IBM-Illinois Discovery Accelerator Institute

Annual Report

AUG 2021 - AUG 2022

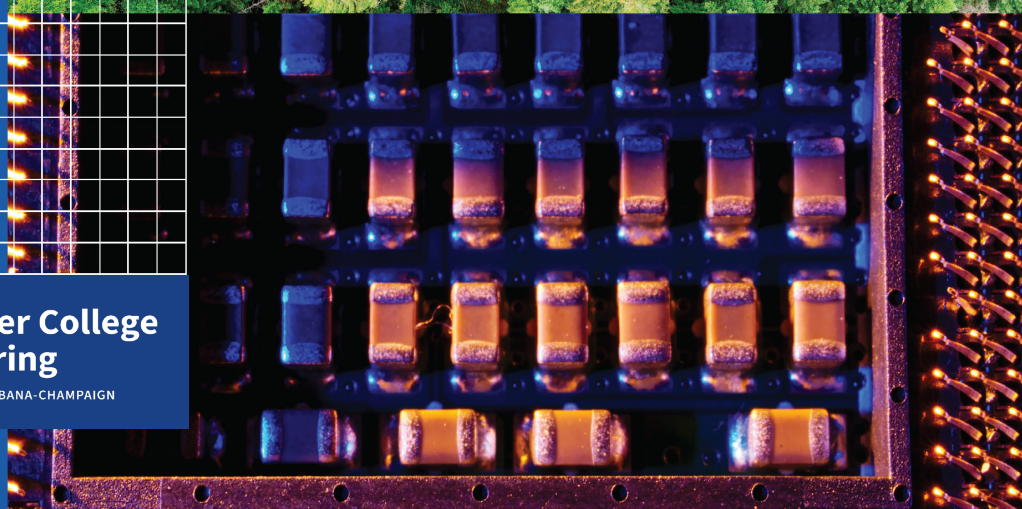

The Grainger College of Engineering

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

Contents

Contents

INTRODUCTION

The IBM-Illinois Discovery Accelerator Institute (IIDAI) is a new model of academic-industry partnership designed to increase access to technology education and skill development to spur breakthroughs in emerging areas of technology, referred to in this document as research thrusts, that include hybrid cloud and AI, quantum computing, accelerated materials discovery, and sustainability to accelerate the discovery of solutions to complex global challenges. Since its inception in the fall of 2021, the focus of the institute has been to cast a wide net around relevant projects within these thrusts to jump start the collaboration. The following is a summary of activities that took place from August 2021 through August 2022. We conclude with a short preview of strategic goals for 2023 and beyond.

1. Hybrid Cloud

1.1. Mission and goals

The Hybrid Cloud and AI thrust is focused primarily on hybrid cloud technologies across the entire stack spanning hardware, infrastructure and systems, operating systems, platform, middleware, software tools and applications. Hybrid cloud technologies are at the core of all computing today and, as such, the thrust has strong synergies with other thrusts in the institute, both as a key enabler as well as providing context and requirements to facilitate accelerated discoveries. Artificial Intelligence (AI) is another computing discipline that is omnipresent today and a key topic in the Hybrid Cloud thrust agenda, again both as an important emerging workload as well as a key enabling technology that promises innovative approaches to improve computing. For instance, AI for IT automation in general is a key focus area within the institute.

Cloud computing today brings together decades of advances and has changed the computing landscape by making it more accessible, efficient, scalable, and available. As more and more workloads move to the cloud, new applications and workloads emerge. General trends towards multi cloud and decentralization are bringing forward new distributed computing challenges in programmability, management, security and compliance, cost, performance, power, and carbon efficiency. The Hybrid Cloud and AI thrust pursues a top-down approach by prioritizing needs of strategic emerging workloads like AI/ML and domains like edge computing in addition to focusing on needs of enterprise workloads to accelerate their move to the cloud. As a result, three priority areas have emerged:

1. Emerging cloud applications, workloads, and environments, which entail developing a deep understanding of emerging needs to drive the evolution of hybrid cloud technologies.

2. Self-driving cloud for safe and resilient automation for the hybrid cloud platform that simplifies management and operations while providing the desired performance, scale, efficiency, and availability.

3. Secure, efficient, and performant hardware, systems, and infrastructure encompassing specialized hardware like accelerators, SmartNICs, programmable software-defined infrastructure, and novel approaches to systems design and virtualization.

In setting goals for the Hybrid Cloud thrust, we considered two critical aspects:

1. Significant research outcomes. These are measurable through publications including academic papers in high quality conferences and journals, blog articles, open-source contributions to strategic projects and, where effective, creation of joint intellectual property through patents.

2. A culture of innovation enabled by open exchange of ideas and a high level of engagement and collaboration between faculty members, students, and IBM researchers. These are measurable through joint publications, number of graduate student theses that were shaped by work supported by the Institute, number of high-bandwidth collaborations enabled through mechanisms like externships, post-doctoral and sabbatical assignments at IBM, IBM Researchers giving invited lectures, teaching classes, and influencing the curriculum, and jointly pursuing external funding opportunities.

1.2. Activities

The 15 projects funded during the period covered by this report can be broadly classified into the three focus areas outlined above: emerging applications and use-cases (edge and AI/ML/HPC), self-driving cloud (automation and optimization for performance, scalability, resiliency, and sustainability), and infrastructure and hardware for better security, performance, and efficiency. 33 UIUC faculty members, 41 IBM Researchers and 68 UIUC students (including Ph.D., M.S., part-time, and undergraduate students) were involved in these 15 hybrid cloud projects in the first year of the Institute. Thirteen Hybrid Cloud students visited IBM during summer 2022 as externs. They worked closely with their IBM mentors, accessed IBM computing resources and data on site when appropriate, and made significant research progress through such externships.

The Hybrid Cloud and AI thrust held two thrust-level workshops in the past year: one at UIUC and one at IBM T.J. Watson Research Center in Yorktown Heights, NY. The workshop at UIUC took place in April 2022. 13 IBM researchers (including IBM VPs, Fellows, and Directors) came to the UIUC campus to join this two-day event. Besides high-level engagement between IBM and UIUC (e.g., the Vice Chancellor for Research and Innovation, and the Dean and Executive Associate Dean of the Grainger College of Engineering all attended this workshop), all 15 UIUC project faculty leaders gave in-depth presentations to report their research findings and results. IBM visitors provided insightful and constructive feedback and comments for the projects and also actively participated in two panels on Future Research Vision and Directions for the Hybrid Cloud Thrust and Strategic Directions and Planning for the Hybrid Cloud Thrust.

The workshop at IBM took place in August 2022. More than 10 UIUC faculty PIs joined this workshop either in person or online, and each of them gave an update on the research progress of their projects and discussed near-term and long-term research and collaboration plans. This workshop included a unique activity: student externs for our thrust held a poster session and presented their research projects with an emphasis on their achievements during their externship at IBM, while working with their IBM mentors.

1.3. Highlights of accomplishments

In 2022, there were 21 papers published and an additional 10 papers either submitted or in preparation for submission (approximately 20 of these with joint authorship between UIUC and IBM), 2 invention disclosures, and 13 student externships completed.

The joint projects have been featured in selective computing forums including OSDI, ASPLOS, ECCV, KubeCon, and Linux Plumber Conference with 11 posters and presentations overall. Open-source contributions have also been made across 8 different streams. Other notable accomplishments include a joint IBM-UIUC led workshop at the IEEE BigData conference, as well as multiple patent applications.

2. Quantum Computing

2.1. Mission and goals

The IIDAI Quantum thrust will advance quantum information science, integrate quantum discoveries with related future technologies, and train a new generation of students to work in quantum-related fields. Quantum Information Science (QIS) is moving quickly and showing promise, but it is still in its early stages with many unknowns. Therefore, the IIDAI Quantum thrust takes an agile approach with bold programs that will change the way industry and academia collaborate on inventing new quantum devices and preparing students to work in QIS.

2.2. Activities

There are four research projects in the quantum thrust, focusing on:

1. Modular quantum computing architectures.
2. Superconducting devices based on high kinetic inductances.
3. Explorations of quantum measurements.
4. First principles defect simulations and quantum embedding.

There is also one education/outreach project on Quantum Education and Training. The activity for Y1 was focused on hiring, and research and program development. An Illinois Quantum Information Science and Technology Center (IQUIST) program manager, Patrick Snyder, is providing part-time support for the IIDIA quantum thrust. A new hire for an IQUIST education coordinator was approved; a search is underway. This position will provide partial FTE support for the IIDAI quantum thrust.

The IIDAI quantum thrust has been highly synergistic across the midwestern quantum ecosystem. It has collaborated with IQUIST and the NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks to support the Quantum Applications Grant Program and the IBM/Illinois Quantum Undergraduate Research Scholars Program.

2.3. Highlights of accomplishments

Quality factors exceeding 100,000 were measured for a newly designed connectorization approach for modular Transmon qubit quantum computing. This is amongst the highest quality-factors ever reported and will be sufficient for a high-fidelity cable gate. A fabrication recipe for high-quality superconducting devices with large kinetic inductances was developed, and quality factors exceeding 200,000 were observed for resonators made from this material. A new quantum variational optimization scheme was used to study a multiple-access communication scenario involving shared classical and quantum correlations. The Frenkel Hamiltonian as a model for a hydrogen fluoride array was simulated using 25 qubits on the IBM quantum simulator.

3. Materials Discovery

3.1. Mission and Goals

IBM and UIUC have a common interest in developing the tools for the future of materials discovery. These include tools for extracting materials information from the scientific literature, frameworks for efficient simulation or virtual screening of materials candidates, AI models for hypothesis generation and materials evaluation, tools for retrosynthesis planning, and automated laboratory systems for the synthesis and analysis of candidate materials. Along with these technologies, there is a parallel need to develop a technical workforce skilled in modern materials discovery.

The Materials Discovery thrust will develop frontier AI/ML, high throughput materials synthesis and characterization, and advanced laboratory automation tools. We will focus on the following specific technologies, all aligned with IBM's overall Accelerated Discovery strategy.

1. Autonomous labs—we will explore the limits of AI applied to laboratory automation and robotics. This includes accurate deep learning models for reaction prediction and retrosynthetic analysis as well as models to translate experimental methods to robotic actions, removing the need to program a synthesis robot directly. A key role of this automation will be the rapid generation of large materials property data sets that can be used to train predictive AI models.

2. AI for polymer materials discovery, including reinforcement learning systems for polymer synthesis, expert-in-the-loop systems for efficient chemical decision making, and recommendation systems for synthesis or reaction parameters.

3. Generative models for materials design that leverage both rule-based ML and deep learning to enable inverse design. An additional goal of the thrust is to seed a new discovery-ready workforce through training of graduate students and development of deep academic relationships between IBM and UIUC.

3.2. Activities

The joint UIUC and IBM team have been working on three distinct yet interconnected projects:

1. Development of a chemoenzymatic retrosynthesis tool that combines biological catalysis and chemical catalysis.
2. Construction of a polymer dataset and demonstration of its utility for machine-learning driven materials discovery.
3. Development of a machine learning driven closed-loop workflow for discovering organic photovoltaic molecules meeting the 10/10 challenge: 10% power conversion efficiency (PCE) and 10 years photostability.

3.3. Highlights of Accomplishments

The major accomplishment of the thrust was the development of an ongoing and active collaboration between IBM researchers and UIUC students. Starting in 2022, there were monthly meetings where the students provided research updates, asked questions, and interacted with the IBM team. In addition, the thrust nucleated one new seed project in the area of foundation AI models for materials.

4. Climate and Sustainability

4.1. Mission and Goals

The Sustainability thrust is focused on creating impactful work that aligns with the strategic vision of IBM and fits with the academic goals of the UIUC faculty and students in areas relevant to climate change and sustainability. IBM has been focusing on Sustainability as a key area of growth both for IBM business units and for the IBM Research Division. IBM Research has a sharp focus on accelerating discovery in areas that are critical to reducing greenhouse gas emissions, decarbonizing operations, and analyzing the impact of the changing climate on businesses and communities.

The goals of the Sustainability thrust reflect strategic objectives of the IBM Research Climate and Sustainability team: to accelerate scientific discovery across areas of materials for CO2 capture, sequestration, natural solutions for carbon sequestration, and climate impact modeling and prediction.

Additionally, we continue to encourage and support teams to pursue external funding to grow and sustain their IBM-Illinois collaborative endeavors. As we head into 2023, another two goals will be added for the thrust: (1) create a significant cross-thrust initiative focused on the intersection of Sustainability and Hybrid Cloud and (2) enable open-source tools that can be leveraged by the global ecosystem of sustainability researchers, students, and early-stage companies.

4.2. Activities

The key research activities of Y1 center on the two funded research projects that include accelerating discovery of climate-related impacts across sectors and understanding extreme events and their impact on society; and materials discovery to support carbon dioxide conversion and utilization.

4.3. Highlights of Accomplishments

Accomplishments for the Sustainability thrust include:

1. the on-boarding of two seed projects.
2. substantive engagement of IBM with the Midwest AgTech Engine proposal (NSF Regional Innovation Engine) led by UIUC.
3. collaboration between UIUC and IBM researchers in the submittal of a NOAA proposal.
4. four publications submitted by the two funded projects (Award No. 107357; Award No. 107359).

5. Diversity, Equity, and Inclusion Cross-cut

5.1. Mission and Goals

The DEI cross-cut in the IIDAI is aimed at leveraging and enhancing programming that enables broader participation in activities of the Institute. The goal is to increase participation and retention of students from underrepresented groups through recruitment and mentoring by those involved in IIDAI. While direct involvement in IIDAI activities is desired, the importance of building the overall interest of students from a variety of backgrounds in STEM is a recognized goal of this cross-cut. Efforts in this area take the “long view” and consistently engage with student populations to encourage their participation.

5.2. Activities

5.2.1. ISUR program

In collaboration with the Illinois Scholars Undergraduate Research (ISUR) program, we are supporting undergraduate researchers involved in 3 hybrid cloud, 1 hybrid cloud+AI, 1 hybrid cloud+sustainability, 1 materials discovery, 3 QIS, and 1 sustainability projects. Participating IIDAI researchers include Professors Kindratenko, McHenry, Iyer, Chang, Schleife, Chitambar, Pfaff, Leditzky, and Wuebbles.

Sixteen students were selected as IIDAI Research Scholars to work on these projects. Students were selected from 36 applicants and were matched based on mutual interest between student and researcher.

The demographic makeup of the selected scholars is as follows: 27% women, 6% first-generation college students, 20% Hispanic, 0% African-American; 47% seniors, 40% juniors, 13% sophomores. Student majors break down as follows: 53% CS, 20% EE, 13% CompE, 13% Physics. Students will continue on their research projects through the spring semester and will present their work at the Illinois Scholars Undergraduate Research (ISUR) Expo on April 25, 2023.

5.2.2. GIANT projects

Five projects were co-funded by IIDAI and the Inclusion, Diversity, Equity & Access (IDEA) Institute: pre-college family engagement (project 1), skill development and career awareness for undergraduate students (projects 2 and 3), and building community and belonging for graduate students (projects 4 and 5). Two Grassroots Initiatives to Address Needs Together (GIANT) projects are led by teams of graduate students. These students are developing project management and leadership skills by carrying out all aspects of their projects (e.g., writing the proposal, leading a team, designing the research survey and getting Institutional Review Board (IRB) approval, running the initiative, and managing the budget).

5.3. Highlights of accomplishments

The DEI cross-cut supported 5 projects through the Grassroots Initiatives to Address Needs Together (GIANT) program (https://idea.illinois.edu/giant/giant-project-selections-2022) and supported 16 undergraduate students to participate in the Illinois Scholars Undergraduate Research (ISUR) program working on IIDAI-focused research.

6. Concluding Remarks

The first 2 years of IIDAI have seen the successful launch of the four technical thrusts described in this report, helping the institute reach its initial objective of developing strong bonds between UIUC and IBM Research while driving technical agendas in strategic areas. As we approach the completion of the initial set of projects defined during the launch, the institute is preparing itself to enter a new phase with a focus on a smaller set of larger research initiatives aiming at producing institute wide external impact and eminence. Areas under consideration include the development of a platform for accelerated discovery powered by state-of-the-art multi-cloud technologies, bringing together classical and quantum computing seamlessly to support sustainability and material discovery research. During the first half of 2023, steering committee members are defining and translating these objectives into thrust level activities that will drive the solicitation and admission of a new set of 2 year research proposals by 3Q23.

IIDAI Annual Highlights Report informs alumni, industry partners, peers, friends, faculty, students, staff and other stakeholders about the institute's accomplishments and newsworthy activity. This issue covers the fiscal year 2021-2022.

Appendix: Papers Published

Papers Published by Hybrid Cloud

1. Yifan Yuan, Jinghan Huang, Yan Sun, Tianchen Wang, Jacob Nelson, Dan Ports, Yipeng Wang, Ren Wang, Charlie Tai, and Nam Sung Kim. RAMBDA: RDMA-driven acceleration framework for memory-intensive us-scale datacenter applications. IEEE International Symposium on High-Performance Computer Architecture (HPCA), February 2023.

2. Ren, Wei, Junhao Pan, and Deming Chen. "AccGuard: Secure and trusted computation on remote FPGA accelerators." In 2021 IEEE International Symposium on Smart Electronic Systems (iSES), pp. 378-383. IEEE, 2021.

3. Pritom Saha Akash, Jie Huang, Kevin Chang, Yunyao Li, Lucian Popa, and ChengXiang Zhai. Domain Representative Keywords Selection: A Probabilistic Approach. In Findings of the Association for Computational Linguistics: ACL 2022, pages 679–692, Dublin, Ireland. Association for Computational Linguistics.

4. Bhavya, Bhavya, Jinjun Xiong, and Chengxiang Zhai. "Analogy Generation by Prompting Large Language Models: A Case Study of InstructGPT." 15th International Conference on Natural Language Generation (INLG 2022).

5. Zhang, Yu, Yunyi Zhang, Yucheng Jiang, Martin Michalski, Yu Deng, Lucian Popa, ChengXiang Zhai, and Jiawei Han. "Entity Set Co-Expansion in StackOverflow." IEEE BigData 2022 Workshop on Knowledge Discovery and Data Mining in IT Operations (BigData-IT).

6. Shin, Jungeun, Diana Arroyo, Asser Tantawi, Chen Wang, Alaa Youssef, and Rakesh Nagi. "Cloud-native workflow scheduling using a hybrid priority rule and dynamic task parallelism." In Proceedings of the 13th Symposium on Cloud Computing, pp. 72-77. 2022.

7. Archit Patke, Haoran Qiu, Saurabh Jha, Srikumar Venugopal, Michele Gazzetti, Christian Pinto, Zbigniew T. Kalbarczyk, and Ravishankar K. Iyer. Evaluating Hardware Memory Disaggregation under Delay and Contention. In Proceedings of the 1st Workshop on Composable Systems Co-located with IPDPS 2022 [CompSys 2022].

8. Guo, Hongpeng, Haotian Gu, Zhe Yang, Xiaoyang Wang, Eun Kyung Lee, Nandhini Chandramoorthy, Tamar Eilam, Deming Chen, and Klara Nahrstedt. "BoFL: bayesian optimized local training pace control for energy efficient federated learning." In Proceedings of the 23rd conference on 23rd ACM/IFIP International Middleware Conference, pp. 188-201. 2022.

9. Ashok, Sachin, P. Brighten Godfrey, and Radhika Mittal. "Leveraging Service Meshes as a New Network Layer." In Proceedings of the Twentieth ACM Workshop on Hot Topics in Networks, pp. 229-236. 2021.

10. Serif Yesil, José E. Moreira, and Josep Torrellas. Dense dynamic blocks: optimizing SpMM for processors with vector and matrix units using machine learning techniques. In Proceedings of the 36th ACM International Conference on Supercomputing (ICS '22). Association for Computing Machinery, New York, NY, USA, Article 27, 1–14, 2022. https://doi.org/10.1145/3524059.3532369.

11. Qiu, Haoran, Weichao Mao, Archit Patke, Chen Wang, Hubertus Franke, Zbigniew T. Kalbarczyk, Tamer Başar, and Ravishankar K. Iyer. “SIMPPO: a scalable and incremental online learning framework for serverless resource management.” In Proceedings of the 13th Symposium on Cloud Computing, pp. 306-322. 2022.

12. Mao, Weichao, Haoran Qiu, Chen Wang, Hubertus Franke, Zbigniew Kalbarczyk, Ravi Iyer, and Tamer Basar. “A Mean-Field Game Approach to Cloud Resource Management with Function Approximation.” In Advances in Neural Information Processing Systems. 2022.

13. Uwacu, Diane, Ananya Yammanuru, Marco Morales, and Nancy M. Amato. “Hierarchical Planning With Annotated Skeleton Guidance.” IEEE Robotics and Automation Letters 7, no. 4 (2022): 11055-11061.

14. Hannah Lee, James Motes, Marco Morales, Nancy M. Amato. “Parallel Hierarchical Composition Conflict-Based Search for Optimal Multi-Agent Pathfinding.” IEEE Robotics and Automation Letters (RA-L), vol. 6, no. 4, pp. 7001-7008, 2021. International Conference on Intelligent Robots and Systems (IROS) 2021.

15. Sirui Xu, Yu-Xiong Wang, and Liang-Yan Gui. Diverse human motion prediction guided by multi-level spatial-temporal anchors. In European Conference on Computer Vision (ECCV), 2022.

16. Junkun Chen and Yu-Xiong Wang. PointTree: Transformation-robust point cloud encoder with relaxed K-D trees. In European Conference on Computer Vision (ECCV), 2022.

17. Yuanyi Zhong, Haoran Tang, Junkun Chen, Jian Peng, and Yu-Xiong Wang. Is self-supervised learning more robust than supervised learning? In Proceedings of the International Conference on Machine Learning (ICML) Workshop on Pre-training: Perspectives, Pitfalls, and Paths Forward, 2022.

18. Mingtong Zhang, Shuhong Zheng, Zhipeng Bao, Martial Hebert, and Yu-Xiong Wang. Beyond RGB: Scene-property synthesis with neural radiance fields. In Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV), 2022.

19. Kuan-Ying Lee, Yuanyi Zhong, and Yu-Xiong Wang. Do pre-trained models benefit equally in continual learning? In Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV), 2022.

20. Kai Yan, Alex Schwing, and Yu-Xiong Wang. CEIP: Combining explicit and implicit priors for reinforcement learning with demonstrations. In Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS), 2022.

21. Ram Krishna, Thong Nguyen, Atom O. Watanabe, Dale Becker, Arvind Kumar, and Elyse Rosenbaum. A Methodology to Optimize the Number and Placement of Decoupling Capacitors in a Multilevel Power Delivery Network. IEEE International Conference on Electrical Design of Advanced Packaging and Systems, 2022.

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

IIDAI
104 South Goodwin Avenue MC-230
Urbana, IL 61801
217.333.1370
discoveryacceleratorinstitute.grainger.illinois.edu

NONPROFIT
ORGANIZATION
U.S. POSTAGE
PAID
CHAMPAIGN, IL
PERMIT NO. 453