Breaking 100TFlops at 3KW Power with IBM Power9 and NVIDIA V100 GPUs over NVLink and 200GbE Mesh Interconnect

Ioan Raicu[*‡+], William Scullin[^], Ben Allen[^], Kyle Hale[*], Kyle Chard[‡#], Alexandru Iulian Orhean[*]

[*]Department of Computer Science, Illinois Institute of Technology, Chicago IL, USA
[‡]Mathematics and Computer Science Division, Argonne National Laboratory, Argonne IL, USA
[^]Advanced Leadership Computing Facility, Argonne National Laboratory, Argonne IL, USA
[#]Computation Institute, University of Chicago, Chicago IL, USA

iraicu@cs.iit.edu, {wscullin,bsallen}@alcf.anl.gov, khale@cs.iit.edu, chard@uchicago.edu, aorhean@hawk.iit.edu

***Abstract*—** *To achieve the next level of performance from large scale scientific computing, we have been exploring ways to optimize cluster utilities and applications to maximize performance, while minimizing power usage. To do this, we investigated automated ways to manage cluster power consumption through management of processor per-core frequencies (both to over boost and under clock), water cooling, eliminating fans, powering down of accelerators when not in use, and compiler level optimizations. The Chicago Fusion Team is composed of six mentors and six students from various Chicago area institutions, such as Illinois Institute of Technology (IIT), University of Chicago (UChicago), Argonne National Laboratory (ANL), and Adlai E. Stevenson High School (Stevenson). We propose a cluster configuration leveraging three nodes and a direct mesh interconnect, enclosed in a 6U rack. The rack will be water cooled, and contain two IBM Power9 processors per node (22-cores per socket), four NVIDIA V100 NVLink GPUs per node, NVMe SSDs, and Mellanox ConnectX6 200GbE network interconnect. We estimate that this 6U rack will deliver over 150 double precision TFlops/sec with less than 6kW of power. With aggressive power management and compiler optimizations, this hybrid cluster will fit in the 3kW power envelope and achieve nearly 100 TFlops/sec. Due to its variety of hardware, we expect it will be extremely flexible to efficiently support a wide range of real applications (beyond Linpack). Sponsorship will be from IBM, Argonne National Laboratory, and Illinois Institute of Technology.*

I. Chicago Fusion Team

A. Mentors

The new Chicago Fusion Team has 6 mentors from various Chicago area institutions, such as Illinois Institute of Technology (IIT), University of Chicago (UChicago), and Argonne National Laboratory (ANL); these mentors include faculty members who work in scientific computing, data science, resource management at extreme scales, distributed storage systems and run-time systems, parallel languages, operating systems, and virtual machines, as well as system administrators who run the Advanced Leadership Computing Facilities (ALCF) at ANL.

Ioan Raicu is an Associate Professor in CS at IIT, a guest research faculty in MCS at ANL, as well as a guest faculty in EECS at Northwestern. He received his PhD from UChicago in 2009. He is the recipient of the NSF/CRA CIFellow and the NSF CAREER awards. He is the PI of a NSF REU site engaging undergraduate students since 2015 in the BigDataX summer program. He is also the PI for the NSF-funded Mystic testbed in reconfigurable computing. His research interests are in distributed systems, emphasizing large-scale resource management in supercomputing, cloud computing, and many-core computing. He will serve as the team primary advisor.

William Scullin is a member of the Catalyst team at the ALCF at ANL. William is a computational generalist who enables scientific discovery through systems and software engineering at scale. With a strong background in both systems administration and computational science, he helps the team find resources and guides their research.

Ben Allen is the lead administrator for the Joint Laboratory for Systems Evaluation at the ALCF at ANL. He designs and maintains an environment that serves as a testbed for new and interesting hardware and software and supports HPC application and performance engineers using it. Ben has been key in helping the team realize their hardware plans in a safe and sane manner.

Kyle C. Hale is an Assistant Professor in CS at IIT. His research interests lie at the intersection of operating systems, parallel computing, and computer architecture, with a focus on building experimental systems. His work ranges from Networks-on-Chip to specialized operating systems and parallel languages. He earned his Ph.D. in CS from Northwestern University in 2016. He will help tune the OS parameters for best performance, as well as explore the potential use of research light-weight OS for the Intel Xeon Phi.

Kyle Chard is a Senior Researcher and Fellow in the Computation Institute at UChicago and ANL. His research focuses on applying computational and data-intensive approaches to solve scientific problems. His background has primarily focused on distributed systems, related to scheduling of resources in Grid and Cloud

environments. He received his PhD from Victoria University of Wellington in 2011. He will mentor students in data-management related topics.

Alexandru I. Orhean is a 3rd year PhD student in the DataSys laboratory at IIT. His research interests are in distributed systems, with a focus on data management, indexing, and search in large-scale data repositories. He will mentor students on a daily basis with hands-on activities.

B. Students

The mentors have identified 6 students to participate in the initial summer program to prepare for this competition, from two institutions: Illinois Institute of Technology (IIT) in Chicago Illinois and Adlai E. Stevenson High School (Stevenson) in Lincolnshire, IL. The students are ordered below by seniority.

Alexander Ballmer (IIT) is a senior in CS at IIT. He participated in the IIT's SCC team in 2014, 2015, and 2017. He is a CAMRAS scholar with a full ride scholarship. He has worked in the DataSys lab since June 2014. He has interned at Argonne National Laboratory in 2015 and 2016. He has attended the SC conference in 2014, 2015, 2016, and 2017. His interests in research are distributed systems, HPC file systems, and peer-to-peer networking. Other interests include building R/C aircraft, model rocketry, and filmmaking. He has interests in pursuing a Master degree in CS.

Parker Stephen Joncus (IIT) is a junior in Applied Mathematics Data Science at IIT. He has interests in pursuing a PhD degree in Data Science. He has worked with many programming languages such as Java, Python, Matlab, C, R, SQL, and Macaulay2. He is eager to learn new ideas and concepts, how to apply them, and how to program them.

Keith Alex Bateman (IIT) is a junior in CS at IIT. He has strong algorithms and coding foundations which he tested through his participation in the ACM ICPC competitions in 2017 and 2018. He is fiercely self-motivated and is determined to make an impact. He has interests in pursuing graduate school.

Rohit Mandava (Stevenson) is junior at Stevenson High. He has received numerous awards such as a perfect composite ACT score, Gold Medal Recipient at History Fair State Finals, and participated in the USAPHO (U.S.A Physics Olympiad). He is proficient in both Java and Python programming languages. He loves to play basketball and run track & field.

Katherine Zeng (Stevenson) is junior at Stevenson High. She goes by the name of Kitty, how fun is that? She is proficient in Java and Python, has strong algebra and calculus background, physics background, language arts and composition, as well as creative writing. She has participated in several Science Olympiads from 2016 to 2018, and won several 1st places at both regional and state competitions.

Andrew J. Gan (Stevenson) is junior at Stevenson High. He has participated in various national competitions, such as Math Madness High School National Championship (1st place), TEAMS (Tests of Engineering Aptitude, Mathematics, and Science) National Competition, Science Olympiad National Tournament, and the Harvard-MIT Math Tournament. His skillsets include C++, Python, Java, JavaScript, Hadoop, R, Express, MongoDB, and Calculus.

C. Prior Participation

IIT had three teams that competed in 2014 [4, 7], 2015 [5, 6, 8], and 2016 [cite]. The IIT team came in 9th (from 12 teams in 2014), 4th (from 9 teams in 2015), and 6th (from 16 teams in 2017). With all the lessons learned from the prior three competitions, the new Chicago Fusion Team that spans 4 institutions aims for 1st place. Except for IIT, none of the other institutions have had teams compete in SCC in the past.

Some of the lessons we learned from our prior competitions are:

1. Having the entire student team working on preparing for the SCC competition during the summer is critical due to the likely avalanche of student commitments when the Fall semester/quarters begin
2. Having the actual hardware from the vendors available to the students starting as early as the summer is critical to intimately understanding the hardware and how to tune it; we plan to make heavy use of the NSF funded Chameleon and Mystic testbeds prior to receiving the competition hardware
3. Having additional mentors specializing in various parts of the HPC software stack is important to best tune these applications on a variety of different hardware

Figure 1: IIT Chicago Fusion teams from 2014, 2015, and 2017 at SCC@SC [5, 6, 8]

II. Hardware and Software

A. Hardware

We propose a cluster configuration leveraging three nodes, enclosed in a 6U rack. Each node will have dual socket IBM Power9 processors (22 cores @ 3.1 GHz), 384GB RAM, 1.5TB Intel Optane NVMe SSDs (split over 4 drives), and a 200GbE switchless meshed interconnect. One of the nodes will be a multi-purpose node, that will serve as the head node and resource manager node. We estimate that this 6U rack will deliver approximately 150 double precision TFlops/sec with about 6kW of power.

We will use aggressive power management techniques as well as research-rooted compiler optimizations to reach these targets while using as much of the hardware as possible. We will use water cooling to reduce cooling power costs, and we will use redundant power supplies to load balance power consumption across the two PDUs. Given the 6x IBM Power9 CPUs and 18x V100 GPUs, the system will be extremely flexible and useful for a wide range of real applications (beyond Linpack). We expect to be able to keep the power under control to not exceed the maximum 3.1kW during the competition with aggressive power throttling.

Some interesting information about our Power9 systems, which lead us to choose IBM Power9 over Intel Xeon and AMD Epyc processors. The Power9 architecture has faster I/O, up to 5.6x more I/O bandwidth than x86 servers. The AC922 includes a variety of next-generation I/O architectures, including: PCIe Gen4, CAPI 2.0, OpenCAPI and NVLINK. These interconnects provide 2 to 5.6 times as much bandwidth for today's data-intensive workloads compared with the PCIe3 Bus Gen3 found in x86 servers.

Table 1: Summary of proposed cluster hardware

	Description	Aggregate over a 3-node system
Chassis	IBM Power AC922 2U chassis	3x nodes
Computing	2x 22-core POWER9 CPU 3.1GHz	132-cores @ 3.1GHz
GPGPUs	6x NVIDIA TESLA V100 32GB SXM3	18x V100 GPUs: 140.4 DP TFlops and 576GB HBM2 DRAM
Memory	24x16GB DDR4 RAM	1152GB DDR4 RAM
Storage	4x375GB NVMe storage	4.5TB SSD with => 24GB/sec I/O using GPFS parallel file system
Network	Mellanox ConnectX6 Dual Port 200GbE	800Gb/sec bisection bandwidth (switchless design)
Cooling	Water cooling on CPUs and GPUs	
Power	2000 watts per node	Total power for the 3-node 6U cluster should be 6kW at full blast

In the prior competitions (2014, 2015, and 2017), we had an IB switch that consumed about 200 watts of power under load, and it caused significant power imbalance. In order to reduce this 200 watts power loss, we decided to create a direct mesh interconnect with 2 Ethernet ports at 200Gb/s each per node, and direct cables from every node to every other node. This not only reduces power consumption, but will also deliver 4X the bandwidth compared to the more traditional switch based 100GbE topology, and it should also reduce network latency for small packets.

B. Software

The software stack includes the use of Linux Redhat, Warewulf (cluster management), Slurm (job management), IBM GPFS (parallel file system), MVAPICH2 MPI, and Allinea power monitoring. We plan on setting up a small power efficient Ethernet network with a 12-port switch and a MinnowBoard MAX acting as a provisioning agent, host for the scheduler, and monitoring host.

C. Power Management Optimizations

To manage power per node, we will explore dynamic frequency scaling, processor state management, and elimination of cooling fans, and water cooling. We also plan on using advanced compiler research to find the best compiler optimizations that delivers the best flops/watt. One such research project we will investigate is Softener, a compiler for C/C++ nondeterministic applications.

D. Applications

Linpack Benchmark: We will work with IBM and NVIDIA to obtain optimized HPL binaries, as well as guidance for users looking to compile the software for themselves. After carefully choosing problems sizes and parameters, and controlling for power consumption, we could still outperform an all Intel solution or a PCIE GPU solution on a per-watt basis. With a tuned binary, we suspect that we could further out perform CPU-based competitors. We plan to leverage NVIDIA V100 GPUs for all the applications, even the ones that do not support GPUs; we plan on porting as much of the applications to CUDA as time permits.

Parallel Deep Learning with Horovod: We believe that our GPU heavy platform will give us a leg-up on the Horovod application. Given the included 18 NVIDIA V100 NVLink GPUs, we expect that our deep learning tuned cluster will work exceptionally well. The NVLink GPU interconnect and the dual port 200GbE network will be 4X~6X faster than more traditional approaches utilizing x86 servers and switches. This significant bump in interconnect performance will give our cluster a big edge over the competition.

OpenMC: The OpenMC application is a monte carlo particle transport code that supports HDF5. Although this application has been explored on traditional processor architectures, there seems to be GPU supports based on recently published work from 2016 [18]. We intend to aggressively pursue the CUDA codebase to leverage the huge computing capabilities of the V100 GPUs.

SeisSol: This application is part of the reproducibility challenge in which high-resolution simulations of the 2004 Sumatra-Andaman earthquake are run and evaluated. It is note-worthy that the Chicago Fusion team in 2017 received the highest score on the reproducibility component. We plan to study this SC17 paper in detail and complete the reproducibility challenge with top notch scores again.

Mystery Application: We believe that our hybrid platform will give us a leg-up on the mystery application. SCC organizers have historically chosen mystery applications with reasonable CPU and GPU support. With ample compute power in both CPUs and GPUs, we expect to be able to support a wide range of applications well.

III. Vendor/Institution Support

ANL is in the early stages of confirming sponsorship from IBM. NVIDIA has sponsored the Chicago Fusion team in 2014 and 2017, and we expect them to sponsor the team again in 2018. Argonne National Laboratory and Illinois Institute of Technology have already confirmed sponsorship for the Chicago Fusion team. We are aiming for the competition hardware to be available to the students in June 2018. Our ANL sponsor points of contacts are William Scullin and Ben Allen. The team primary advisor is Ioan Raicu from IIT.

IV. Diversity

The team is diverse in the sense that it includes high-school students, and undergrads across different years, from 2nd year to 4th years, as well as 1 female student. Most of the students are straight A students, and some of the undergrads have a full-ride scholarships.

Nearly all the students have Linux experience through research in the DataSys laboratory, internships, or coursework. The undergrads have been exposed to a variety of programming models such as multi-threading, OpenMP, MPI, CUDA, OpenCL, MapReduce, workflows, client/server architectures, sockets, and event-driven concurrent programming. The high-school students will be given an intensive 1 week training on all these topics at the beginning of the summer. All students have been working for many years with C/C++ as well as Java. Most have been exposed to batch schedulers (e.g. Slurm and SGE) and are proficient in bash scripting, low level OS kernel tuning for process management and network tuning, and using profiling tools to analyze performance bottlenecks and issues. They have also been exposed to a variety of clouds from Google, Microsoft, and Amazon, and are familiar with everything from user-level virtualization, to para-virtualization, to hardware-based network virtualization. They have also used both Ethernet and Infiniband networks and are familiar with advanced features that could affect network performance (e.g. frame size in Ethernet, Single Root Input/Output Virtualization SRIOV for Infiniband). One of the students (Alex Ballmer) has also participated in SCC in 2014, 2015, and 2017, and has attended the SC conference every year since 2014; this would be his last attempt in 2018.

V. Team Preparation

The mentors have organized a 10-week summer program that will run at IIT from May 29th 2018 to August 3rd 2018 [3], where the 6 identified students will spend 40 hours a week preparing for the SCC competition (paid hourly with funds from a NSF REU supplement). The students will be mentored by the 5 mentors as well as graduate students in the DataSys lab at IIT. The students will spend the summer studying the announced applications, in doing code review of other scientific

applications, and in setting up and configuring the proposed cluster, including installing and configuring Linux, a shared file system, MPICH, HPL, and all the announced applications. Students will also use ANL resources such as Chameleon, Cooley, Theta, and the 2014/2015/2017 SCC clusters as additional testbeds to help us better understand the application performance and scalability. A new testbed that is coming online this summer is the NSF-funded Mystic testbed, which will have a variety of hardware (e.g. Intel/AMD x86, ARM, IBM Power9, NVIDIA V100, Intel FPGA) for the students to explore. They will explore a large parameter space that governs the possible configurations that will yield the best performance per watt.

The team mentors firmly believe in live visualizations as one of the best mechanisms to understand the performance and bottlenecks of an application. We will use a combination of monitoring tools, such as the Darshan project [17] being developed at Argonne National Laboratory. We will also leverage the innovative work in distributed filesystems FusionFS [13] to outperform more traditional HPC filesystems such as PVFS or GPFS. Dr. Raicu's group has made much progress in the design and implementation of distributed key/value storage systems (ZHT [14]) which might come in valuable to further accelerate the respective data-intensive HPC applications. Much of the research [9, 10, 11, 12, 15, 16] happening in the DataSys laboratory could be put to the test in accelerating these HPC applications.

In both 2014 and 2015, the IIT team built their own custom system status display from LED strips – a unique feature of our rack that earned a lot of attention from the SC14 and SC15 attendees. IIT has a first-rate architecture school and we intend to build on what we learned from the prior year's status display, suggestions from architecture faculty and students, and the industrial design of historical systems to produce a system that will be visually engaging and functional.

The students will participate in daily meetings with graduate students, and weekly meetings with their mentors. They will join in various activities from a parallel running NSF REU program at IIT called "BigDataX: From theory to practice in Big Data computing at eXtreme scales" [2]; some of these activities are "Life as a PhD Student", "Careers in Research", "Women in Computing", picnic with other REU programs in Chicago, and a day-long field trip to ANL where students can meet researchers in a variety of fields using computing to solve grand challenges.

The final students who will make up the Chicago Fusion Team will continue training in the Fall semester. They will take 3-credits of independent study to allow them to have sufficient time in their schedule to continue preparing for SCC with at least 10 hours/week. Over the course of 6 months (June to November), each student would have invested nearly 500 hours in this competition!

VI. WHY THE CHICAGO FUSION TEAM WILL WIN

In 2017, we ran through a series of questions on why we believed we would win. All of our answers from last year carry forward, with some small edits.

Why do you think your team will have an edge in the event (what's your secret sauce)?

A steady supply of candy, but there's a rumor that there will be cake at the end if we can just solve all the puzzles...

What will you do to prepare that makes your team unique?

Weekly meetings all summer and all fall will come with acknowledgement of great successes and great failures in the form of candy. Some teams might be afraid of a trojan horse, we know it's only really trouble if there's a piñata.

What are you most looking forward to in the competition?

Partying all night long, two nights in a row, with a half a million-dollar toy doesn't sound too shabby.

How long will it take you to get to Dallas?

About eight months from start to finish - if we practice, practice, practice.

What do your team members do for fun?

What is this fun? Compiler optimization flags, here we come! Now if we could only remember which ones might break our code…

If your team had a theme song, what would it be?

As we're from Chicago and on a mission from God, we nominate "Everybody Needs Somebody To Love" by The Blues Brothers.

ACKNOLEDGEMENT

This work is supported in part by NSF OCI-1054974 CAREER award (REU supplement) and NSF CNS-1730689 CRI award (Mystic Testbed). This work will also use resources of the ALCF at ANL, which is supported by the DOE contract DE-AC02-06CH11357. Furthermore, the Chicago Fusion team will make use of the NSF-funded Chameleon testbed at Argonne National Laboratory. The team is grateful for the generous support by IBM, NVIDIA, Argonne National laboratory, Illinois Institute of Technology, and the National Science Foundation.

REFERENCES

[1] Alexander Ballmer, et al. "Scalable High-performance Computing in a 3000-Watt Power Envelope", Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2017

[2] NSF REU Site: BigDataX: From theory to practice in Big Data computing at eXtreme scales; http://datasys.cs.iit.edu/grants/BigDataX/index.html, 2018

[3] BigDataX SCC Program -- Summer 2017; http://datasys.cs.iit.edu/grants/BigDataX/2017/scc.html, 2018

[4] Student Cluster Competition (SCC) @ IEEE/ACM Supercomputing/SC 2014; http://datasys.cs.iit.edu/events/SCC-SC14/index.html, 2018

[5] Student Cluster Competition (SCC) @ IEEE/ACM Supercomputing/SC 2015; http://datasys.cs.iit.edu/events/SCC-SC15/index.html, 2018

[6] Ben Walters, et al. "15 TFlops Haswell vs. 60 TFlops Knight Landing for HPC Scientific Computing Applications", Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2015

[7] Kevin Brandstatter, et al. "Delivering 3.5 Double Precision GFlops/Watt and 200Gb/sec Bi-Section Bandwidth with Intel Xeon Phi-based Cisco Servers", Student Cluster Competition (SCC), IEEE/ACM Supercomputing/SC 2014

[8] Kevin Brandstatter, et al. "Experiences in Optimizing Cluster Performance For Scientific Applications: Controlling Configuration, Utilization, and Power Consumption", GCASR 2015

[9] Dongfang Zhao, Da Zhang, Ke Wang, Ioan Raicu. "Exploring Reliability of Exascale Systems through Simulations", ACM HPC 2013

[10] Ioan Raicu, et al. "Towards Data Intensive Many-Task Computing", book chapter in "Data Intensive Distributed Computing: Challenges and Solutions for Large-Scale Information Management", IGI Global Publishers, 2012

[11] Dongfang Zhao, Ioan Raicu. "HyCache: A User-Level Caching Middleware for Distributed File Systems", IEEE HPDIC 2013

[12] Michael Wilde, et al. "Extreme-scale scripting: Opportunities for large task-parallel applications on petascale computers", Scientific Discovery through Advanced Computing Conference (SciDAC09) 2009

[13] Dongfang Zhao, et al. "Towards Exploring Data-Intensive Scientific Applications at Extreme Scales through Systems and Simulations", IEEE Transaction on Parallel and Distributed Systems (TPDS) Journal 2015

[14] Tonglin Li, et al. "A Convergence of Key-Value Storage Systems from Clouds to Supercomputers", Concurrency and Computation: Practice and Experience (CCPE) Journal 2015

[15] Ke Wang, et al. "Exploring the Design Tradeoffs for Extreme-Scale High-Performance Computing System Software", IEEE Transaction on Parallel and Distributed Systems (TPDS) 2015

[16] Dongfang Zhao, Ioan Raicu. "Distributed File Systems for Exascale Computing", Doctoral Showcase, IEEE/ACM Supercomputing/SC 2012

[17] Darhsan, http://www.mcs.anl.gov/research/projects/darshan/, 2017

[18] Timothy P. Burke, et al. "GPU Acceleration of Kernel Density Estimators in Monte Carlo Neutron Transport Simulations," Trans. Am. Nucl. Soc., 115, 531-534 (2016)