvesting Organelle at Atomic Detail.** M. Sener, J. E. Stone, A. Barragan, A. Singharoy, I. Teo, K. L. Vandivort, B. Isralewitz, B. Liu, B. Goh, J. C. Phillips, L. F. Kourkoutis, C. N. Hunter, and K. Schulten. SC'14 Visualization and Data Analytics Showcase, 2014.
  ******Winner of the SC'14 Visualization and Data Analytics Showcase***
- **Runtime and Architecture Support for Efficient Data Exchange in Multi-Accelerator Applications.** J. Cabezas, I. Gelado, J. E. Stone, N. Navarro, D. B. Kirk, and W. Hwu. IEEE Transactions on Parallel and Distributed Systems, 2014. **(In press)**
- **Unlocking the Full Potential of the Cray XK7 Accelerator.** M. D. Klein and J. E. Stone. Cray Users Group, Lugano Switzerland, May 2014.
- **GPU-Accelerated Analysis and Visualization of Large Structures Solved by Molecular Dynamics Flexible Fitting.** J. E. Stone, R. McGreevy, B. Isralewitz, and K. Schulten. Faraday Discussions, 169:265-283, 2014.
- **Simulation of reaction diffusion processes over biologically relevant size and time scales using multi-GPU workstations**. M. J. Hallock, J. E. Stone, E. Roberts, C. Fry, and Z. Luthey-Schulten. Journal of Parallel Computing, 40:86-99, 2014.

Related Publications

http://www.ks.uiuc.edu/Research/gpu/

- **GPU-Accelerated Molecular Visualization on Petascale Supercomputing Platforms.** J. Stone, K. L. Vandivort, and K. Schulten. UltraVis'13: Proceedings of the 8th International Workshop on Ultrascale Visualization, pp. 6:1-6:8, 2013.
- **Early Experiences Scaling VMD Molecular Visualization and Analysis Jobs on Blue Waters.** J. Stone, B. Isralewitz, and K. Schulten. In proceedings, Extreme Scaling Workshop, 2013.
- **Lattice Microbes: High-performance stochastic simulation method for the reaction-diffusion master equation.** E. Roberts, J. Stone, and Z. Luthey-Schulten. J. Computational Chemistry 34 (3), 245-255, 2013.
- **Fast Visualization of Gaussian Density Surfaces for Molecular Dynamics and Particle System Trajectories.** M. Krone, J. Stone, T. Ertl, and K. Schulten. *EuroVis Short Papers,* pp. 67-71, 2012.
- **Immersive Out-of-Core Visualization of Large-Size and Long-Timescale Molecular Dynamics Trajectories.** J. Stone, K. L. Vandivort, and K. Schulten. G. Bebis et al. (Eds.): *7th International Symposium on Visual Computing (ISVC 2011)*, LNCS 6939, pp. 1-12, 2011.
- **Fast Analysis of Molecular Dynamics Trajectories with Graphics Processing Units – Radial Distribution Functions.** B. Levine, J. Stone, and A. Kohlmeyer. *J. Comp. Physics*, 230(9):3556-3569, 2011.

Related Publications

http://www.ks.uiuc.edu/Research/gpu/

- **Quantifying the Impact of GPUs on Performance and Energy Efficiency in HPC Clusters.** J. Enos, C. Steffen, J. Fullop, M. Showerman, G. Shi, K. Esler, V. Kindratenko, J. Stone, J Phillips. International Conference on Green Computing, pp. 317-324, 2010.
- **GPU-accelerated molecular modeling coming of age.** J. Stone, D. Hardy, I. Ufimtsev, K. Schulten. J. Molecular Graphics and Modeling, 29:116-125, 2010.
- **OpenCL: A Parallel Programming Standard for Heterogeneous Computing.** J. Stone, D. Gohara, G. Shi. Computing in Science and Engineering, 12(3):66-73, 2010.
- **An Asymmetric Distributed Shared Memory Model for Heterogeneous Computing Systems**. I. Gelado, J. Stone, J. Cabezas, S. Patel, N. Navarro, W. Hwu. *ASPLOS '10:* Proceedings of the 15th International Conference on Architectural Support for Programming Languages and Operating Systems, pp. 347-358, 2010.

Related Publications
http://www.ks.uiuc.edu/Research/gpu/

- **GPU Clusters for High Performance Computing**. V. Kindratenko, J. Enos, G. Shi, M. Showerman, G. Arnold, J. Stone, J. Phillips, W. Hwu. Workshop on Parallel Programming on Accelerator Clusters (PPAC), In Proceedings IEEE Cluster 2009, pp. 1-8, Aug. 2009.
- **Long time-scale simulations of in vivo diffusion using GPU hardware**. E. Roberts, J. Stone, L. Sepulveda, W. Hwu, Z. Luthey-Schulten. In IPDPS'09: Proceedings of the 2009 IEEE International Symposium on Parallel & Distributed Computing, pp. 1-8, 2009.
- **High Performance Computation and Interactive Display of Molecular Orbitals on GPUs and Multi-core CPUs**. J. Stone, J. Saam, D. Hardy, K. Vandivort, W. Hwu, K. Schulten, 2nd Workshop on General-Purpose Computation on Graphics Pricessing Units (GPGPU-2), ACM International Conference Proceeding Series, volume 383, pp. 9-18, 2009.
- **Probing Biomolecular Machines with Graphics Processors**. J. Phillips, J. Stone. Communications of the ACM, 52(10):34-41, 2009.
- **Multilevel summation of electrostatic potentials using graphics processing units**. D. Hardy, J. Stone, K. Schulten. J. Parallel Computing, 35:164-177, 2009.

Related Publications

http://www.ks.uiuc.edu/Research/gpu/

- **Adapting a message-driven parallel application to GPU-accelerated clusters**. J. Phillips, J. Stone, K. Schulten. Proceedings of the 2008 ACM/IEEE Conference on Supercomputing, IEEE Press, 2008.
- **GPU acceleration of cutoff pair potentials for molecular modeling applications**. C. Rodrigues, D. Hardy, J. Stone, K. Schulten, and W. Hwu. Proceedings of the 2008 Conference On Computing Frontiers, pp. 273-282, 2008.
- **GPU computing**. J. Owens, M. Houston, D. Luebke, S. Green, J. Stone, J. Phillips. Proceedings of the IEEE, 96:879-899, 2008.
- **Accelerating molecular modeling applications with graphics processors**. J. Stone, J. Phillips, P. Freddolino, D. Hardy, L. Trabuco, K. Schulten. J. Comp. Chem., 28:2618-2640, 2007.
- **Continuous fluorescence microphotolysis and correlation spectroscopy**. A. Arkhipov, J. Hüve, M. Kahms, R. Peters, K. Schulten. Biophysical Journal, 93:4006-4017, 2007.