Benchmarking in the Data Center: Expanding to the Cloud (BID’21)

Workshop held in conjunction with PPoPP 2021:

Principles and Practice of Parallel Programming 2021

https://ppopp21.sigplan.org/track/PPoPP-2021-workshops-and-tutorials#program

https://parallel.computer/index.html

Program

Feb 27, 17:00 - 21:00, UTC-5, Eastern Time (US & Canada)

17:00: Welcome Remark, 2”, Juan (Jenny) Chen	
Keynote Session (17:10 - 19:45)	
17:10-17:35	Mr. Sven Breuner, VAST Data & Dr. Chin Fang, Zettar *Title: elbencho – A new storage benchmark for AI et al* *Moderator: Benson Muite*
17:35-17:45	Q&A *Moderator: Benson Muite*
17:45-18:10	Dr. Wanling Gao, Institute of Computing Technology, Chinese Academy of Sciences *Title: AIBench Scenario: Scenario-distilling AI Benchmarking* *Moderator: Juan (Jenny) Chen*
18:10-18:20	Q&A *Moderator: Juan (Jenny) Chen*
18:20-18:45	Dr. Jason (Zhixiang) Ren, Peng Cheng Laboratory, China *Title: AIPerf: AutoML as an AI-HPC benchmark* *Moderator: Kevin Brown*
18:45-18:55	Q&A *Moderator: Kevin Brown*
18:55 - 19:10, Coffee Break	
19:10-19:30	Dr. Dan Huang, Sun Yat-sen University, China *Title: A Comprehensive Study of In-Memory Computing on Large HPC Systems* *Moderator: Samar Aseeri*
19:35-19:45	Q&A *Moderator: Samar Aseeri*

Panel Session	
19:45-20:30	Industry Panel Discussion *Panelists*: Dr. Sreeram Potluri, NVIDIA, USA Dr. Jithin Jose, Microsoft, USA Ms. Verónica G. Melesse Vergara, Oakridge Leadership Computing Facility, USA *Moderators: Ammar Awan, Shahzeb Siddiqui*
20:30-20:45	General Discussion
Close Remark	

Keynote Speakers

Mr. Sven Breuner

Bio:
Sven Breuner is the creator of the award-winning BeeGFS parallel file system. As the former CEO and CTO of ThinkParQ, he helped customers with the adoption of BeeGFS in data centers around the globe and including some of the fastest supercomputers of our time. In 2019, he decided to focus on enabling adoption of low latency flash storage technology and joined Excelero as Field CTO for the past 2 years before soon moving on to VAST Data. Sven holds a M.Sc. degree in Computer Science.

Dr. Chin Fang

Bio:
Chin Fang holds a master's and a doctoral engineering degree from Stanford University. He is the founder & CEO of Zettar Inc. (Zettar). The company's mission is to create a simple, scale-out capable, and efficient software data mover for moving data at scale and speed. Targeted users include modern distributed data-intensive enterprises and large science projects such as the Linac Coherent Light Source II (LCLS-II). In March 2019, Zettar won the Supercomputing Asia 2019 Data Mover Challenge, a grueling two-month long international competition at the highest level. Dr. Fang led the Zettar team to beat out six other elite national teams from the U.S. and Japan by a wide margin. Since 2015, Zettar has been engaged to support the ambitious data movement requirements of LCLS-II, a premier U.S. DOE Exascale Computing preparation project. Zettar is supported by its revenue and has been awarded highly competitive grants by two U.S. federal agencies, the National Science Foundation and the Department of Energy.

Title:
elbencho - A new storage benchmark for AI et al

Abstract:
A misunderstanding of storage system characteristics is often the reason behind an inability to nicely scale critical jobs to a required level of performance - be it in DeepLearning or other areas, on-prem or in the cloud. To make it easy to characterize or validate a modern storage system for a certain workload, Sven Breuner, the creator of the BeeGFS parallel file system, has published elbencho. Elbencho is a new vendor-neutral storage benchmark that supports a wide variety of access patterns, ranging from lots of small files over random IOPS in large shared files up to low level testing of underlying block devices. For DeepLearning, elbencho can even include GPUs in the data access. Elbencho is available on github for everyone to use and contribute. A major contribution, the storage sweep tools, came from Zettar Inc to test perfomance for a wide variety of file sizes and generate a result chart with a simple push-button method. The contribution was motivated by the need of a storage benchmarking tool for an important U.S. DOE Energy Sciences Network (ESnet) ESnet6 initiative. It considers only hyperscale data sets, i.e. each set has >= 1M files, or >= 1TB in size, or both. The file size range considered is wide, ranging from 1KiB - 1TiB. Elbencho was the only tool that could meet such strigent requirements. To automate the benchmarking, the storage sweep tools were born.

Dr. Wanling Gao

Bio:
Wanling Gao is an Assistant Professor in computer science at the Institute of Computing Technology, Chinese Academy of Sciences and University of Chinese Academy of Sciences. Her research interests focus on big data and AI benchmarking, computer architecture, workload characterization and simulation. She received her B.S. degree in 2012 from Huazhong University of Science and Technology and her PhD degree in 2019 from Institute of Computing Technology, Chinese Academy of Sciences and University of Chinese Academy of Sciences in China.

Title:
AIBench Scenario: Scenario-distilling AI Benchmarking

Abstract:
Modern real-world application scenarios like Internet services not only consist of diversity of AI and non-AI modules with very long and complex execution paths, but also have huge code size, which raises serious benchmarking or evaluating challenges. Using AI components or micro benchmarks alone can lead to error-prone conclusions. Together with seventeen industry partners, we extract nine typical application scenarios, identify the primary components, and propose a comprehensive and representative AI benchmark suite—AIBench, including scenario, component (training and inference), and micro benchmarks, to fulfil different benchmarking requirements. As the proxy to real-world applications, the AIBench scenario benchmarks let the software and hardware designers obtain the overall system performance and find out the key components within the critical path. AIBench Training and Inference contains seventeen representative AI tasks with the state-of-the-art models to guarantee the diversity and representativeness for workload characterization. Furthermore, considering benchmarking affordability, we identify AIBench Training Subset for performance ranking. AIBench benchmarks are available from https://www.benchcouncil.org/aibenchmark.html.

Dr. Zhixiang Ren

Bio:
My main research interest is high performance artificial intelligence and scientific computing. I obtained my doctoral degree in high-energy astrophysics from the University of New Mexico in the United States, and continued my research as a post-doctoral fellow. I have published more than 30 papers with experience in the applications of high-performance scientific computing, machine learning and deep learning in science.

Title:
AIPerf: AutoML as an AI-HPC benchmark

Abstract:
The plethora of complex artificial intelligence (AI) algorithms and available high performance computing (HPC) power stimulates the expeditious development of AI components with heterogeneous designs. Consequently, the need for cross-stack performance benchmarking of AI-HPC systems emerges rapidly. The de facto HPC benchmark LINPACK can not reflect AI computing power and I/O performance without representative workload. The current popular AI benchmarks like MLPerf have fixed problem size therefore limited scalability. To address these issues, we take a "pencil-and-paper" manner and propose an end-to-end benchmark suite utilizing automated machine learning (AutoML), which not only represents real AI scenarios, but also is auto-adaptively scalable to various scales of machines. We implement the algorithms in a highly parallel and flexible way to ensure the efficiency and optimization potential on diverse systems with customizable configurations. To learn from the success of LINPACK, we utilize operations per second (Ops), that is measured in an analytical and systematic approach, as the major metric to quantify the AI performance. We perform evaluations on various systems to ensure the benchmark's stability and scalability, from 4 nodes with 32 NVIDIA Tesla T4 (56.1 Tera-Ops measured), up to 50 nodes with 400 Huawei Ascend 910 (16.99 Peta-Ops measured) and 50 nodes with 400 NVIDIA V100 (9.1 Peta-Ops measured). With flexible workload and single metric measurement, our benchmark can scale and rank AI-HPC easily. The source code, specifications and detailed procedures are publicly accessible on GitHub\footnote{AIPerf: https://github.com/AI-HPC-Research-Team/AIPerf}.

Dr. Dan Huang

Bio:
Dr. Dan Huang currently is an associate professor in School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou China. He received Ph.D. in computer engineering at University of Central Florida. Before this, he received master and bachelor degrees in Southeast University and Jilin University respectively. His research interests are scientific data management, in-memory computing, parallel programming model, distributed storage systems, virtualization technology and the I/O of distributed system. In addition, he worked in Oak Ridge National Lab, USA(ORNL) as a researcher for about ten months. The responsibilities of this research includes collaborating with ORNL's research

scientists on software-defined storage technologies and researching, developing and evaluating new storage solutions for Department of Energy applications. His researches have been published in some top-tier conferences and journals, including IEEE TC, IEEE TPDS, IEEE ICDCS, IEEE IPDPS and IEEE DAC.

Title:
A Comprehensive Study of In-Memory Computing on Large HPC Systems

Abstract:
With the increasing fidelity and resolution enabled by high-performance computing systems, simulation-based scientific discovery is able to model and understand microscopic physical phenomena at a level that was not possible in the past. A grand challenge that the HPC community is faced with is how to handle the large amounts of analysis data generated from simulations. In-memory computing, among others, is recognized to be a viable path forward and has experienced tremendous success in the past decade. Nevertheless, there has been a lack of a complete study and understanding of in-memory computing as a whole on HPC systems. This paper presents a comprehensive study, which goes well beyond the typical performance metrics. In particular, we assess the in-memory computing with regard to its usability, portability, robustness and internal design trade-offs, which are the key factors that of interest to domain scientists. We use two realistic scientific workflows, LAMMPS and Laplace, to conduct comprehensive studies on state-of-the-art in-memory computing libraries, including DataSpaces, DIMES, Flexpath and Decaf. We conduct cross-platform experiments at scale on two leading supercomputers, Titan at ORNL and Cori at NERSC, and summarize our key findings in this critical area.