COMPUTE-ENERGY NEXUS WORKSHOP REPORT

Klara Nahrstedt (Co-Chair)
Coordinated Science Laboratory
Siebel School of Computing and Data Science
Grainger College of Engineering
University of Illinois Urbana-Champaign

Philip Krein (Co-Chair)
Coordinated Science Laboratory
Electrical and Computer Engineering Department
Grainger College of Engineering
University of Illinois Urbana-Champaign

Technical Program Committee:
Deming Chen, Jian Huang, Rakesh Kumar
Elyse Rosenbaum, and Andrew Stillwell
(Electrical and Computer Engineering Department)

Local Organization Committee:
Baub Alred, Jenny Applequist, Karen Cromwell, Clint
Harper, Erica Kennedy, and Dr. Normand Paquin
(Coordinated Science Laboratory)

1. INTRODUCTION AND ORGANIZATION

On November 6, 2024, researchers from the Coordinated Science Laboratory (CSL), various academic units of The Grainger College of Engineering, and the National Center for Supercomputing Applications (NCSA) at the University of Illinois Urbana-Champaign gathered with industry representatives in Urbana, Illinois to discuss challenges, solutions, and research directions at the nexus of computing and energy.

Dean Rashid Bashir opened the workshop on Compute-Energy Nexus (CEN) topics and acknowledged the timeliness of the workshop. He pointed out to the audience that Grainger College of Engineering researchers are always at the forefront of grand challenges such as the one we are now facing, that of considering computing together with energy needs. Data are currently the fuel for artificial intelligence (AI), and data centers are critical to support AI workloads with sustainable and energy-efficient solutions. Furthermore, he stressed that the CEN topic is particularly important to the state of Illinois, the Governor of Illinois, and Chicago partners. The Governor of Illinois has already put forward the “Innovate Illinois” effort, which includes four initiatives: Beyond Silicon, Beyond Carbon, Beyond Biology, and Beyond Boundaries. This workshop greatly contributes towards the goals of the Governor’s initiative.

CSL Director Klara Nahrstedt presented the goals of the workshop. She pointed to news coverage highlighting that many leading computing companies have a very high demand for energy for their data centers. For example, news from Goldman Sachs indicates that AI is poised to drive a 160% increase in data center power demands [1], NPR said that “AI brings soaring emissions for Google and Microsoft, a major contributor to climate change” [2], and a New York Times editorial stated that “noisy, hungry data centers are catching communities by surprise” [3]. Many studies show a major increase of energy consumption between 2016 and 2022 due to the need to serve traditional data centers, cryptocurrency mining, and AI-dedicated data centers [5].

Nahrstedt stressed that the primary goals of the workshop would be to have discussions and engage in team building for major funding opportunities and to promote CSL, the Grainger College of Engineering, and U. of I. thought and research leadership in the CEN areas. The trends, published by IDC [4], clearly show that we need CEN conversations about (a) data center’ energy consumption, IT power, and carbon emissions forecasts; (b) decision processes and criteria for deploying workloads across hybrid multi-cloud architectures; (c) the impact of high-performance computing (HPC) and new AI workloads on data center resources; (d) adoption of high-density racks to support HPC and GenAI; (e) power scarcity and strategies to mitigate it (e.g., microgrids); (f) use of predictive and proactive maintenance to transform data center management; and (g) adoption of modular, containerized, and microdata centers, sometimes called ‘edge’ data centers, among many other topics.

After the introductory remarks of Dean Bashir and CSL Director Nahrstedt, the workshop *focused holistically* on two major CEN topics: (a) the compute-energy nexus inside of data centers, especially the role of application workloads, hardware and software co-design, and cooling to advance energy-efficient solutions; and (b) the compute-energy nexus outside of data centers, especially the role smart power grids can play in providing energy for data centers. Given the interdisciplinarity of CEN challenges, topics were addressed by a diverse group of speakers and panelists spanning many departments and interdisciplinary research units in the Grainger College of Engineering and the University of Illinois.

The morning was dedicated to the CEN topics inside data centers. Keynote speaker Wen-mei Hwu from NVIDIA, short-talk speakers Rakesh Kumar (CSL/Electrical and Computer Engineering (ECE)), Josep Torrellas (CSL/Siebel School of Computing and Data Science (SCDS)), Sarita Adve (CSL/SCDS), and Nenad Miljkovic (Mechanical Science & Engineering (MechSE)), and panelists Kathy Huff (Nuclear Engineering), Wen-mei Hwu (NVIDIA), Brett Bode (NCSA), and Naresh Shanbhag (CSL/ECE) presented and discussed the need for new energy-efficient rack designs, computer architecture designs, energy aspects of application workloads, energy associated with network interconnections, and many other issues.

The afternoon focused on CEN topics related to smart grid support for data centers. Keynote speaker Chris Gladwin (Ocient), short-talk speakers Andrew Stillwell (ECE), Philip Krein (CSL/ECE), Jian Huang (CSL/ECE), and Pavan Hanumolu (CSL/ECE), and panelists William Schaumann (Burns & McDonnell), Joe Jablonski (Ocient), Alejandro Domínguez-García (CSL/ECE), and David Nicol (CSL/ECE) presented and discussed the need to address the compute-energy nexus holistically, and to pay great attention to new directions in the needs for power grids, the limitations of power grid infrastructure, and many other challenges and potential research directions.

The report discusses both CEN topics and insights gained from the keynotes, short talks, and panels. In Section 2, we will present the CEN discussions that addressed needs, challenges, directions, and insights in the morning sessions. In Section 3, we will describe the CEN discussions from the afternoon sessions. We will conclude the report with a call for action and next steps in Section 4 that reflects the ideas of the workshop speakers, panelists, and audience participants.

2. COMPUTE-ENERGY NEXUS INSIDE DATA CENTERS

Energy usage in data centers is increasing at a rapid pace, and there is a major need for approaches to mitigate energy consumption and related cost. Keynote speaker Wen-mei Hwu from NVIDIA, together with the short-talk speakers and other panelists, addressed the needs and potential approaches to meeting them.

a. Keynote Address with the Major Questions “Why Do AI Data Centers Grow So Fast, and What Do We Do About It?”

The NVIDIA keynote speaker, Wen-mei Hwu, argued that the scale of AI computer clusters has been growing rapidly because of model training, and that data centers will continue to grow in the foreseeable future. Hwu said that the very first barrier AI data centers will face is the power wall, followed by the data scarcity wall, latency wall, chips production availability wall, and other constraints and their “walls.”

He pointed out that over the years, we have been, are, and will be struggling with connected clusters in data centers. In 2011, the Blue Waters high-performance cluster had to interconnect thousands of CPUs and GPUs, but even that interconnectivity was not sufficient for AI training because of power delivery. This led a Toronto team, using NVIDIA GPUs, to parallelize data and train AI models on individual GPUs, and hence exchange data relatively infrequently, and thus making progress in power delivery and speed for AI training. In 2016, NVIDIA correctly predicted that the density of GPUs needed to increase and introduced DGX-1 with NVLInk connectivity. This further solidified multi-GPU parallelism. In 2018, “transformer” AI models were introduced, that led to a substantial reduction of data labeling needs. However, it also led to an increased need to support matrix multiplications for attention training, and NVIDIA’s NVLInk supported distributed matrix multiplications, increasing the demand for NVIDIA products. In 2022, ChatGPT came out, requiring multiple tens of thousands of GPUs to train AI models. This development led to an increased gap between government/academia and industrial GPU infrastructures with respect to the numbers of GPUs available to train new AI models. In 2024, multimodal large language models required multiple hundreds of thousands of GPUs for training of AI models. This calls for not only for a high density of GPUs interconnected within a rack of cluster computers, but also high-speed connectivity among rack clusters. Options for the latter include InfiniBand and close relatives such as Slingshot networks. All these requirements lead to increased power demands, increased cost, and increased cluster sizes.

Hwu pointed out several challenges:

» Deep-learning AI models perform well or better when mimicking humans in terms of image and speech recognition, but do not perform very well for more nuanced language interpretation. One needs much more data to train AI models for these functions.

» We will exhaust the real data that humans generate using publicly available data. It is expected that by 2028, we will be using the entirety of publicly available data. There will be a need to create synthetic data to train future AI models.

» There will be a need to be able to complete multiple data passes over data to train future AI models as the datasets get bigger and bigger.

» Power will be the primary limiter for AI training beyond 2030 and the first big bottleneck.

» One of the major challenges is the scaling required for AI training.

> **Data Parallelism:** When we train a model, we have many pieces of data on which to train. To execute the training, we take a subset of images and form a batch. However, as data grow, one must divide each batch into sub-batches and assign them to individual GPUs. But if one divides batches, one must also adjust training parameters, which leads to more latency. Hence, one cannot indefinitely increase batch size, since delay adjustments can cause non-convergence of the overall training process. This places a limit on the use of data parallelism.
>
> **Model Parallelism:** Whenever we have a batch, we may split the model into different pieces and run the batch through these pieces. One may also do pipeline parallelism, but this requires the use of distributed matrix multiplication across GPUs. That in turns, translates into increased communication and high bandwidth capacity to support this type of parallelism.

» The high interconnectivity of GPUs inside of racks (cluster computers) and between clusters (racks) will require high bandwidth interconnecting networks. Fiberoptics would be highly desirable due to high bandwidth availability. However, introducing fiberoptics everywhere is expensive. For example, consider a 2-million-GPU data center with 2 levels of interconnects, with L1 switches connecting GPUs within a rack (cluster of computers), and L2 switches connecting GPUs across racks. If fiberoptics was used for both L1 and L2 interconnectivity, the estimated cost could be as high as $7.2 billion at this scale. On the other hand, if one were to implement L1 interconnectivity via copper and L2 interconnectivity via fiberoptics, the cost could come down to as low as $860 million.

» The density of GPUs brings not only increased power challenges, but also thermal and cooling challenges. The growth of GPU speeds and the mechanical density (number of GPUs per rack) involve growing generation of heat, and there is a pressing need to dissipate heat more effectively. Air cooling in larger and larger data centers will not be sufficient, and one must switch to liquid cooling. Today we can accommodate 23–64 GPUs per rack, but this is expected to increase to 100+ GPUs per rack creating tremendous energy demands and requiring much greater heat dissipation capabilities.

Note: L1 domains will grow with GPU rack density, and this direction will support increased model parallelism. If fiber optics costs come down, we may see the L1 domain interconnected via optics in the future, which would enable an even further increase of GPUs per rack. Another factor to consider, however, is the current lower durability and reliability of fiber optic connectors as compared to copper connectors.

Trends in AI training suggest that the scale of data centers will increase tremendously. In the next 4–5 years, we will be moving towards 300,000 to 5 million GPUs in each cluster requiring 500 megawatts (MW) to 10 gigawatts (GW) of power.

The challenges of providing such a power supply include:

» The speed at which we can approve and build new power infrastructure? Power grid changes, e.g., changes to the power infrastructure to increase capacity, take approximately 10 years.

» Will nuclear and renewables provide the needed power supply? Coal energy cannot grow; it will actually shrink. With nuclear energy, while "clean," one must consider safety and potential major impacts if disasters occur. Renewable energy involves the challenge of variability due to daylight cycles and weather changes.

Insights and directions to improve power efficiency from the compute side are:

» **To increase performance per watt in AI infrastructure:** do chip-level reduced precision, do rack-scale optimization, and reduce data movement.

» **To reduce energy consumption via algorithmic approaches (e.g., distillation, pruning).**

» **To reduce energy consumption with new innovations in data and training model approaches.**
For example, current training assumes data are given, and every time there are new data, models are retrained from scratch. It is desired to redefine and put structures/relations into training data so that even if new data arrive, the models do not need to be retrained from scratch. If one has higher-quality data with less redundancy and noise, this will also reduce computation and energy needs.

Hwu concluded with several discussion points:
(a) One should understand AI training paradigms in terms of power usage (e.g., conventional data centers use a cyclic power paradigm, but AI data centers do not). (b) We must better understand how the human brain works and how we learn. For example, if a person needs to know something, there is likely to be a reference that the individual will seek to find the needed information on the fly. Current AI models do not reference needed information on demand to train better; rather, AI models grow bigger and bigger. (c) One should consider model performance for different domains and their data scarcity (data wall). For example, for medical applications, limited training data will not be a problem for a few years. What we will need is specialized training models. (d) AI inference is not a big problem in terms of computation since it is less compute intensive. However, inference must be cheap.

Hence, we need better business models.

b. Short Talks

Speaker Rakesh Kumar addressed data-center concerns through large-scale integration. One can greatly improve energy efficiency through a significant increase in the number of compute components integrated on a single substrate. For example, significant energy efficiency benefits are possible from integrating a large number of components on a water. However, doing so also introduces challenges, such as the delivery of power to each compute component, and cooling, heating, and other issues. Wafer-scale computing was considered in the early 1980s for improving performance, reliability, and reduction of jumbles of wires. In 2018, Kumar and his team revisited wafer-scale computing to improve the energy-per-bit performance, arguing that the cost of communication between components on the wafer is orders of magnitude lower than the cost of communication which, in turn, is much cheaper than communicating across data center racks.

However, wafer-scale computing presents the following challenges:

» If one puts a large number of compute dies on a wafer, how does one cool the system and deliver the power to it, and interconnect them? If one starts to consider liquid cooling, thermal solutions, and power delivery, one could end up potentially supporting only a small number of dies on a wafer, losing the integration benefits.

» Another challenge is internal and external bandwidth density. Let's say that one is building network switches to interconnect GPUs. How does one build network switches with large radices (where a radix is the number of bidirectional ports on a switch)? One direction is to integrate many smaller switches onto a wafer. This could give the illusion of a single large switch and potentially get a radix as high as 8,192. However, this solution requires a large number of wires on the wafer to connect these switches – the internal wiring itself may become a bottleneck. Similarly, the external wiring that connects this network switch to outside needs high bandwidth to keep the switch busy.

Insights and directions for energy efficiency are as follows:

» The very large integration of compute and communication components is very promising for energy efficiency and for bringing down system-level costs. However, power delivery, cooling, and interconnects are major bottlenecks to the realization of the promise.

» We must think of system cost, not device cost. A wafer-scale system will cost millions of dollars, but if a system is well designed, it will reduce the overall cost even if the compute cost may be much higher.

Speaker Josep Torrellas addressed energy inefficiencies and potential solutions in data-center infrastructure. One must leverage hardware acceleration and integration of compute and communication components, minimize data movement, co-design hardware and software innovations, and integrate security and correctness from the ground up. To achieve these goals, we must have evolvable computing, which means we must consider the design of composable accelerator hardware, composable and flexible compute acceleration, orchestration of ensembles of accelerators, automatic generation of compilers, and distribution and orchestration of codes for different accelerators. This kind of evolvable computing can deliver energy efficiency.

However, evolvable computing brings the following challenges:

» We need near and in-memory storage (IMS) acceleration with ubiquitous IMS blocks in the memory hierarchy, network switches, and SSDs. Torrellas gave an example of a conceivable way to go towards a solution, the XFM: Near-Memory Acceleration of Software-Defined Far Memory.

» We must consider flexible and reconfigurable networking and storage since hardware is often underutilized, consuming energy for nothing and delaying execution. Reconfigurable optical interconnects can adjust network topology based on application requirements. A logically centralized storage manager can steer I/O to specific SSDs based on knowledge of network topology and reconfiguration delays. Torrellas gave an example of a solution, the CC-NIC: A Cache-Coherent Network Interface Card (NIC), to greatly improve the efficiency of NIC cards, since PCIe communication represents 80% of intra-rack round-trip latency.

» We need computing in network switches, since networks could then monitor and analyze network traffic, run security codes, and detect anomalies. Torrellas specified solutions, such as Homunculus to automatically generate efficient ML models to run on switches, and Caravan to learn online and check for malware.

» We need self-balancing systems during runtime such that accelerators do not wait for work and waste energy. If we had self-balancing systems, this could minimize data movement.

» We need energy-efficient and accelerated distributed transactions such as key-value stores, log management, and lock managers. Torrellas presented an application-customized networking stack with eBPF, called Electrode. eBPF accelerates distributed system network software while maintaining safety. Electrode accelerates consensus protocols with new communication primitives in eBPF.

Torrellas pointed out several insights and directions. (a) Limiting the energy consumption of distributed computing is a pressing research problem. (b) Evolvable computing must consider multiple cross-layer specialization and adaptation techniques. (c) Evolvability will need to take a fresh look at the programmability of these systems.

Speaker Sarita Adve addressed energy challenges related to immersive computing. Immersive computing is about the seamless integration of virtual worlds with physical environments. The virtual world may include virtual reality, augmented reality, and mixed reality (which together are “extended reality,” or XR).

However, immersive computing brings the following challenges:

» XR technologies must be comfortable, mobile, trustworthy, able to last all day, and able to offer rich experiences. For XR technologies, power is a first-order constraint along with thermal comfort, battery life, and form factor.

» To achieve XR needs, we need hardware acceleration, co-design approximation, and distributed XR with computation offloading to a co-designed device/edge/cloud ecosystem.

  » Hardware acceleration would have great benefits for immersive computing, but it cannot support the speed at which we would like algorithms and models to evolve.

  » Co-design approximation has great potential for XR, but it is not clear how to model and measure the impact on user experience.

  » Distributed XR can respond better to algorithms and model evolution and would avoid investment in new hardware, but it would impact network latency and bandwidth.

» One of the major challenges for immersive computing is the lack of open reference systems and benchmarks. This is a barrier to immersive computing research. Adve discussed her ILLIXR (Illinois Extended Reality) testbed to support XR research.
Adve pointed out several insights and directions. (a) The XR goal is to provide comfortable, mobile, all-day trustworthy end devices and an ecosystem providing rich experiences for billions of users. (b) This XR goal faces insatiable, unsustainable demand for compute, and therefore energy, in data centers. (c) We need to explore co-designed device-edge-cloud research to address latency, bandwidth, quality, and energy issues along with co-designed hardware runtime and network models.

Speaker Nenad Miljkovic addressed holistic rack-to-processor power and thermal co-design for future servers. He pointed out that in the past, computing designers designed systems without thinking of thermal problems. He argued that the energy consumption of conventional data centers is already growing at an extremely fast pace, and AI data centers will only add to demands for power and increase the need for effective thermal handling solutions. He concluded that we must co-design for computing and thermal problems together. He showed graphs that made key points, e.g.: (a) Average power usage effectiveness (PUE) plummeted between 2007 and 2018. (b) Typically, 40% of power consumption is directed to cooling in air cooled data centers if used for cooling, and the rest for compute. (c) Air cooling via HVAC itself (large chillers) needs a lot of energy to cool the IT load (up to 40% of the power coming into the DC). (d) Air cooling is inefficient (e.g., for the new NVIDIA Blackwell chip, air cooling cannot be used); hence, we need architectures with liquid cooling and Nvidia is recommending this. Otherwise, the chips are massively throttled back and are not being used at their full potential.

However, to achieve efficient liquid cooling, we must solve the following challenges:

» Can we more efficiently direct heat rejection to the ambient outdoor environment, even in hot summer weather? Right now, we use up to 40% of the incoming power to achieve this, power that could be used for AI workloads and compute, that is being used for cooling instead.

» How can we reuse heat if we move away from air cooling? For example, we could use it for heating water, or other usages in a neighborhood in close proximity to the data centers. This is already being done in Europe to some extent, where massive district heating and cooling systems are already prevalent.

» How does one bring cooling closer to chips to minimize thermal resistance while avoiding reliability issues? The lower thermal resistance allows heat capture at higher temperatures, reduced energy consumption for cooling, and reduced thermal energy wastage in the case of heat re-use. Miljkovic discussed his solutions which are being developed for the rising trend of multichip modules such as the Nvidia Blackwell. These use passive cooling (no input energy to drive the liquid flow) and achieve self-regulating thermal management with buoyancy-directed flow and high temperature heat capture for heat reuse. This project was being conducted with Nokia and Carrier.

Miljkovic mentioned the following insights and directions in this area. (a) One should also consider water quality. (b) Immersion cooling is coming, but the main issue is compatibility and component (e.g., capacitor) stability where fluids interact. If we move to immersion cooling, the whole data center and rack structure would have to change. (c) One can consider moving to liquid cooling just by retrofitting and modifying the cooling system of an existing air-cooled data center (for example, with Air Assisted Liquid Cooling (AALC)), but this is not possible with immersion cooling due to the scale of required retrofitting and incompatibility with existing data center cooling architectures. Hence why the Bitcoin mining sector has fully adopted immersion cooling with open arms (many new and capital-intensive Bitcoin mining data centers have been built in the last decade), while data center operators have not (many existing and very large data centers that are already air cooled and operating).

c.Industry and Academia Panel on Energy-Efficient Data Centers

Panelists were asked to reflect on some questions given by the organizers as well as some of their own questions:

» What can we do with software/hardware to make data centers energy efficient and dynamic?

» How do we build system infrastructure to support large language models in an energy-efficient manner and at extreme scale?

» What are the barriers to carbon-free data centers?

» What workforce issues are involved in enabling design and implementation of future energy-efficient and carbon-free data centers?

» In the end, data centers generate heat. How do we manage data centers, thermal loads, and how do we enhance their performance 100× in an energy sense?

Wen-mei Hwu opening remarks: One of the important criteria for industry today is how quickly a customer needs additional computing, say another 1,000 GPUs. There is an important measure of how many days it takes from an order until the system is up and running. NVIDIA spends years working on that number and ways to reduce it); it is an important part of the success of any company in this domain.

He provided several of his own questions:

» How many days after a Blackwell chip comes out does it take for an entire system to be delivered to customers, up and running? Right now, we can do it quickly. But the power infrastructure is a challenge: if a new system needs any additional power supply, even a new power feed, that is an order of magnitude more latency. People literally use truck-mounted generators, but it is very expensive and inefficient.

» He would really like to understand how quickly we could move energy around. It is a fundamental problem we have not thought about sufficiently in the past.

» In addition, we have a huge amount of inefficiency in heat and power supply conversion. We should be able to cut down energy wasted in that process. How do we store energy? How do we convert energy waste into an asset?

Brett Bode opening remarks: He started working at the University of Illinois in 2008 and joined Blue Waters at NCSA. The Blue Waters building is still nice, but not what they need today. They're still working on various challenges, such as planning for increased loads beyond 24 MW capability and running campus chilled water systems for cooling.

» They have a lot of unused capacity in the facility. For the National Petascale Computing Facility, at the time, they did not build for the water capacity the facility needed.

» They have problems with below-dewpoint water, but do not have funds to fix it.

» Delta and DeltaAI use 480V power systems for input. It is a modest-scale system that uses four water-cooling technologies. Blue Waters was basically an air-cooled machine with retrofitting. Delta has direct-to-chip water cooling, but they extract only 70–75% of the heat load with it, so they must provide air cooling as well. It is basically a water-cooled air-cooled system. The CPU cabinets are very dense 60 kW cabinets. A100 cabinets are way less dense. DeltaAI computer racks use no fans; they are 100% water-cooled. This same technology went into Frontier, etc., which has had problems, but the bugs seem to have been worked out before it came here. NCSA still has lots of room in the building.

» Funding is an issue. They are working to put other high-density projects in there. The total capacity of Delta and DeltaAI is about 3 times what Blue Waters could provide.

Katy Huff opening remarks: She works on provisioning nuclear energy rather than power use in data centers. She was in the Biden-Harris Administration and helped establish the net zero-energy by 2050 goal. This is a very ambitious goal with a need for 550–770 new gigawatts of clean stable power, of which about 200 GW would come from nuclear energy. This would require about double the current U.S. nuclear capacity. Nuclear today provides about 20% of electricity in the U.S. We have a bigger mountain to climb from data centers just since 2021. Data centers could represent as much as 24% of global electricity capacity soon. The appetite for power has grown.

She advises Amazon on deploying nuclear for their data centers. Amazon is a good example of ambitions for nuclear energy. Amazon is spending $500M on the development of small modular reactors. Nuclear power typically runs all day, every day, running at approximately 92% of the time or more. Existing plants must shut down for 1.5 months every 1.5–2 years to replace fuel. If we could instead rotate little pebbles of fuel through a unit, we could extend the time the power is up. Nuclear has the highest capacity factor [fraction of time when electricity is actually delivered] of any existing source. Any new energy plant deployment, however, takes years, in part to obtain required regulatory approvals.

Naresh Shanbhag opening remarks: He pointed out many energy efficiency challenges in data centers. Large AI models need distributed platforms. Connectivity and computation go hand in hand; it is not just about compute. If we take one bit and transmit it, the energy consumption is orders of magnitude higher than on-chip compute. We cannot ignore the cost of communication; we must account for it. There are three pillars of energy efficiency: power delivery, compute (faster GPUs, accelerators, etc.), and connectivity (how we connect them). We should think of each pillar as a full stack, with software, architecture, circuits, and devices. Each pillar's challenges are being addressed in a siloed fashion. The CuBIC JUMP center he is in is looking at connectivity: how we get data across an optical or electrical network as fast as possible with as little energy consumption as possible. The CoCoSys JUMP center he is in is about building new AI algorithms to be efficient, to reduce workload without reducing inference quality. Power delivery is not showing up as a first-class citizen. That is the gap. We must find a way to bring the three pillars together. There are partial solutions, i.e., there is at least one connection between power delivery and compute; but we need a lot more. We need to unify those three pillars and treat them as a unit. Power delivery must be a first-class citizen. The cross-domain and cross-layer system optimization problem is hard; the challenges are both technical and in workforce development.

The panel discussion provided several interesting insights:

» What is our goal in making something useful? The panelists pointed out the following directions: (a) identify pockets of opportunities in data centers to combine energy, computing, and communication pillars; (b) identify metrics such as performance, reliability, carbon impact, and others, then focus on efficiencies and transmission issues, and suggest couplings between energy generation and chips in data centers; and (c) encourage measurements of production systems.

» What are the CEN topics and problems for which academia would not compete with industry? The panelists discussed (a) the potential for a data-center testbed to study new problems, (b) the creation of simulation and emulation models of a data center where one could explore plug-and-play operation with different solutions, and (c) new workloads of tomorrow to be experimented with.

» Is it possible to have software, workloads, and energy resources work together so that data-center workloads can be distributed geographically where resources are available? The panelists offered several points. (a) Reduced scale of nuclear reactors and their smaller designs are possible and would lend themselves for distributed energy sources. However, geography must be chosen very carefully because of water, safety, and other constraints. (b) Location of data centers close to power sources is an option (e.g., see Finland solutions) to reduce loss of energy owing to transmission and carbon emissions. (c) Data transmission can be a bottleneck. (d) Reliability of data centers needs to be considered, since components fail frequently in hybrid data centers. Software-hardware solutions must tolerate failures.

» Is it possible to place data centers under oceans, have floating or underwater nuclear reactors support data centers, or have data centers offshore close to wind turbines? The panelists provided the following insights. (a) Permission to have an underwater or floating nuclear reactor is a big issue. (b) If one puts data centers on the ocean floor or floating, machines need to be stabilized. For example, for mobile data centers, all connectors would need to be redesigned. (c) Placing data centers offshore involves domestic and international issues, such as public objections, placement in international waters (that would include a foreign policy issue), and others.

» How can water scientists help with cooling from micro/nanoscale levels to the facility scale level? The panelists offered the following input. (a) Data centers are moving to water cooling, but that is exposing a lot of issues with mineralization and organics. (b) Advanced nuclear reactors are using coolants other than water, but the secondary loop is creating steam and spinning a turbine. There are still related water quality issues. (c) There is an important line between isolated cooling water versus immersive cooling. Heat sink cooling is problematic, and there are circulation problems and industrial problems. If we really want to solve the water problem, over the long run, the entire subsystem will need to be immersed in water with a modular replaceable cooling system.

3.COMPUTE-ENERGY NEXUS AND THE SMART GRID

Smart grid design and implementation is an integral aspect of next-generation data centers, and keynote speaker Chris Gladwin, short-talk speakers, and panelists addressed the interconnections in smart grid and data-center co-design.

a. Keynote Address on the Growing Data-Center Energy Needs

The afternoon keynote speaker, CEO Chris Gladwin of the Ocient company, acknowledged that computing is the most successful industry ever, and said that its growth seemed to be endless until last year, when the lack of energy became a limiting factor. Energy efficiency suddenly became a big deal. For example, Loudon County, Virginia, has more data centers than Europe, but they are now out of power. Similarly, Dallas, Texas, is down to 1% of power availability. The need for more power is driven by (a) conventional data centers, whose power needs are anticipated to grow by 50% in less than five years (2022 through 2026); (b) cryptocurrency, since bitcoin mining has grown; and (c) AI data centers, which are growing at an accelerated rate.

There are several challenges:

» We must deal with large-scale datasets, complex data preparation, data duplications, platform licenses, and administration.

» There are new governance and security challenges with each target environment.

» Different skillsets are required for data engineering and for each target environment.

» We must deal with multiple clusters where different systems do different things in an energy-efficient manner.

Gladwin offered the following insights and directions: we must (a) develop and promote energy-efficient hardware and software solutions, (b) collaborate with industry partners to drive adoption and awareness of energy-efficient solutions, (c) create common energy-efficient measurements and interfaces, and (d) drive customers to purchase and deploy energy-efficient systems.

The Q&A session with Gladwin brought several insights:

Q: Is the goal to make data centers more energy efficient, or should society aim to bring the embodied carbon footprint down?
A: If one optimized the data centers for energy, the price would be lowered, but then more consumers would come in. Every time the IT industry makes something more efficient, it never happens that people maintain their previous level of use. In fact, people will start to use vastly more. It follows the Jevons paradox in economics, i.e., “when technological advancements make a resource more efficient to use (thereby reducing the amount needed for a single application) … [then] the cost of using the resource drops, [but the] overall demand increases causing total resource consumption to rise.”

Q: Why don’t data-center operators provide consumers with any information on how much energy use/carbon footprint their workloads have?
A: Many data-center providers and operators do not [or cannot] even calculate energy/carbon footprint per workload, let alone optimize for it.

Q: Do we need new undergraduate programs or interdisciplinary programs to upskill the workforce in compute-energy nexus topics?
A: Energy-efficient, large-scale computing is a big deal, but the power grid/energy industry and IT industry have not talked or collaborated closely. Changes are needed, including new directions for education and setting of standards for energy- and carbon-efficient data centers. The Midwest and Illinois could be the center of this synergy and collaboration efforts between these sectors.

b. Short Talks

Speaker Andrew Stillwell discussed the “last mile” of power conversion in data-center power delivery. One must consider the point of load (PoL) converter that supplies power to processors. There has been rapid growth in processor power consumption, driven by generative AI, and it is necessary to innovate on processors working under low-voltage regimes. Thermal load must be also managed since energy is dissipated in the form of heat.

There are several challenges to consider:

» With the conventional approach, one uses an intermediate bus converter architecture to bring the desired power to processors: one starts with about 48 volts and needs to deliver 1 volt to the processor. This conversion usually happens in two stages using a transformer or switched capacitor circuit configuration. Then there is a multiphase regulation stage with many inductors. All these components take space on motherboards and the requirements are very challenging, i.e., one needs to deliver ~2,000 A current at less than 1 volt while meeting slew rate demands and minimizing power conversion losses with a ~2-kilowatt heat load right there.

» The thermal management of power delivery must be solved for delivery in tight spaces.

» With the growth of data centers, there are challenges with point-of-load converters.

» How does one push power into the package? Package power is the power that is consumed by the CPU cores, cache, and other computing elements.

Stillwell pointed out several insights and directions:

» Industry and academia have different approaches to how power conversion and packaging are accomplished. For example, Google uses a multi-stage process for vertical power delivery, whereas academia considers single-stage power delivery with switched-capacitor stepdown, reducing inductor voltage, and other conversions to yield high power system efficiency.

» Funding opportunities are driving thermal management solutions for power delivery.

» Answers must be found to adding more PoL converters and being able package them into small spaces.

---

[1] https://en.wikipedia.org/wiki/Jevons_paradox

Speaker *Philip Krein* presented compute-energy nexus needs in data centers as a system of peer requirements in which energy and information processing need to be treated as peers. This means that there must be a collaborative co-design to obtain the lowest system energy per digital operation and optimal use of resources.

There are several challenges to be solved:

» We are facing a power wall, and optimization makes no sense unless we consider power conversion. However, in many calculations and designs, power conversion is ignored.

» Conventional designs treat energy and information separately, and we deal with conventional silos. For example, in a typical AC architecture, with building-level uninterruptible power, there is a sequence of energy losses that stack up. In a typical DC architecture, with building-level uninterruptible power, there is some reduction of losses and better reliability. To deliver power to digital circuits, one must look at the full stack of conversions where the delivery happens in parallel with multiple cascaded conversion stages. In present systems, the end use involves high current and low voltage.

» Another major challenge in this space is pushing less than 1 volt at thousands of amps when conducting stacking. We are pushing the limits of Kirchhoff's laws by a lot. We are fighting the physics of circuits.

» Barriers and further challenges include the need for on-chip extra steps of isolation and protection, 3D packaging, and the need to push balanced and dynamic computing.

Krein pointed out several insights and directions:

» There are several system opportunities. (a) We need to enable software processes to control their own slew rates (slew rate is the change of voltage or current, or any others electrical or electromagnetic quantity, per unit of time). (b) Data centers need to limit power ramp rates (going to 100 MW/s is huge and even 10 MW/min is high from a grid perspective). (c) We must enable redirecting and rebalancing of loads and be a dynamic participant in an intelligent grid, and we should cross-communicate and track stochastic resources. (d) Processors and modules can be in series.

» In general, we need to get past brute-force silos to find system answers.

» We must coordinate solar and wind power sources and leverage 3D packaging for power management.

» We need to support much wider voltage adaptation ranges.

» We should adapt dynamic software management at the mobile-device scale to address needs at the full data-center scale.

To push through the power wall, the power delivery system in data centers needs to be reorganized. One approach would be to put units in series and achieve high voltage and low current. This would support low-energy computing more directly, given balanced computing. One could organize the units at many levels: the server level, blade level, board level, chips in series, or all levels. However, putting cores in series is nontrivial. Putting chips in series is easier than putting cores in series.

Speakers Jian Huang and his PhD student, Jinghan Sun, discussed modular data centers (MDCs). MDCs are compact, containerized, portable, easy to deploy, and easy to co-locate with renewable energy farms. These are sometimes called 'edge' data centers.

However, there are several challenges to consider:

» Renewable energy is affected by temporal variability of power, i.e., solar power follows diurnal cycles and wind power depends on weather. Batteries can help, but they introduce extra carbon emissions and cost. A challenge is on how to utilize renewable energy in an efficient manner and create/assess "stable farms."

» If one assumes "stable" renewable farms, the question is how one selects the "best" renewable farm, and how one identifies complementary renewable sites and optimizes workloads.

» Another major challenge is how to build a sustainable AI infrastructure; this will require consideration of the server architecture, system virtualization for all chips, AI model/hardware co-design, and carbon-aware operating systems.

Huang and Sun provided the following direction:

» Renewable energy production hours can have good predictability with tools such as the SkyBox framework [6] as a step towards "stable" renewable farms. SkyBox employs a learning-based approach for platform operators to explore the efficient use of renewable energy with MDC deployment across geographical regions. It means that SkyBox can help data-center operators build optimized renewable-based modular data centers by intelligent grouping of renewable energy farm sites that would provide stable aggregate production, as well as smart VM placement and migration of VMs for continuous computation.

Speaker Pavan Hanumolu explored data-center connectivity when running AI models and workloads. Over recent years, we have seen astronomical energy consumption, with 75% going towards computing and 25% towards communication. One must carefully consider distributed computing and storage since communication between processors at all levels is a bottleneck. We note that communication on a chip is very efficient, communication via copper is less efficient; and communication based on optical networks is very costly.

Hanumolu pointed out the following challenges:

- How can we improve edge bandwidth density? We can push more data onto a single wire, but the cost goes up exponentially after some point.
- How do we build interconnects between processors in an energy-efficient manner? Is it even possible? Advanced packaging could help, where we could increase the number of wires by inventing a new package. But that might lead to power delivery issues.
- How do we get optical networks to 400 or 800 GBps, and even to 1.2 terabytes (TB) per second? This challenge will require key research innovations in optical devices, communication integrated circuits (ICs), compute and memory, advanced packaging, power management ICs, and power devices.

Hanumolu engaged in discussion with the audience and the following insights and directions were brought out during the Q&A:

Q: How much of the communication bottleneck is due to processing?
A: The connectivity bottleneck is at all levels. For example, NVIDIA has a memory to GPU bottleneck, and an optical network at a 500+ meter range has a bandwidth constraint.

Q: What are the challenges of switches in data centers?
A: The challenge is that we cannot make switches any bigger. We can only put so many wires around them. Hence, they are limited in terms of their bandwidth (50 TB per second). There is no clear path to scaling.

- There is the notion of “switch-lets,” but it is not clear if one could build a 200-TB-per-second switch using four 50-TB-per-second switches. This brings networking, protocols, and I/O challenges with it.
- One could move to optical networks with electric switches, or purely optical networks with optical switches. But the optical to electrical and electrical to optical (OEO) transitions are expensive. So, one should avoid OEO, and once one is in the optical domain, stay in the optical domain.
- Optical infrastructure is expensive, and one must know traffic patterns in advance; that limits the usage of some applications.
- Switches need to consider power, especially if traffic congestion occurs.
- Optical switches and programmable switches are gaining in importance.

c. Industry and Academia Panel on the Smart Grid and Energy-Efficient Data Centers

Panelists were asked to reflect on the following questions:

- What can we do with software/hardware to make data centers energy efficient and dynamic with renewables and the power grid?
- How do we make AI-centric data centers active load participants in the power grid?
- Where do we see scaling going? What are the general grid growth and impact issues?
- How do we manage a facility’s overall footprint, including co-location of data centers and smart grids, and their impact to minimize environmental impact: on-site renewables, storage, noise management, housing, buffers, and proper neighborhood planning?
- What are the barriers from the smart grid side to achieving carbon-free data centers?
- What are the workforce issues related to achieving a holistic smart grid and data-center designs and implementations?

Alejandro Domínguez-García opening remarks: There is a question as to how data centers can be active load participants. There is lots of control over power generation, but not a lot of control on the load side. If we are trying to integrate renewable power, we need to get some control over load, say for electric vehicles (EVs) or HVAC systems in houses. Industrial loads are inelastic though. There are opportunities to pursue algorithmic design.

A second question is: what are the barriers from the smart power side to achieving carbon-free data centers? In the U.S., there are not a lot of connections/population where wind is available, for example.

Joe Jablonski opening remarks: He talked about a problem Ocient is focusing on, to give a sense of scale and why they have an energy problem. For a customer, they had to tap every fiber line in the country to get 300 terabits per second. One goal was to monitor all the data flows. That is roughly 1 billion metadata records a second. They had to load the data flows, which meant 3.6 trillion rows per hour, 86 trillion rows per day, or 31 quadrillion rows per year. That is roughly 60 petabytes of storage. It only takes ~$13.5 million in hardware to build that out. We can now afford to build systems like that. That system requires 350 kW of power. We can now build these systems at scale. That is what is really driving the economics of this.

Jablonski posed a lot of questions, such as:

» How do we use less power? There are all techniques for bypassing operating system kernels as much as possible and maximizing hardware efficiency. Database systems are idle 20–60% of the time.

» How can we create an execution plan to minimize usage of power during idle times? How do we get power consumption during idle time to 0?

» And finally, how do we time-shift? A big problem is that measuring power in a computer system is incredibly difficult.

» How do we establish a real framework for monitoring power? How do we find out, if we run either Algorithm A or Algorithm B on a query, which consumes less power?

» And how do we time-shift workloads to manage peak power demands? The scale and the cost of hardware are driving new large use cases that drive power consumption.

David Nicol opening remarks: A question that jumped out for him: how to make data centers dynamic and efficient? We want better living through models, algorithms, and value-based optimization. When we have a problem domain and we find optimizations to do anything better or faster, we are taking advantage of some characteristic of that domain. What is characteristic of the kinds of things we are talking about with renewables and data centers: *variability in load and controlling the load and taking advantage of time-shifting.*

If we are considering *value-based* computation, we need to talk about efficiency in energy consumption, operations per joule or flops per joules, and how much computation we get. We often assume that all flops are equal, which they are not. If we have a global optimization algorithm, the longer we run, the better the answer we get. But if we put a value on the extra return we obtained from running longer, early flops are much more valuable than later ones. There is a trade-off in quality/resolution of the answer vs. the amount of energy we put into it. Think of utility per joule. Efficiency is about organizing things to increase that metric. AI is emerging as dominant. LLMs consume a tremendous amount of resources. There must be space in there for working on quality of solutions vs. amount or type of training we do: there must be a knob in this space.

Bill Schaumann opening remarks: He does a lot of work in data-center design and implementation. He answered several questions.

Where do we see scaling going? And what are the general grid growth and impact issues?

» It depends on the growth of the market. Currently, there is a rush to deploy mega AI data centers (1 GW+ data centers on a single campus). It is expected to continue at this rate for 3–5 years, but there's disagreement on when it will taper off; but we expect the colocation data-center market to adopt some of the AI data-center load, much as it did in the cloud computing market. Colocation centers are provisioning now for AI. 60 MW and 80 MW AI deployments will be common on smaller campuses.

» We will not typically see just one data center. We will see 3, 4, or 5 on a campus. No one wants to build all that infrastructure just for 1 building. We will still see 500 MW campuses going in all over the place, but usually in clusters, and smaller campuses scattered like cloud colocation data centers will start to become more prevalent. The overall demand will still grow, but it will not be as centralized after the next few years. It will be more spread out. We are going to be hunting and pecking to find the power; they need to find sites where they can put the power. We are clustering where the power is available.

How do we manage a facility's overall footprint, including co-location of data centers and smart grids, to minimize environmental impact: on-site renewables, storage, noise management, housing, buffers, and proper neighborhood planning? This all comes down to how we build our data center. What are your parameters?

» We need to define the requirements of an AI-centric data center. There has been an increased densification of servers. When Schaumann started doing data centers 18 years ago, a 5-kW rack was a heavy rack. Then we went to 15 and then 50 kW rack; that was last year. Now with water-cooled, we are seeing 150- and 400-kW racks. But the buildings are not any smaller, because of the required physical infrastructure.

» There's the issue of what we do for reliability. Are we at three 9s (up 99.9% of the time) or five 9s (up 99.999% of the time)? For five 9s, we may need to have generators, active harmonic filters, uninterruptible power supplies, super capacitors, etc.; all these things depend on your load profile, etc.

» Then there are cooling technology options (water to the chip, immersion cooling, rear door heat exchanges, air- and water-cooled chillers, fan arrays, cooling towers). All that affects the outdoor footprint.

» People do not realize: the noise from a generator is nothing compared to the noise from a chiller or a cooling tower.

» Then we must look at the power profile: flicker and cycling of loads. Depending on our load profile, we must consider all these things.

» It depends on how our client is deploying their AI server. Some vendors say they do this through a server, by modulating, by going to a heat sink; some are saying other things. For now, we are expecting clients to handle it on their end through server modulation somehow. But it is one of the unknowns for large deployments, because it has only been beta-tested at low power levels.

The panel responded to questions from the audience and provided the following insights:

Q: What is the projection of future data center growth? Does it include renewable energy and EVs? Is future grid development part of the projection?
A: From a data-center perspective, the current projections include data centers and the grid. Everyone knows that we need to build up, especially the grid, i.e., transmission lines need to build up, and substations are the biggest issues. Another issue with data centers' growth is proximity/location with respect to power sources and Internet networking, since many networking lines are needed to bring the workloads in and out.

Q: Are there any new approaches industry is trying for scaling up the power grid as data-center scale is growing?
A: Industry is changing its approaches because of chip manufacturers and site and power requirements. This means the goal is to build flexible data centers that can change on the fly.

Q: Do you see data-center providers taking charge of power generation and distribution just for their own purposes?
A: Until recently, data-center providers have never engineered power or cared to take charge of power generation. There was always enough power. But today we are operating in a power-constrained environment with a political problem of who gets power and who does not. Another point is that owning a power grid brings a lot of responsibilities that data-center providers are not aware of.

Q: Power is transmitted to meet the peak rate, and data center loads are very bursty. Are there smart-grid/data-center solutions for provisioning dynamic power loads or operating on an average power rate with battery supplements to satisfy dynamic compute loads?
A: The peak-rate power supply does not change, and payment is based in part on the peak power rate. Hence, one must work with to time-shift compute workloads. Another important aspect is the SLA (service level agreement) with the customer, which determines the usage of the data center.

Q: Does an AI data center need a different kind of cybersecurity solutions from other, traditional data centers? Does it need more power?
A: Privacy might be a factor in AI data centers. Especially if one deals with HIPAA data, financial data, and other sensitive data over which AI analysis is performed.

Q: What are the biggest concerns about the data-center industry moving into the future?
A: For companies that acquire data centers, the biggest concern is general availability of data centers. Another issue is power rights. We are starting to see the big hyperscalars acquiring power rights like Microsoft's Three Mile Island deal. The data-center providers will ask questions such as: Do I start to acquire power rights? Do I buy options on power?
A: For companies who build data centers, the biggest concern is scale. How big can we go? Servers are getting so dense that the question comes up: how do we get the power needed into that box?
A: Software development is a power cost. New software developers have never touched hardware before. They do not think in terms of power efficiency or power cost. We need software engineers and architects who understand efficiency of hardware and how algorithms impact power.

4.CALL FOR ACTION!

The keynote speakers, short-talk speakers, and panelists suggested various actions to bring us closer to solving the compute-energy nexus challenges.

» Accelerate the construction of power plants.

» Consider the integration of renewable power-generation modalities.

» Consider innovative energy storage and transmission technologies.

- Drastically improve circuit efficiency.
- Develop new computing models and schema for AI, including the training of new models without having to retrain them from scratch every time.
- Limit energy consumption of distributed computing.
- Develop multiple cross-layer specialization and adaptation techniques.
- Improve evolvability and programmability of hardware and software components.
- Foster collaboration between the power grid/energy industry and the IT industry.
- Foster true team co-design with energy and information as peers. Suitable partnerships must be established.
- Remove heat and enhance cooling since the top bottleneck to compute is cooling.

Chips can go faster, and compute can be better. But the ability to remove heat is what is limiting the technology. There are two ways to do this: to make the compute more efficient (i.e., reduce the J of energy required for a compute function), or to enhance cooling and heat rejection. For example, the B200 GPU can run more than 3 kW of IT load. It's limited to 1kW simply because those are the liquid cooling limitations.

References:
[1] Goldman-Sachs, AI is poised to drive 160% increase in data center power demand, May 14, 2024

[2] Dara Kerr, AI brings soaring emissions for Google and Microsoft, a major contributor to climate change, National Public Radio (NPR), July 12, 2024

[3] Sean Patrick Cooper, Noisy, hungry data centers are catching communities by surprise, The New York Times, September 15, 2024

[4] IDC, Cloud to edge datacenter trends, 2024

[5] https://time.com/6987773/ai-data-centers-energy-usage-climate-change/

[6] Jinghan Sun et al, " Exploring the Efficiency of Renewable Energy-based Modular Data Centers at Scale | Proceedings of the 2024 ACM Symposium on Cloud Computing", ACM SoCC 2024.

Appendix A: Organization of the Workshop
Chairs: Klara Nahrstedt and Philip Krein
Program Committee: Deming Chen, Jian Huang, Rakesh Kumar, Elyse Rosenbaum, José Schutt-Ainé, Andrew Stillwell
Local Organization Committee: Baub Alred, Jenny Applequist, Karen Cromwell, Clint Harper, Erica Kennedy, Normand Paquin

Participating Departments, School, and Interdisciplinary Research Units (IRUs):
» Bioengineering Department
» Coordinated Science Laboratory
» Electrical and Computer Engineering Department
» Mechanical Science and Engineering Department
» National Center for Supercomputing Applications (NCSA)
» Nuclear, Plasma and Radiological Engineering Department
» Siebel School of Computing and Data Science

Participating Industry:
» NVIDIA
» Ocient
» Burns & McDonnell

Appendix B: Workshop Speakers' Abstracts:

Morning Talks' Speakers:
» Keynote: Wen-mei Hwu – Compute-Energy Nexus Workshop
» Rakesh Kumar – Compute-Energy Nexus Workshop
» Josep Torrellas – Compute-Energy Nexus Workshop
» Sarita Adve – Compute-Energy Nexus Workshop
» Nenad Miljkovic – Compute-Energy Nexus Workshop

Afternoon Talks' Speakers:
» Keynote: Chris Gladwin – Compute-Energy Nexus Workshop
» Andrew Stillwell – Compute-Energy Nexus Workshop
» Philip T Krein – Compute-Energy Nexus Workshop
» Jian Huang – Compute-Energy Nexus Workshop
» Pavan Hanumolu – Compute-Energy Nexus Workshop

Appendix C: Speakers' and Panelists' Biographies

Speakers:

Sarita Adve is the Richard T. Cheng Professor of Computer Science at the University of Illinois Urbana-Champaign. Her research interests span the system stack, ranging from hardware to applications. She directs IMMERSE, the Center for Immersive Computing, which brings together expertise in immersive technologies, applications, and human experience to perform research, educate a new workforce, and build infrastructure to enable a new era of immersive computing. Her group released ILLIXR (Illinois Extended Reality testbed), an open-source end-to-end extended reality system and research testbed, to democratize XR research, development, and benchmarking. Her work on data-race-free, C++, and Java memory consistency models forms the foundation for memory models used in most hardware and software systems today. She is also known for her work on heterogeneous architectures and software-driven hardware resiliency. She is a member of the American Academy of Arts and Sciences, a fellow of the ACM, IEEE, and AAAS, and a recipient of the ACM/IEEE-CS Ken Kennedy award, the ACM SIGARCH Maurice Wilkes award, and the IIT-Bombay Distinguished Alumnus/Alumna award. As ACM SIGARCH chair, she co-founded the CARES movement, winner of the CRA distinguished service award, to address discrimination and harassment in computer science research events. She received her PhD from the University of Wisconsin–Madison and her B.Tech. from the Indian Institute of Technology, Bombay.

Rashid Bashir is Professor of Bioengineering, the Grainger Distinguished Chair in Engineering, and is currently the 15th Dean of Grainger College of Engineering at the University of Illinois at Urbana-Champaign. He received his B.S. degree in Electrical Engineering from Texas Tech University and his M.S and Ph.D. in Electrical Engineering from Purdue University in 1992. He is an internationally renowned scholar in micro-fluidics and nanotechnology based diagnostic technologies for precision and personalized medicine, and 3D bio-fabrication of multi-cellular engineered living systems for biological soft robotics. Bashir received the 2012 IEEE EMBS Technical Achievement award and the 2018 Pritzker Distinguished Lectureship Award from the Biomedical Engineering Society (BMES). He also received the 2021 Professional Impact Award for Education from the American Institute of Medical and Biological Engineering (AIMBE). He was on the founding team of the Chan Zuckerberg Biohub Chicago, awarded in 2023 and now is a member of the Executive Advisory Committee. He is a fellow of IEEE, BMES, AIMBE, APS, IAMBE, RSC, and AAAS. He was elected to the National Academy of Inventors in 2018, National Academy of Medicine in 2023, and to the American Academy of Arts and Sciences in 2024. He is also academic co-founder of Prenosis, Inc. and VedaBio, Inc.

Chris Gladwin is the CEO and Co-Founder of Ocient, a company providing the software platform for the largest data-analyzing systems in the world. He also is the Chair and Co-Founder of P33, an organization that is helping transform Chicago into a globally top-tier tech region. And Chris is the Chair and Co-Founder of The Forge, which provides amazing recreational and environmental experiences along with ecosystem revitalization. In 2004, Chris founded Cleversafe, which became the largest and most strategic object storage vendor in the world. He raised $100M and led the company to a $1.4B exit in 2015 when IBM acquired the company. The technology Chris invented and Cleversafe commercialized is used by most people in the digital world every day and became one of the ten most powerful patent portfolios in the world. Prior to Cleversafe, Chris was the Founding CEO of MusicNow, which provided the first digital music services in the U.S., and was the founding CEO of Cruise Technologies, the leading provider of wireless thin client computers. Previously, he managed product strategy for Zenith Data Systems and was a database programmer at Lockheed Martin. Chris has an engineering degree from MIT and a PhD from Illinois Tech and is an inventor on over 280 patents.
Pavan Hanumolu is an Electrical and Computer Engineering Professor at the University of Illinois Urbana-Champaign. His research focuses on energy-efficient wireline communication links and integrated power management circuits. He is an IEEE Fellow and a former Editor-in-Chief of the *IEEE Journal of Solid-State Circuits.*

Jian Huang is an Associate Professor and Y. T. Lo Faculty Fellow in the ECE department at the University of Illinois Urbana-Champaign. His research interests include computer systems and architecture, AI infrastructure, memory/storage systems, systems security, and especially the intersections of them. Most recently, he has been working on sustainable AI infrastructures. His research contributions have been published at computer architecture and systems conferences such as ISCA, MICRO, ASPLOS, OSDI, and SOSP. He received the inaugural SIGMICRO Early Career Award, NSF CAREER Award, NSF CRII Award, USENIX Best Paper Award, MICRO Top Picks, MICRO Hall of Fame, Dean's Award for Early Innovation, and a few industry awards. He is also the founder of the workshop on Hot Topics in System Infrastructure (HotInfra).

Wen-mei W. Hwu is a Senior Distinguished Research Scientist and Senior Director of Research at NVIDIA. He is also a Professor Emeritus and the Sanders-AMD Endowed Chair Emeritus of ECE at the University of Illinois Urbana-Champaign after 34 years of service. His research is in the parallel architecture, algorithms, and infrastructure software for data-intensive and computational intelligence applications. He served as the Illinois co-director of the IBM-Illinois Center for Cognitive Computing Systems Research Center (c3sr.com) from 2016 to 2020. He was a PI of the NSF Blue Waters supercomputer project and its associated PAID program for training science teams to use GPUs. For his research contributions, he received the ACM/IEEE Eckert-Mauchly Award, the ACM SigArch Maurice Wilkes Award, the ACM Grace Murray Hopper Award, the IEEE Computer Society Charles Babbage Award, the ISCA Influential Paper Award, the MICRO Test-of-Time Award, the IEEE Computer Society B. R. Rau Award, the CGO Test-of-Time Award, several best paper awards, and the Distinguished Alumni Award in CS of the University of California, Berkeley. He is a Fellow of the IEEE and ACM.

Philip T. Krein holds the Grainger Endowed Chair Emeritus in Electric Machinery and Electromechanics at the University of Illinois Urbana-Champaign. He received the IEEE William E. Newell Power Electronics Award in 2003 and the IEEE Transportation Technologies Award in 2021. He holds 42 U.S. patents and is a Registered Professional Engineer in Illinois and Oregon, a Fellow of the U.S. National Academy of Inventors, and a member of the U.S. National Academy of Engineering.
Rakesh Kumar is an Electrical and Computer Engineering Professor at the University of Illinois Urbana-Champaign. His research focuses on computer architecture and system-level design automation. He is particularly excited these days about the promise and challenges of large-scale integration.

Nenad Miljkovic is the Founder Professor of Mechanical Science and Engineering at the University of Illinois Urbana-Champaign. He has courtesy appointments in Electrical and Computer Engineering and the Materials Research Laboratory. He is the Director of the Air Conditioning and Refrigeration Center (ACRC), which is supported by 21 industrial partners. His group's research intersects the multidisciplinary fields of thermo-fluid science, interfacial phenomena, scalable nanomanufacturing, and renewable energy. He is a recipient of the NSF CAREER Award, the ACS PRF DNI Award, the ONR YIP Award, the ASME ICNMM Young Faculty Award, the ASME Pi Tau Sigma Gold Medal, the CERL R&D Technical Achievement Award, the US Army Corps of Engineers ERDC R&D Achievement Award, the SME Young Faculty Award, the Bergles-Rohsenow Young Investigator Award in Heat Transfer, and the ASME EPPD Early Career Award, and is an ASME Fellow.

Klara Nahrstedt is the Swanlund Endowed Chair and Professor in the Siebel School of Computing and Data Science at the University of Illinois at Urbana-Champaign and the Director of the Coordinated Science Laboratory. Her research interests are in trustworthy multimedia distributed systems, quality of service (QoS) and resource management, trustworthy power grids, and advanced edge-cloud-based cyber-infrastructures. She is the recipient of the IEEE Computer Society Technical Achievement Award, the ACM Special Interest Group on Multimedia Technical Achievement Award, and others. Nahrstedt received her Diploma in mathematics from Humboldt University, Berlin, Germany, in 1985. In 1995, she received her Ph.D. from the Department of Computer and Information Science at the University of Pennsylvania. She is ACM, IEEE and AAS Fellow, member of the Leopoldina German National Academy of Sciences, and member of the USA National Academy of Engineering.

Andrew Stillwell received dual B.S. degrees in Electrical Engineering and Computer Engineering from the University of Missouri, Columbia, MO, in 2005. He spent seven years at National Instruments before arriving at the University of Illinois Urbana-Champaign, where he received M.S. and Ph.D. degrees in electrical and computer engineering in 2015 and 2019, respectively. He is currently an Assistant Professor in the Electrical and Computer Engineering Department with the U. of I. His research interests include renewable energy applications, hybrid switched-capacitor converters, power electronics design optimization, high-density and high-efficiency power converters, and advanced control techniques for multilevel converters.

Jinghan Sun is a final-year CS Ph.D. student at the University of Illinois at Urbana-Champaign, advised by Prof. Marc Snir and Prof. Jian Huang. His research interests are at the intersection of machine learning (ML) and systems, with a particular focus on building efficient and reliable storage systems using machine learning. He works across the entire storage stack, from storage hardware development (e.g., solid-state drives) to system software and cloud storage platforms.

Josep Torrellas is the Saburo Muroga Professor of Computer Science at the University of Illinois Urbana-Champaign. He is the Director of the SRC/DARPA ACE Center for Evolvable Computing, past Co-Leader of an Intel Strategic Research Alliance (ISRA) on Computer Security, and past Director of the Illinois-Intel Parallelism Center (I2PC). His research interests are computer architectures for shared-memory multiprocessors and parallel computing. Some of his contributions include thread-level speculation (TLS) architectures, the Bulk Multiprocessor concept, deterministic record and replay mechanisms, process variation mitigation techniques, and hardware defenses against speculative execution attacks. In addition, he has contributed to several experimental multiprocessor designs such as IBM's PERCS Multiprocessor, Intel's Runnemede Extreme-Scale Multiprocessor, Illinois Cedar, and Stanford DASH. He has graduated 49 Ph.D.s, who are now leaders in academia and industry.

Panelists Biographies:
Brett Bode, National Center for Supercomputing Applications, University of Illinois Urbana-Champaign. Manager of the system and tools software development for the Blue Waters project. In 2012, he took on the role of division director for the newly formed Advanced Digital Services Directorate.

Alejandro Domínguez-García, Electrical & Computer Engineering, University of Illinois Urbana-Champaign. M. Stanley Helm Professor.

Katy Huff, Nuclear, Plasma, & Radiological Engineering, University of Illinois Urbana-Champaign. Was Assistant Secretary for the Office of Nuclear Energy from 2022 to 2024. Is an associate professor at the University of Illinois Urbana-Champaign.

Wen-mei Hwu, NVIDIA. Senior Distinguished Research Scientist and Senior Director of Research at NVIDIA; professor emeritus and the Sanders-AMD Endowed Chair Emeritus of ECE at the University of Illinois Urbana-Champaign after 34 years of service.

Joe Jablonski, Ocient. Co-founder and Chief Product Officer.

David Nicol, Electrical & Computer Engineering, University of Illinois Urbana-Champaign. Herman M. Dieckamp Endowed Chair in Engineering and Director of the Information Trust Institute.

William Schaumann, Burns & McDonnell. Senior Associate Electrical Engineer.

Naresh Shanbhag, Electrical & Computer Engineering, University of Illinois Urbana-Champaign. Jack Kilby Professor. Was the director of the Systems on Nanoscale Information fabriCs (SONIC) Center.

Appendix D: Workshop Agenda
Workshop Website Link: Agenda – Compute-Energy Nexus Workshop

8:00–8:15 a.m., Rashid Bashir (Dean of Grainger College of Engineering at U. of I.)
"Welcome and Importance of the Workshop"

8:15–8:30 a.m., Klara Nahrstedt (Director of CSL at U. of I.)
"Setting Goals of the Workshop"

8:30–9:30 a.m., Keynote
Session Chair: Klara Nahrstedt

Wen-mei Hwu (NVIDIA)
Title: "Power and Energy Efficiency in Hyper-Scale AI Data Centers"

9:30–9:45 a.m., Break

9:45–10:45 a.m., Short Presentations
Session Chair: Roy Campbell

Rakesh Kumar (ECE at U. of I.)
Title: "Addressing Data Center energy concerns through Large-scale integration: Opportunities and Challenges"

Josep Torrellas (SCDS at U. of I.)
Title: "Addressing Energy Inefficiencies in Datacenter Infrastructure"

Sarita Adve (SCDS at U. of I.)
Title: "The Energy Challenge of Immersive Computing"

Nenad Miljkovic (MechSE at U. of I.)
Title: "Holistic Rack-to-Processor Power and Thermal Co-Design for Future Servers and Data Centers"

10:45 a.m.–12:15 p.m., Industry & Academia Panel

Moderator: Klara Nahrstedt
Wen-mei Hwu (NVIDIA)
Kathy Huff (Nuclear Engineering at U. of I.)
Brett Bode (NCSA at U. of I.)
Naresh Shanbhag (ECE at U. of I.)

12:15–1:30 p.m., Lunch & Posters

1:30–2:30 p.m., Keynote
Session Chair: Philip Krein

Chris Gladwin (Ocient)
“Growing Data Center Energy: A Global Threat and Opportunity”

2:30–3:30 p.m., Short Presentations
Session Chair: José Schutt-Ainé

Andrew Stillwell (ECE at U. of I.)
Title: “Point-of-Load Converters: Status and Future”

Philip Krein (ECE at U. of I.)
Title: “Teaming for the Computing/Energy Nexus”

Jian Huang (ECE at U. of I.)
Title: “Exploring the Efficiency of Renewable Energy-based Modular Data Centers at Scale”

Pavan Hanumolu (ECE at U. of I.)
Title: Solving Connectivity Bottlenecks for Efficient Data Centers

3:30–3:45 p.m., Break

3:45–5:00 p.m., Industry & Academia Panel
Moderator: Philip Krein

William Schaumann (Burns & McDonnell)
Joe Jablonski (Ocient)
Alejandro Domínguez-García (ECE at U. of I.)
David Nicol (ECE at U. of I.)

Appendix E: Posters Presented

- “Current-sourced Hybrid Switched-Capacitor Converter for Data Center Power Delivery,” Aria K. Delmar, William Vavrik, and Andrew Stillwell (University of Illinois Urbana-Champaign).
- “EcoCell: Energy-aware Traffic Shaping for Cellular Radio Access Networks,” Seoyul Oh, Zikun Liu, Bill Tao, and Deepak Vasisht (University of Illinois Urbana-Champaign).
- “Eliminating Memory Overheads with Data-Centric Computing,” Minh S. Q. Truong, Yiqiu Sun, Arjun Tyagi, Ryan Wong, Sudhanshu Agarwal, Yilin Shen, Dawei Xiong, Pratik R. Sampat, Akash Pardeshi, Shreya Sharma, and Saugata Ghose (University of Illinois Urbana-Champaign).
- “Energy-efficient Machine Learning Systems,” Jiaxi Li and Klara Nahrstedt (University of Illinois Urbana-Champaign).
- “Formal Verification for Smart Grid Security Challenges,” Mckenzy Heavlin, Klara Nahrstedt, and David Nicol (University of Illinois Urbana-Champaign).
- “H-GREEN: Hierarchical GNN-RL-based Energy-efficient Resource Management in Multi-cloud,” Jinghua Wang, Asser Tantawi, Olivier Tardieu, Alaa S. Youssef, Tamar Eilam, Chen Wang, Eun Kyung Lee, Pradip Bose, Klara Nahrstedt, and Deming Chen (University of Illinois Urbana-Champaign).
- “Power Electronics and Computing as Peers in Future Data Centers and Supercomputers,” Pradeep S. Shenoy (Texas Instruments) and Philip T. Krein (University of Illinois Urbana-Champaign).
- “Reconductoring: Facilitating Grid Evolution,” Leslie M. Gioja (University of Illinois Urbana-Champaign).
- “A Resilient Power Grid Leveraging Internet Infrastructure,” Ashish Kashinath, Radhika Niranjan Mysore, and Sibin Mohan (University of Illinois Urbana-Champaign)