

E-NEWS LETTER

Master of Computer Applications

EDITORIAL BOARD

Editor in Chief

Dr. Anuj Kumar
(Head of Department)

Editor

Mr. Vijay Kumar
(Assistant Professor)

Co-Editors

Mr. Arvind Kumar Mishra
(Assistant Professor)

Student Coordinators

Mr. Nitin Kumar
Ms. Priyanka Gangwar

CONTENTS:

Latest Update World	01-04
Faculty Arena	04-05
Students Corner	05-05
College Update	05-06

LATEST UPDATES WORLD

Supercomputing center dataset aims to accelerate AI research into optimizing high-performance computing systems

When the MIT Lincoln Laboratory Supercomputing Center (LLSC) unveiled its TX-GAIA supercomputer in 2019, it provided the MIT community a powerful new resource for applying artificial intelligence to their research. Anyone at MIT can submit a job to the system, which churns through trillions of operations per second to train models for diverse applications, such as spotting tumors in medical images, discovering new drugs, or modeling climate effects.



But with this great power comes the great responsibility of managing and operating it in a sustainable manner—and the team is looking for ways to improve.

"We have these powerful computational tools that let researchers build intricate models to solve problems, but they can essentially be used as black boxes. What gets lost in there is whether we are actually using the hardware as effectively as we can," says Siddharth Samsi, a research scientist in the LLSC.

To gain insight into this challenge, the LLSC has been collecting detailed data on TX-GAIA usage over the past year. More than a million user jobs later, the team has released the dataset open source to the computing community.

Their goal is to empower computer scientists and data center operators to better understand avenues for data center optimization—an important task as processing needs continue to grow. They also see potential for leveraging AI in the data center itself, by using the data to develop models for predicting failure points, optimizing job scheduling, and improving energy efficiency. While cloud providers are actively working on optimizing their data centers, they do not often make their data or models available for the broader high-performance computing (HPC) community to leverage. The release of this dataset and associated code seeks to fill this space.

"Data centers are changing. We have an explosion of hardware platforms, the types of workloads are evolving, and the types of people who are using data centers is changing," says Vijay Gadepally, a senior researcher at the LLSC. "Until now, there hasn't been a great way to analyze the impact to data centers. We see this research and dataset as a big step toward coming up with a principled approach to understanding how these variables interact with each other and then applying AI for insights and improvements."

Papers describing the dataset and potential applications have been accepted to a number of venues, including the IEEE International Symposium on High-Performance Computer Architecture, the IEEE International Parallel

and Distributed Processing Symposium, the Annual Conference of the North American Chapter of the Association for Computational Linguistics, the IEEE High-Performance and Embedded Computing Conference, and International Conference for High Performance Computing, Networking, Storage and Analysis.

Workload classification

Among the world's TOP500 supercomputers, TX-GAIA combines traditional computing hardware (central processing units, or CPUs) with nearly 900 graphics processing unit (GPU) accelerators. These NVIDIA GPUs are specialized for deep learning, the class of AI that has given rise to speech recognition and computer vision.

The dataset covers CPU, GPU, and memory usage by job; scheduling logs; and physical monitoring data. Compared to similar datasets, such as those from Google and Microsoft, the LLSC dataset offers "labeled data, a variety of known AI workloads, and more detailed time series data compared with prior datasets. To our knowledge, it's one of the most comprehensive and fine-grained datasets available," Gadepally says.

Notably, the team collected time-series data at an unprecedented level of detail: 100-millisecond intervals on every GPU and 10-second intervals on every CPU, as the machines processed more than 3,000 known deep-learning jobs. One of the first goals is to use this labeled dataset to characterize the workloads that different types of deep-learning jobs place on the system. This process would extract features that reveal differences in how the hardware processes natural language models versus image classification or materials design models, for example.

The team has now launched the MIT Datacenter Challenge to mobilize this research. The challenge invites researchers to use AI techniques to identify with 95 percent accuracy

the type of job that was run, using their labeled time-series data as ground truth.

Such insights could enable data centers to better match a user's job request with the hardware best suited for it, potentially conserving energy and improving system performance. Classifying workloads could also allow operators to quickly notice discrepancies resulting from hardware failures, inefficient data access patterns, or unauthorized usage.

Too many choices

Today, the LLSC offers tools that let users submit their job and select the processors they want to use, "but it's a lot of guesswork on the part of users," Samsi says. "Somebody might want to use the latest GPU, but maybe their computation doesn't actually need it and they could get just as impressive results on CPUs, or lower-powered machines."

Professor Devesh Tiwari at Northeastern University is working with the LLSC team to develop techniques that can help users match their workloads to appropriate hardware. Tiwari explains that the emergence of different types of AI accelerators, GPUs, and CPUs has left users suffering from too many choices. Without the right tools to take advantage of this heterogeneity, they are missing out on the benefits: better performance, lower costs, and greater productivity.

"We are fixing this very capability gap—making users more productive and helping users do science better and faster without worrying about managing heterogeneous hardware," says Tiwari. "My Ph.D. student, Baolin Li, is building new capabilities and tools to help HPC users leverage heterogeneity near-optimally without user intervention, using techniques grounded in Bayesian optimization and other learning-based optimization methods. But, this is just the beginning. We are looking into ways to introduce heterogeneity in our data centers in a principled approach to help our users achieve the maximum

advantage of heterogeneity autonomously and cost-effectively."

Workload classification is the first of many problems to be posed through the Datacenter Challenge. Others include developing AI techniques to predict job failures, conserve energy, or create job scheduling approaches that improve data center cooling efficiencies.

Energy conservation

To mobilize research into greener computing, the team is also planning to release an environmental dataset of TX-GAIA operations, containing rack temperature, power consumption, and other relevant data.

According to the researchers, huge opportunities exist to improve the power efficiency of HPC systems being used for AI processing. As one example, recent work in the LLSC determined that simple hardware tuning, such as limiting the amount of power an individual GPU can draw, could reduce the energy cost of training an AI model by 20 percent, with only modest increases in computing time. "This reduction translates to approximately an entire week's worth of household energy for a mere three-hour time increase," Gadepally says.

They have also been developing techniques to predict model accuracy, so that users can quickly terminate experiments that are unlikely to yield meaningful results, saving energy. The Datacenter Challenge will share relevant data to enable researchers to explore other opportunities to conserve energy.

The team expects that lessons learned from this research can be applied to the thousands of data centers operated by the U.S. Department of Defense.

Other collaborators include researchers at MIT Computer Science and Artificial Intelligence Laboratory (CSAIL). Professor Charles Leiserson's Supertech Research Group is

investigating performance-enhancing techniques for parallel computing, and research scientist Neil Thompson is designing studies on ways to nudge data center users toward climate-friendly behavior.

Samsi presented this work at the inaugural AI for Datacenter Optimization (ADOPT'22) workshop last spring as part of the IEEE International Parallel and Distributed Processing Symposium. The workshop officially introduced their Datacenter Challenge to the HPC community.

FACULTY ARENA

“Azadi Ka Amrit Mahotsav”



This year India is celebrating 76th Independence Day on 15th August 2022 to mark the country's freedom from nearly two centuries of British colonial rule. Prime Minister Narendra Modi is leading the celebrations from Red Fort in New Delhi and is addressing the nation from Red Fort in New Delhi. This is his ninth address as the prime minister. To celebrate the 75 years of Independence the Government of India has started an initiative called “Azadi Ka Amrit Mahotsav”.

On March 12, 2021, Prime Minister Narendra Modi has launched the “Azadi Ka Amrit Mahotsav” to celebrate and commemorate 75 years of independence that started a 75-



week countdown to our anniversary of independence.

This means that India will celebrate its 76th Independence Day on August 15, 2022, and mark the end of its 75 years of freedom.



Indian Independence Day 2022: History

The Englishmen entered Surat and Gujarat in 1619 for trading purposes. After the victory of the English East India Company at the Battle of Plassey in 1757, the Britishers made their control over India.

The British Government dominated over the people of India for almost 200 years starting from 1757. Legendary freedom fighters and leaders like Bhagat Singh, Chandrashekhar Azad, Netaji Subhas Chandra Bose, Sardar Vallabhbhai Patel, Mahatma Gandhi and others sacrificed everything to see India as a free nation.

The Indian Independence movement began during World War I and was led by Mohandas Karamchand Gandhi. On August 15, 1947, India got its freedom, ending an almost 200-year British rule. India's first prime minister Jawaharlal Nehru raised the Indian national flag above the Lahori Gate of Red Fort in Delhi on August 15, 1947.

It is a tradition that has since been followed by the incumbent Prime Minister, followed by an address to the country.

Vijay Kumar Dubey
Assistant professor, MCA

STUDENT'S CORNER

मेरा भारत

हरी भरी धरती हो
नीला आसमान रहे
फहराता तिरंगा,
चाँद तारों के समान रहे।

त्याग शूर वीरता
महानता का मंत्र है
मेरा यह देश
एक अभिनव गणतंत्र है

शांति अमन चैन रहे,
खुशहाली छाये
बच्चों को बूढ़ों को
सबको हर्षाये

हम सबके चेहरो पर
फैली मुस्कान रहे
फहराता तिरंगा चाँद
तारों के समान रहे।

Priyanka Gangwar,
MCA, 2021 Batch

COLLEGE UPDATE

➤ **15th August-2022 (स्वतंत्रता दिवस)**

आज़ादी का अमृत महोत्सव की शुभकामनाएँ

15th अगस्त, 2022 को एसआरएमएस कॉलेज ऑफ इंजीनियरिंग एंड टेक्नोलॉजी में स्वतंत्रता दिवस समारोह में मन की स्वतंत्रता, शब्दों में विश्वास और आत्माओं में गर्व देखा गया।



प्रातः 8:55 बजे तक सभा का आयोजन हुआ और उसके बाद सम्मानित अतिथियों का स्वागत किया गया। अध्यक्ष सर श्री देव मूर्ति जी, ट्रस्ट सचिव श्री आदित्य मूर्ति जी, डीन शिक्षाविद डॉ. प्रभाकर गुप्ता और विभिन्न विभागों के प्रमुखों का स्वागत किया गया।

हमारे सम्माननीय अध्यक्ष श्री देव मूर्ति जी के सम्मानित हाथों से गौरवशाली राष्ट्रगान के साथ हमारा विशाल तिरंगा झंडा फहराया गया। तत्पश्चात, गायक मंडली ने अपनी

सुरीली आवाजों में देशभक्ति का गीत गाया, जिसके बाद आकाश सक्सेना ने पूरे दिल से भाषण दिया।

इसके बाद, वातावरण में एक और देशभक्ति गीत ने भारतीय दिलों को भक्ति से भर दिया। तत्पश्चात, डीन शिक्षाविद, डॉ. प्रभाकर गुप्ता ने सभी को स्वतंत्रता दिवस की शुभकामनाएं देते हुए अपना भाषण दिया।

इसके बाद, अध्यक्ष महोदय ने हार्दिक शुभकामनाएं देकर अपने भाषण की शुरुआत की और हमारे स्वतंत्रता सेनानियों और उनके संघर्षों की याद दिलाने वाले प्रेरक और शिक्षाप्रद शब्दों को जारी रखा।

अंत में टायरो की सचिव वंशिका गुप्ता द्वारा धन्यवाद प्रस्ताव दिया गया और मिष्ठान वितरण के साथ समारोह संपन्न हुआ।



➤ 19th August-2022 (हैप्पी जन्माष्टमी)

श्री कृष्ण जन्मोत्सव 19 अगस्त "2022 को बड़ी भक्ति और आनन्द के साथ मनाया गया। इस स्मरणोत्सव का आयोजन श्री राम मूर्ति स्मारक कॉलेज ऑफ इंजीनियरिंग एंड टेक्नोलॉजी द्वारा कॉलेज के राधा मोहन मंदिर में किया गया इस पावन अवसर पर सभी डॉक्टर्स, फैकल्टी, स्टाफ़ और स्टूडेंट्स को सादर आमंत्रित किया गया।

भक्तों द्वारा विलक्षण भक्ति के साथ भगवान कृष्ण की मूर्ति की पूजा की गई। पूरे मंदिर परिसर को फूलों, आम के पत्तों के साथ-साथ चमकदार रोशनी और दीयों से सजाया गया था, जो पर्यावरण को आध्यात्मिक सकारात्मकता और निष्ठा से भर देते थे।

यह पवित्र अवसर लोगों को एक साथ लाता है, जो एकता और विश्वास का प्रतीक है। एसआरएमएस रिद्धिमा ने विभिन्न नृत्य कथाओं और प्रदर्शनों द्वारा उत्साह और खुशी की भावना के साथ पर्यावरण को प्रफुल्लित किया।

माननीय सभापति श्री देव मूर्ति जी द्वारा महाभिषेक (देवता के ऊपर पानी और दूध का मिश्रण डालना) किया गया। पंजिरी और पंचामृत भगवान कृष्ण के आशीर्वाद से भरे दो दूध आधारित व्यंजन हैं जो सभी भक्तों को दिया जाने वाला आम प्रसाद था।

