Workshop Report:
2021 International Symposium of Quantitative Codesign of Supercomputers

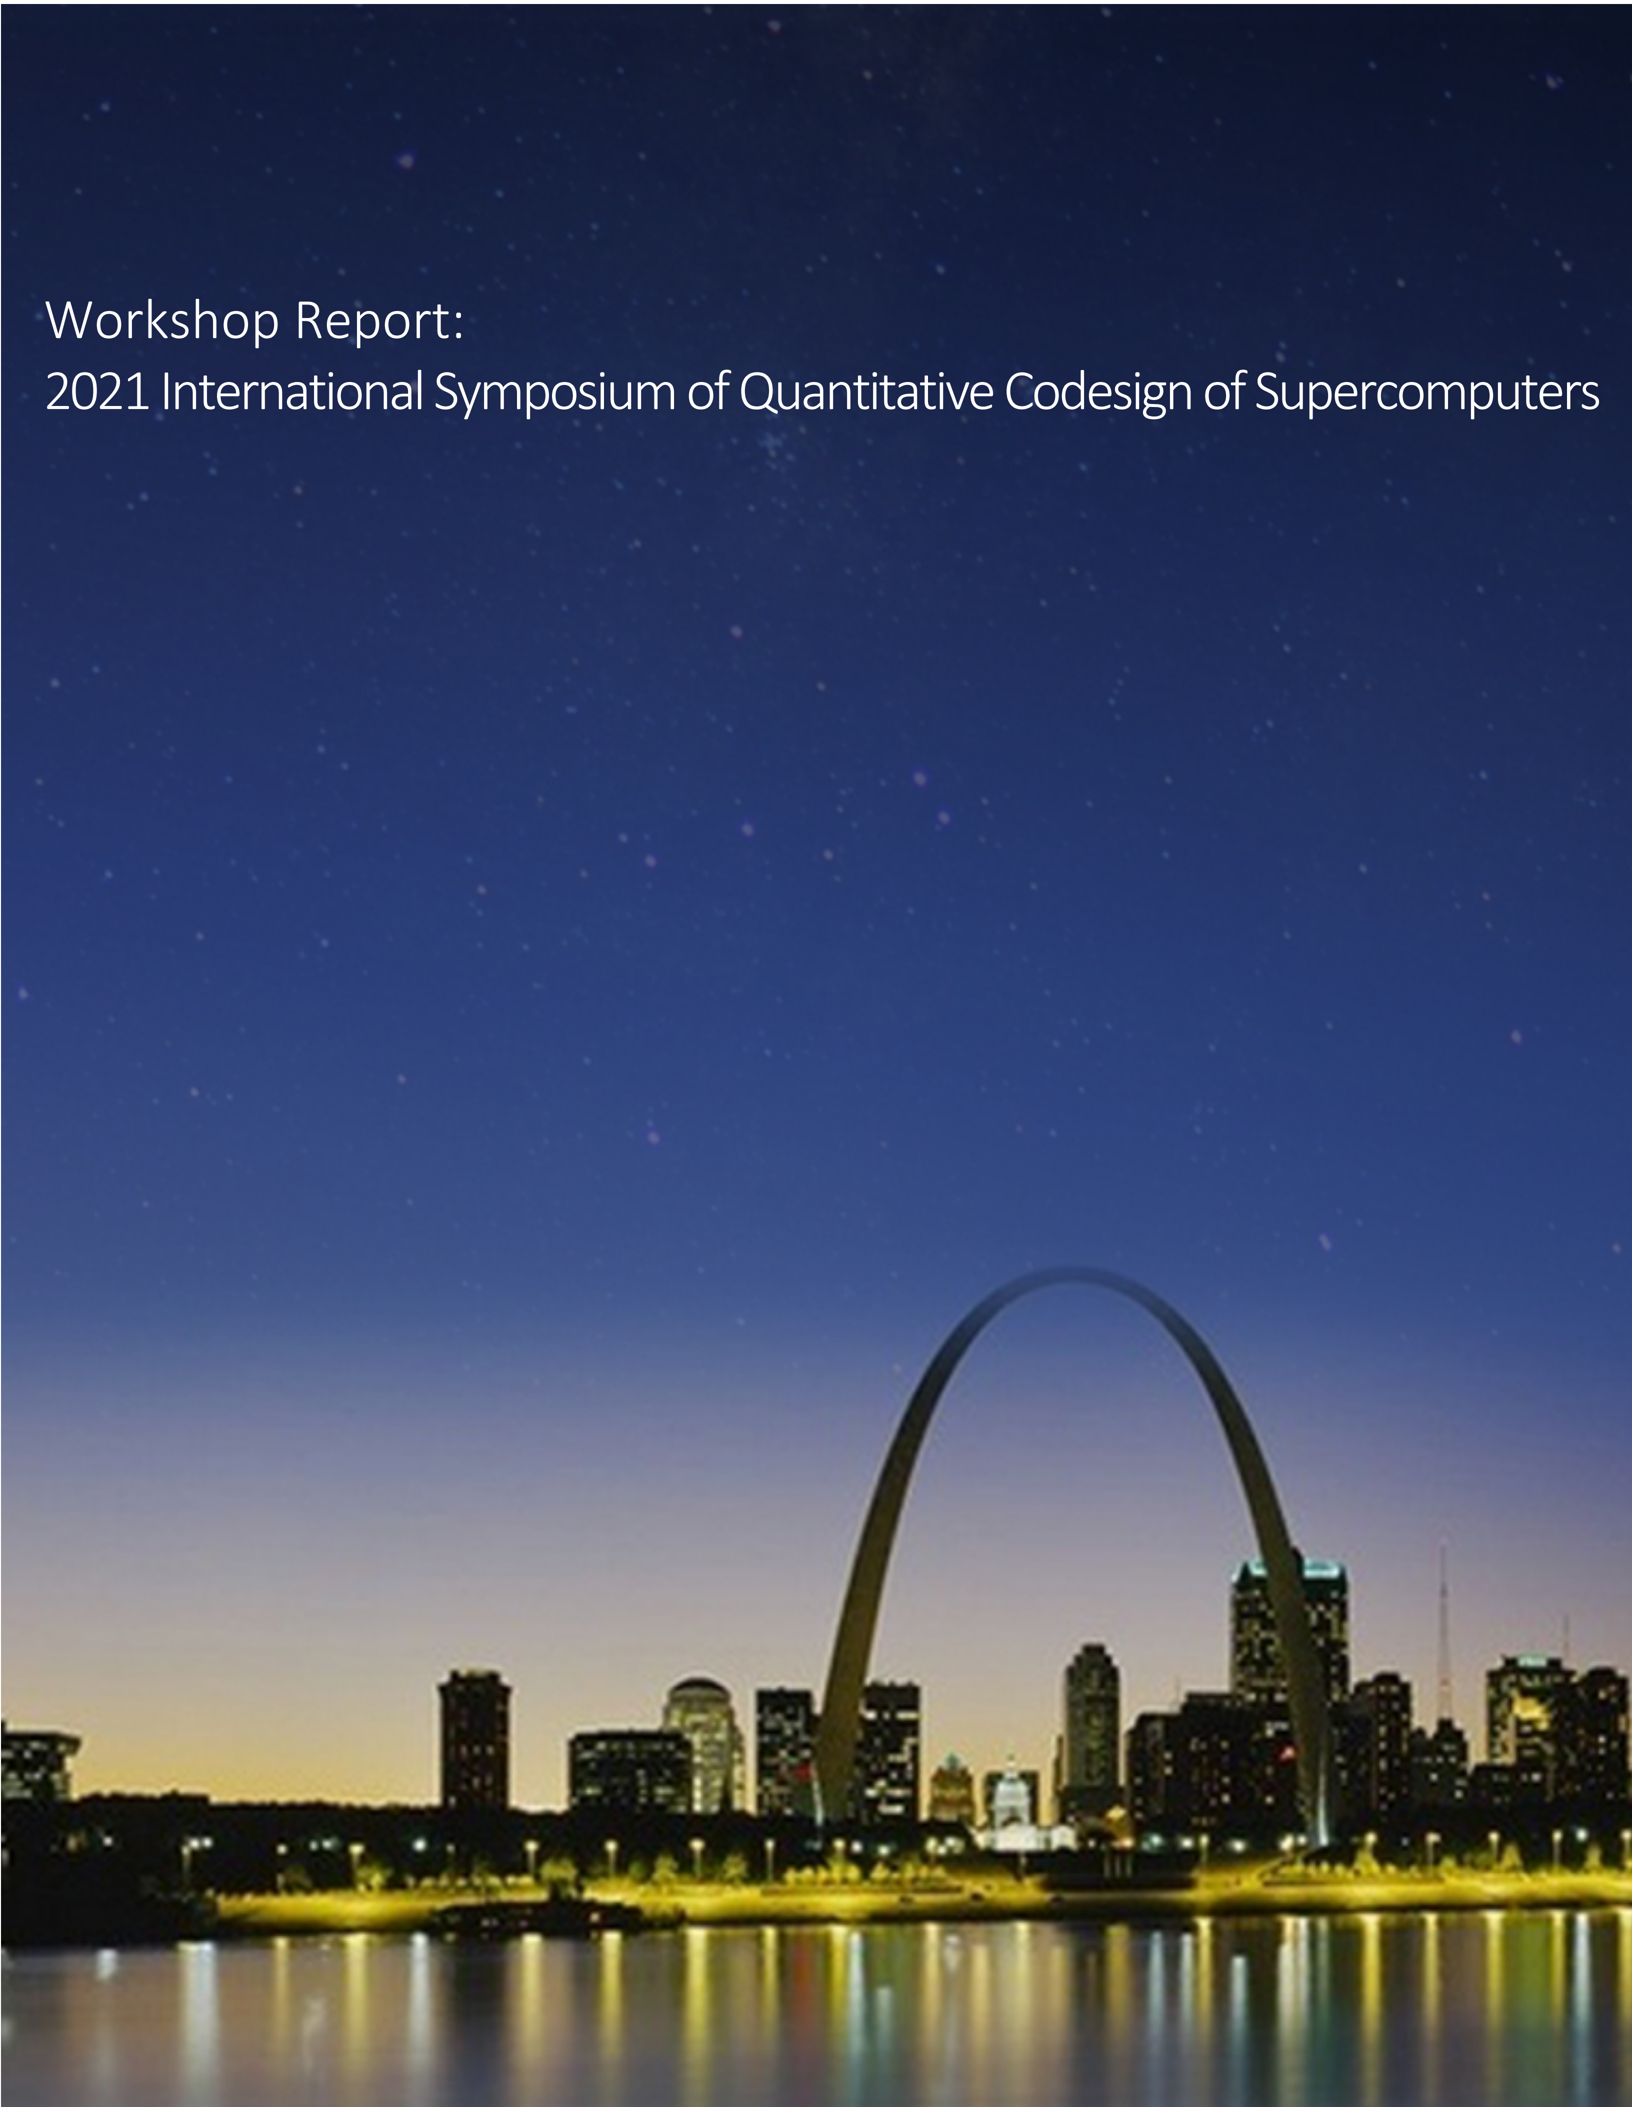

Acknowledgements

We would like to acknowledge the efforts of the following people who not only made our symposium possible, but worked to make our symposium outstanding: SC'21 workshop committee chair Anshu Dubey, vice chair Mohamed Wahib, and proceedings chair Sivasankaran Rajamanickam; the anonymous workshop proposal reviewers who provided helpful guidance; SC'21 planning committee chair Lori Diachin; virtual logistics operations subcommittee Verónica G. Melesse Vergara, our web designer Laura Chavez-Becker; and the audio-visual team that provided great support for our symposium.

Thank You!

2021 International Symposium on the Quantitative Design of Supercomputers
Held in conjunction with Supercomputing '21
St Louis, MO – November 19, 2021

1. Background to Workshop

The Quantitative Codesign of Supercomputers symposium is an annual workshop series that aims to significantly improve the effectiveness of high-performance computing through bringing about increased understanding of current limitations and improved development processes.

This symposium considers combining two methodologies—collaborative codesign and data-driven analysis—to realize the full potential of supercomputing. For full potential of supercomputing we consider everything pertaining to output production, including but not limited to the performance of applications, system software, workflows, health of hardware. Our centers store vast sums of information, yet using this data presents demanding challenges. Much of the quantitative challenge has to do with discovering, accessing, and analyzing the right data. Codesign also presents formidable challenges. For example, how can a codesign development use the data collected on current systems to facilitate the design of next-generation supercomputers and successfully support our upcoming environments. Quantitative codesign offers a collaborative evidence-based approach to address our existing needs and our upcoming ambitions. This symposium was created to bring together leaders in the field to review current efforts across centers and discuss areas that show potential. For further information on the topic of Quantitative Codesign, please refer to Appendix 1 – Workshop Scope, Context & Related Activities.

2. Purpose of the Workshop

The purpose of the workshop was to build the necessary community support to bring about improved and increased numbers of instances of quantitative codesign. As architectural options expand in type and complexity, the need for quantitative bases to drive architectural directions becomes increasingly urgent. We do not have the primary mission to raise awareness of an individual's research; rather we wish to bring more wide-ranging interactions highlighting vision and positions and stimulating discussions.

3. Workshop Process and Timeline

The Quantitative Codesign of Supercomputers symposium took place during the fifth day of the 2021 Supercomputing conference. Due to COVID conditions at the time, the workshop was held as a hybrid model with both in-person and virtual attendees and speakers. The workshop was framed in the Symposium format to achieve the kind of deep interactions that lead to change within HPC. Our preference for audience interaction was in response to the state of the field (which we see as in its infancy). Given our desire to bring more wide-ranging interactions highlighting vision and positions and stimulating discussions, we developed a schedule designed to facilitate these interactions (see Table 1). In particular:

- The keynote speaker was chosen based on his long history in HPC with work that spans all areas of codesign including novel architectures, system and application software, tool development, performance diagnostics and more, in both lab and academic environments.
- Three speakers were chosen who, as an aggregate, provided codesign perspectives of proxy-applications, instrumentation, and hardware design.
- A panel with diverse expertise in codesign, HPC system software and middleware research, center wide monitoring and operational aspects, bringing HPC products to market, and application / library development presented perspectives in How Better Data Can Result In Better Codesign.
- A moderated discussion of audience, speakers, and panelists was included to enable both technical discussions and community-building.

Table 1 – Symposium Agenda

Duration	Speaker/Panelist		Abstract
5 mins		Terry Jones	Welcome
30 mins	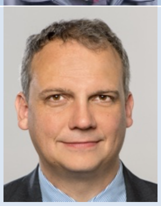	Matin Schulz	On the Experience of Long Time Data Collection on SuperMUC to Drive Energy Efficiency HPC systems are starting to push the limits in energy and power consumption. Therefore, increasing energy efficiency is one of the most important design goals of HPC architectures. First, we need to understand the energy and power consumption in architectures and their applications. We implemented a comprehensive system. I will discuss our infrastructure, the data collected, the insights we were able to gain, and how we were able to use the data to affect machine design and operation.
20 mins		Duncan Roweth	HPE Network Codesign of Slingshot Interconnect This talk will describe network codesign activities undertaken as part of the Cray PathForward project: characterization of ECP applications and workloads, development of Slingshot device models, enhancements to the SST simulator, and system scale simulation. We will also discuss a novel hybrid emulation/simulation environment developed as part of this work. We will present results of simulation together with results measured on HPE Cray EX systems being shipped.
20 mins		Mitsuhisa Sato	Bright Side and Dark Side of Codesign: Lessons Learned from Codesign for Fugaku We have carried out the project to develop the Japanese exascale supercomputer "Fugaku." In the project, we designed and developed an original manycore processor, as well as a system including interconnect and a storage subsystem. The codesign of the system and applications was a key to making it power efficient and high performance. We will present our codesign effort and discuss the "bright" and "dark" sides of the codesign methodology as lessons learned from our project.
30 mins	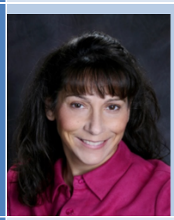	Jeanine Cook	What's Needed for Effective Proxies for Codesign Proxy applications have become a prevalent co-design tool for next-generation systems and for system and testbed procurements. They must accurately represent the behavior of their parent among other qualities to ensure community adoption. In this talk, we will present quantitative methods to determine proxy fidelity with respect to parent and we will discuss practices to improve the usability of proxy apps based on lessons learned from the ECP proxy app project.
40 mins	 stance-statemen "How Better Data Can Result In Better Codesign"	Jim Brandt (Sandia) Moderator	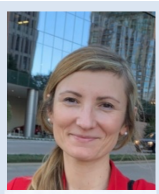   Florina Ciorba Univ of Basel, (TOP LEFT) Esteban Meneses CRNHC, (TOP MIDDLE) Larry Kaplan HPE, (TOP RIGHT)    Jesus Labarta SBC, (BOTTOM LEFT) Hatem Ltaief, KAUST, (BOTTOM MID) Gabe Rockefeller, LANL, (BOTTOM RIGHT)
40 mins		Jim Brandt (Sandia) Moderator	In-person and Virtual Audience participation "Technical challenges, future opportunities, community building, etc."
5 mins	Terry Jones (ORNL)		Closing Remarks

Panelists addressed the stance-statement "How Better Data Can Result In Better Codesign" followed by a moderated discussion with both virtual and in-person attendees. The moderator then moved attention more fully toward the audience in the final session giving additional opportunities to introduce new aspects or to return to an earlier topic. This allowed for maximum interaction among group members supported by the facilitating moderator. The symposium concluding with a few final remarks by the organizing chair and an encouragement for attendees to follow-up with email and/or downloading the workshop report from the symposium website.

The COVID pandemic had an impact on the format and character of the workshop. This was the first time for the SC series of conferences to ever have a hybrid format: SC21 supported both in person attendees at the America's Center in St. Louis Missouri and remote attendees though the SC21 HUBB platform, Zoom and Sli.do. The role of the session chair and organizer remained largely the same as in previous years with some adjustments and increased responsibilities to account for remote participation by speakers and attendees. SC21 produced nearly 400 pages in HUBB spanning all program areas. Workshops were able to request several formats. The Quantitative Codesign of Supercomputers symposium chose the ***Live stream sessions*** option. Under this format, content recorded by AV technicians at the convention center and sent to remote participants in real time via Vimeo; remote presenters supported via zoom (see Figure 1). For all remote symposium presenters, we arranged for an internet assessment on the day of the symposium prior to the symposium start. This was used to ensure no fall-back measures were needed. All remote speakers were able to participate as planned.

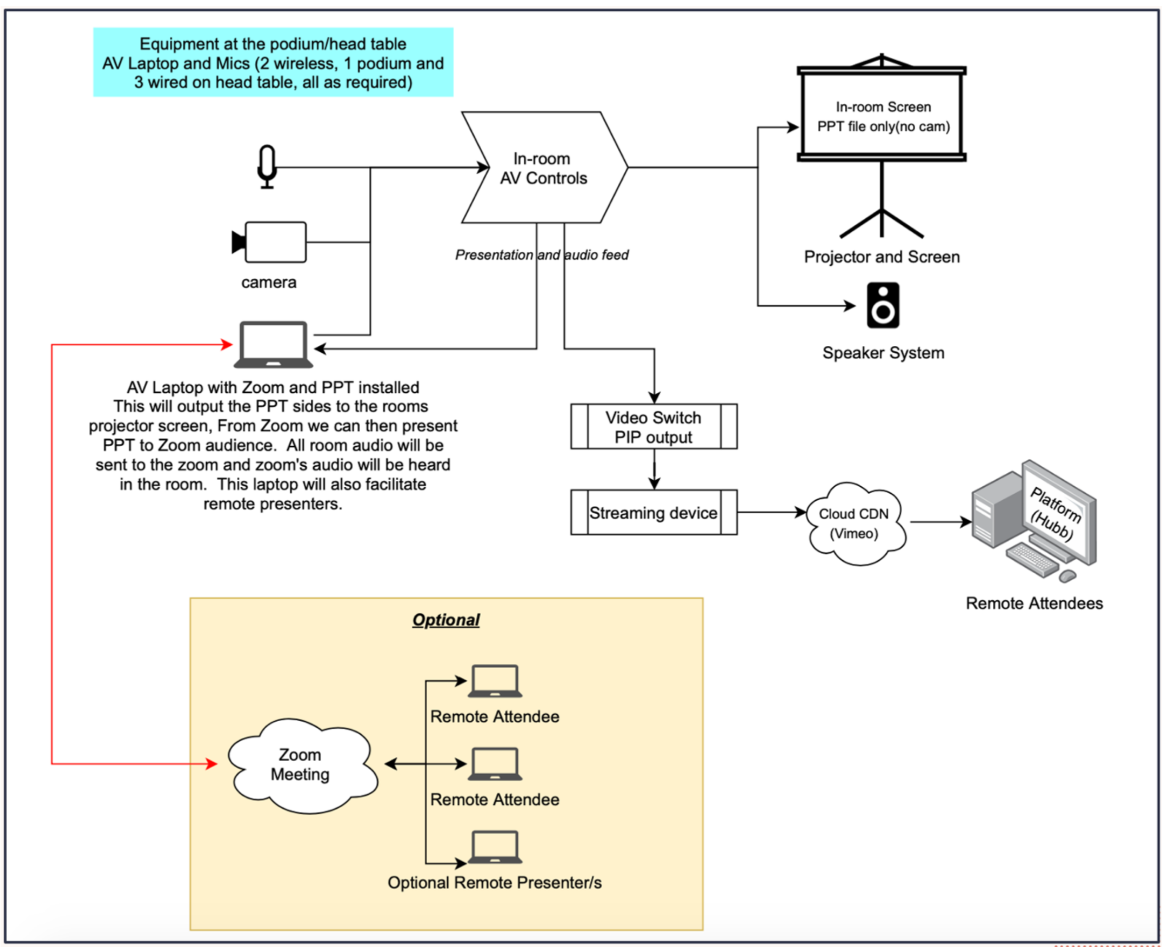

Figure 1 – Logistical Setup of Live Stream Format Used By Symposium. [source: SC'21 Session Chair & Organizers Training presentation]

4. Hybrid Process Experience

The Symposium's program committee considered SC'21's Live Stream hybrid format a success. Prior to the meeting, we wondered how smoothly the Hubb virtual interface and the coordination of in-person presentations working in concert with virtual zoom interface would proceed given Hubb was somewhat unfamiliar to some participants. Our experience was positive. Indeed, many had already experienced the technologies earlier in the week, so our Friday timeslot was a benefit in this regard and we had no significant issues.

5. Workshop Outputs

The following positive results have been achieved with from the workshop:

- A large group of high performance computing professionals came together to pursue community building
- Monitoring journals (outcome and strategy) were discussed and templates provided to guide the process of data collection and the use of these data
- Videos of the invited talks and panels were recorded by SC's Live Stream AV team
- Discussion on Vision and Possibilities of Quantitative Codesign of Supercomputers were discussed, and ideas for future work were identified
- A workshop report was written to document the results

Table 2 provides a list of the talking points from the symposium's concluding audience-wide discussion.

Table 2 – Captured Discussion Points from Moderated Discussion

	Discussion Points
1	What is the hardest problem to overcome to achieve Quantitative Codesign as presented here: • Analysis Paralysis -- progress is stymied while we continually wait for the better answer available with additional data • Too much data is burdened with OUO, proprietary or sharing concerns. There is no good way to anonymize data from supercomputer centers or vendors. (This was a common point, as it is a blocker to analysis, even within a site (e.g., user vs admin accesss). • Lack of Ownership and Leadership. To work, codesign requires many communities coming together which will never happen because there is no single legislating oversight entity over all the needed communities. Note that a single key missing player can cause an inordinate amount of challenges. • Making data FAIR (Findable, Accessible, Interoperable, Reusable) • Overhead of collecting and measuring data. Overcoming the perception that applications will be impacted too much. • Budget. There is no foreseeable way to fund the time and resources required to save the amount of data envisioned by Quantitative Codesign.
2	What makes this the worst problem from the above? • As technocrats, we're comfortable with tackling the technical problems associated with FAIR principles organizing a community with leadership & direction, and garnering support for funding, But overcoming problems with data release (including declassification of helpful information about classified environments) is a hard challenge. • Good discussion on this point.
3	How would you attack this worst problem? • If the worst problem is the release of data from all parties involved (including potentially anonymization of data from some applications running on our supercomputers), developing a plan of attack would require a large codesign-type discussion and subsequent plan of action for those stakeholders. [Note that Quantitative Codesign can accomplish much even with this aspect progressing more slowly than other aspects that are addressed via technical and process-oriented means.]
4	Will Quantitative Codesign of Supercomputers ever be realized? • Those attending SC'21 were upbeat about the need and prospects of Quantitative Codesign for Supercomputers, given new interest in Reimagining CoDesign (ASCR workshop) and new analysis and instrumentation possibilities. However, clearly we are a long way from the vision.

In addition, monitoring journals (outcome and strategy) were discussed and templates provided to guide the process of data collection and the use of this data.

6. Post-Workshop Recommendations and "Next step" Strategies

There are a number of recommended "next steps" that should be followed to increase the usability of quantitative codesign of supercomputers.

a. **Continue the website (Link)**
   Provide ongoing support to Quantitative Codesign of Supercomputers website. This web presence becomes an anchor for announcements and a source to discover resources are pertinent email addresses.
b. **Disseminate the Workshop Report**
   Providing this post-workshop report of the event will be an important resource for the symposium's community building objective.
c. **Organizing & Program Committee Members are to watch for mission furthering opportunities.**
   This follow-up activity is to ensure that a wide segment of high performance computing is monitored for events, interactions and publications for opportunities to advance high performance computing through quantitative codesign concepts.
d. **Advance the Quantitative Codesign agenda with a 2022 Symposium**
   Finally, we are encouraged to repeat the workshop in 2022. This second workshop should consider emphasizing how quantitative codesign can assist in system software development. We would include middleware including runtime systems in our scope. In post discussions of the organizing and program committee, there was a consensus to propose the 360 minute format at SC'22, which would allow additional engagement of the community through position papers and extension of the discussion period, while stipulating that we could fall-back to the 180 minute format if required.

Appendix 1 – Workshop Scope, Context & Related Activities

Scope

This symposium will explore strategies, techniques, and frameworks that combine codesign methodologies with data-driven analysis to realize the full potential of supercomputing. For full potential of supercomputing we consider everything pertaining to output production, including, but not limited to, runtime data of applications, system software and workflows as well as configuration and health of hardware. While modern High Performance Computing (HPC) computing centers track and store an immense amount of data regarding their own operation, using this information to improve the design process still presents significant challenges. Finding the right set of data to collect, filter, and analyze is itself a difficult task, and efforts to address this problem are often fragmented and not applicable across different groups and centers. Even with an appropriate data collection approach, architects and engineers still lack common tools and strategies to aggregate, understand, and incorporate the available information during the design process. As a result, solutions that rely on data collection and analysis are often based on intuition, rather than well-founded theory. Without a clear and standard set of best practices in place, it's no surprise that HPC centers struggle to realize the potential of data-driven codesign across the wide-ranging ecosystem.

Any shortcoming here impacts the whole spectrum of stakeholders. Whether providing hardware architectures, system software, application programming environments, or production run-time environments, having the appropriate knowledge to optimize the interaction and configuration of all of these critical components as well as the evolution of the HPC ecosystem is critical to continued growth. The rapidly changing HPC landscape demands a codesign that effectively uses the data collected on previous and current systems to facilitate the design of next-generation supercomputers and successfully support our upcoming environments.

Context

- **Challenges** We have important issues to resolve, but we're not starting from scratch. HPC computing centers already collect a wealth of information on the health, usage, and efficiency of our machines, workflows and programming environments. While collection and analysis of this information has evolved and improved over the years, there are still severe gaps that have left us unable to provide the knowledge that is needed by hardware and software vendors, system operations staff, application developers, and user groups to create and operate highly efficient and secure large scale HPC systems. Would-be users of this information face significant challenges in obtaining effective analysis in a timely manner and efforts to provide it are currently fragmented across centers both at national and international levels. The infrastructure to collect, store, share and analyze the volumes of available information is a core capability—yet, many challenging barriers remain due in large part to the many stakeholders and insufficient coordination, but also due to data privacy and security issues. With many new potential information sources in future systems, we must quickly identify and address critical requirements and gaps across the various stakeholders. Doing so will enable us to create collective and collaborative solutions that address both existing challenges and emerging needs and effectively support our upcoming HPC environments. The nature of this challenge suggests that it is an excellent opportunity for a codesign approach: codesign is defined as the process of jointly designing interoperating components of a computer system—in particular: applications, algorithms, programming models, system software, as well as the hardware on which they run and the facilities they run in. Design solutions that rely on intelligence derived from the data collection and analysis processes described above are henceforth referred to as Quantitative Codesign of Supercomputers.

  Opportunities abound. Making progress at the highest end of HPC without access to the needed data can be compared to being asked to fly an airplane at night without sufficient instrumentation. Vendors are provided with example applications to target, but often lack a true understanding of where inefficiencies manifest on full scale workloads. Furthermore, architectural simulators do not incorporate the critical performance-killing attributes of current generation technologies and their integration. Hence the vendors miss opportunities for improvement. Moreover, users often only have feedback on operating efficiency at the granularity of total application execution time. Low-level interactions frequently cause substantial performance degradations that users are unable to explain. Likewise, operations staff often lack knowledge of application resource utilization and cannot diagnose the longer run times experienced by the users. In addition, operations staff cannot ensure secure operations without an understanding of normal (expected) behavior and anomalies that deviate from that. Since root causes go undiagnosed on current systems, next generation systems will also fail to address the very same problems.

- Opportunity: First and foremost, we wish to discuss the merits of a coordinated effort to bring together the helpful data from each stakeholder in the codesign space into a framework where data discovery and access is straightforward regardless of data source while respecting data privacy and security concerns. The envisioned Quantitative Codesign environment would pull together data traditionally held by disjointed communities (e.g., sysadmins, application teams, vendors, and so on) into a framework where the needed data is easily accessible. This framework would provide flexible but secure mechanisms for data providers who wish to share their data with others including application teams, vendors, facilities, operations, and system software researchers. In many cases, we seek to bring together data that is currently being produced although not generally known or utilized for a variety of reasons; in a few instances, we seek to extend and provide new data collection capabilities.

  For example, one area that is ripe for integration with Quantitative Codesign processes is the intersection of application development and run-time environments. In the past few years Continuous Integration (CI) has been widely adopted by development teams to continuously test development efforts. As part of these CI efforts, developers test across a variety of platforms on a daily basis and typically provide a pass/fail result for each. Introducing targeted run-time data collection (e.g., memory, application & hardware counters, MPI, OpenMP, GPGPU, I/O, energy consumption) and quantitative analysis into this process would enable feedback to users and identify issues within applications, compiler capabilities, runtimes, and differences across platform architectures that ultimately would drive improvements across the spectrum of stakeholders.

  Integrating Quantitative Codesign capabilities with existing design processes will enable more effective solutions across the computing stack. Information derived from monitoring and analysis would provide valuable insight for users, application developers, system architects, and facility designers as to how, and why, applications make use of the underlying system resources. Furthermore, by identifying the appropriate stakeholders and introducing them to information originating from diverse collection regimes, this symposium seeks to facilitate the discovery and sharing of potentially useful intelligence among larger teams and communities. In doing so, this approach also has the potential to spark further discussions and research on how to collect, employ and share this information more effectively. Thus, there is significant opportunity for discoveries that will not only increase application performance, but also benefit the broader HPC and scientific communities.

- Timeliness: Over the past decade, there has been a growing awareness of the multi-faceted benefits we can derive from data-driven strategies like Quantitative Codesign. This increasing awareness, along with improvements in Machine Learning (ML) technologies, have driven vendors, operations staff, and application developers to espouse integrating an ever-increasing level of instrumentation into their products. The time is ripe for turning this vast trove of available information and the incredible advances in analysis technologies it represents into appropriate knowledge and understanding. Doing so would create a feedback loop that could assist vendors and software developers in their designs. The recent National Strategic Computing Initiative Update Report has recommended that we promote timely access for developers of technologies, architectures, and systems to carry out the research needed to create the future computing software ecosystem, and Quantitative Codesign provides a solution to the ‘access problem’ of these extremely rare machines. If the future envisioned by the CSESSP report is to be realized, our software base will require significant investment in both modified and new code — an activity enormously assisted by Quantitative Codesign. There is no disagreement that more knowledge is good though there is still lack of concurrence across HPC stakeholders as to the cost/benefit tradeoff for varying fidelities of information collection and long term storage. The benefits of Quantitative Codesign will come through integrating design processes with more detailed knowledge of the interactions of the various components within the HPC ecosystem.

Quantitative Codesign is also essential for addressing challenges brought about by the recent trend of increasing heterogeneity and varied accelerators in HPC architectures. For example, many HPC machines now incorporate alternative types of memory alongside conventional DDR SDRAM. Technologies such as "on-package" or "die-stacked" DRAM as well as non-volatile RAMs can provide distinct advantages compared to conventional DRAM, including higher performance as well as cheaper and more energy efficient storage per byte. Each of these technologies also comes with its own limitations, such as smaller capacity or less bandwidth for reads and writes. Further complications arise because some of these new technologies can interface directly with processor caches, while others can only be accessed through peripheral devices, such as GPUs or other accelerators.

Quantitative Codesign could mitigate many of the current problems with allocating and managing such heterogeneous resources effectively. Detailed knowledge of application demands will enable architects to make better decisions about how to select and organize computing hardware. This approach can also help system software, including operating systems, compilers, and runtime software, distribute the available hardware resources among applications more effectively. Codesigned system software could utilize improved knowledge from new data sources for better energy management and workflow management. Integrating high-level profiling and analysis with low-level resource management routines will enable these systems to implement new policies that respond flexibly to changes in application demands and could potentially expose powerful new efficiencies on platforms with heterogeneous hardware.

Related Activities

Among the related activities which we wish to augment are the following:

- The Center and Application Monitoring Session held during the ECP Annual Meeting.
- The International Workshop on Monitoring and Operational Data Analytics (MODA) held with the annual ISC High Performance conference.
- The Workshop on Monitoring and Analysis for High Performance Computing Systems Plus Applications (HPCMASPA) held with the annual IEEE Cluster conference.
- The Workshop on Performance Monitoring and Analysis of Cluster Systems (PMACS) held with the annual Euro-Par conference.

Each of these related activities share an interest in the wealth of information exposed by these systems about how the system resources are being utilized. Our Symposium is unique in its emphasis on applying data to improve the codesign process. The Quantitative Codesign Symposium also has a distinguishing format and venue.

Appendix 2 – Speaker Biographies

Terry Jones

- Terry Jones is a Senior Research Staff member at Oak Ridge National Laboratory (ORNL) where he has worked since 2008 in the Computer Science and Mathematics Division (CSMD) as a Computer Scientist.
- Prior to that, he held a Computer Scientist position at Lawrence Livermore National Laboratory (LLNL).
- Terry earned a Master of Computer Science degree from Stanford University.
- Terry's research interests include system software for high performance computing, runtime systems and middleware, parallel and distributed architectures; performance monitoring; memory and storage systems; distributed clock synchronization, and resilience for complex distributed systems.

Martin Schulz

- Martin Schulz (male) is a Full Professor and Chair for Computer Architecture and Parallel Systems at the Technische Universität München (TUM), which he joined in 2017, as well as a member of the board of directors at the Leibniz Supercomputing Centre.
- Prior to that, he held positions at the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory (LLNL) and Cornell University.
- He earned his Doctorate in Computer Science in 2001 from TUM and a Master of Science in Computer Science from UIUC.
- Martin's research interests include parallel and distributed architectures and applications; performance monitoring, modeling and analysis; memory system optimization; parallel programming paradigms; tool support for parallel programming; power-aware parallel computing; and fault tolerance at the application and system level, as well as quantum computing and quantum computing architectures and programming, with a special focus on HPC and QC integration.

Jeanine Cook

- Dr. Jeanine Cook is a Principal Member of Technical Staff at Sandia National Laboratories.
- After earning a BS in Electrical Engineering and a MS in Computer Science from the University of Colorado, Jeanine received a PhD in Computer Science from New Mexico State University and joined the faculty in Electrical and Computer Engineering.
- During eleven-year tenure she graduated eight PhD and twelve masters, and in 2009 received the Presidential Early Career Award for Scientists and Engineers.
- Jeanine is a leading researcher in High-Performance Computing, performance characterization and modeling, hardware accelerator technologies, and large-scale system monitoring and data analytics. Additionally, she focuses on diversity issues, instantiating and participating in numerous NSF Broadening Participation In Computing projects, and serving on the Board of Directors for the Center for Minorities and People with Disabilities in Informational Technology. In 2020 Jeanine received the Richard A Tapia Achievement Award.

Mitsuhisa Sato

- Mitsuhisa Sato is a deputy Director of RIKEN Center for Computational Science since 2018.
- Since 2010, he has been appointed to the research team leader of programming environment research team in Advanced Institute of Computational Science (AICS) renamed to R-CCS, RIKEN.
- From 2014 to 2020, he was working as a team leader of architecture development team in FLAGSHIP 2020 project to develop Japanese flagship supercomputer.
- He is a Professor（Cooperative Graduate School Program）and Professor Emeritus of University of Tsukuba.

Duncan Roweth

- Dr. Duncan Roweth is a senior DT at HPE.
- Duncan is in the CTO office working on future HPC networking products
- Duncan came to HPE with the acquisition of Cray, where his main claim to fame is as instigator of the Slingshot program
- Before that he was a founder of Quadrics

PANEL

Jesus Labarta

- Jesus Labarta is professor at the Computer Architecture Department at the technical University of Catalonia and director of the Computer Sciences Department at the Barcelona Supercomputing Center.
- He was the Ken Kennedy Award winner of 2017
- He has been working on performance analysis tools, task based programming models and during the last three years he has coordinated the design of a RISCV Vector processor and programming environment within the EPI project.

Larry Kaplan

- Chief Software Architect for HPC and Distinguished Technologist for HPE, Larry is currently working on the design of the software stack for the next generation Shasta product line.
- Special focus areas for Larry include a more open and modular software ecosystem along with a new advanced network software stack, from drivers up through programming models. Larry is also a contributor to Cray's DOE *Forward programs.
- Larry played a significant role in the design of the Cray XC, XE, XT, XMT, and MTA supercomputers, especially in the areas of runtime, operating and supervisory systems, and hardware/software interface. Past projects have included work in system and application resiliency, virtual and physical memory management systems, and network communication and management software.
- Larry has a master's degree in Computer Architecture from the Courant Institute at NYU and a bachelor's degree in Computer Science and Electrical Engineering from Dartmouth College. He holds over a dozen U.S. patents.

Hatem Ltaief

- Hatem Ltaief is the Principal Research Scientist in the Extreme Computing Research Center at KAUST, where is also advising several KAUST students in their MS and PhD research.
- His research interests include parallel numerical algorithms, parallel programming models, performance optimizations for manycore architectures and high performance computing.
- Hatem received the engineering degree from Polytech Lyon at the University of Claude Bernard Lyon I, the MSc in applied mathematics and the PhD degree in computer science at the University of Houston.
- He has contributed to the integration of numerical algorithms into mainstream vendors' scientific libraries, such as NVIDIA cuBLAS and Cray LibSci. He has been collaborating with domain scientists, i.e., astronomers, statisticians, computational chemists and geophysicists, on leveraging their applications to meet the challenges at exascale.

Gabriel Rockefeller

- Gabe Rockefeller is a research software engineer and software project manager at Los Alamos National Laboratory

Florina Ciorba

- Florina Ciorba is an Associate Professor of High Performance Computing at the University of Basel (UniBas), Switzerland.
- Before joining UniBas, she was a (tenured) senior scientist at the Center for Information Services and High Performance Computing (ZIH) at Technische Universität Dresden, Dresden Germany (2010-2015), where she lead the modeling and simulation efforts of the DFG Collaborative Research Center 912 HAEC.
- Before ZIH, she was a postdoctoral research associate at the Center for Advanced Vehicular Systems at Mississippi State University (MsState), Mississippi State, USA (2008 to 2010) where she worked on performance optimization of codes from computational material science.

- She received her doctoral degree in Computer Engineering in 2008 from National Technical University of Athens, Greece.
- She is a member of the OpenMP ARB, the Message Passing Interface (MPI) Forum, the Energy Efficiency HPC working group, the SPEC High Performance Group (HPG) - where she contributed to the newest SPEChpc 2021 Benchmark Suites, the HiPEAC network.
  Florina and co-authors won best paper awards at the Cluster (2019), IPDPSW ParLearning (2014), ICPDC (2014), and SPDC (2011), and had a top-ranked paper at ISPDC (2019).
- She published numerous peer-reviewed scientific conference and journal articles.
- She was distinguished as IEEE Senior Member in 2020 and is a Life Member of ACM.
- Florina was elected in the Strategy Committee of the German National HPC Association in 2021.
- Her research interests include scalable and robust performance optimization, and system and application operational data analytics.
- More information about Florina and her research group at https://hpc.dmi.unibas.ch/en/people/florina-ciorba/

Esteban Meneses

- Esteban Meneses is the director of the Advanced Computing Laboratory at the Costa Rica High Technology Center.
- He finished his PhD in Computer Science at the University of Illinois at Urbana-Champaign.
- His research work spans several topics of resilience for HPC: data analysis from failure logs of HPC systems, fault tolerance protocol design, and deep learning error sensitivity.

Appendix 3 – Organizing Committee and Program Committee

Workshop Organizing Committee

- Terry Jones - Oak Ridge National Laboratory, USA
- Estela Suarez- Jülich Supercomputing Centre, Germany
- Ann Gentile - Sandia National Laboratories, USA
- Michael Jantz - the University of Tennessee, USA

Workshop Program Committee

- Jim Brandt - Sandia National Laboratories, USA
- Florina Ciorba - University of Basel, Switzerland
- Hal Finkel - US DOE office of Advanced Scientific Computing Research, USA
- Lin Gan - National Supercomputing Center, Wuxi, China
- Maya Gokhale - Lawrence Livermore National Laboratory, USA
- Thomas Gruber - Friedrich-Alexander-University Erlangen-Nuernberg, Germany
- Oscar Hernandez - nVidia, USA
- Jesus Labarta - Barcelona Supercomputing Center, Barcelona, Spain
- Hatem Ltaief, King Abdullah University of Science and Technology (KAUST), Saudi Arabia
- Yutong Lu - Director of National Supercomputing Center in Guangzhou, China
- Esteban Meneses - Costa Rica National High Technology Center, Costa Rica
- Bernd Mohr - Jülich Supercomputing Centre, Germany
- David Montoya - Trenza, USA
- Dirk Pleiter - KTH Royal Institute of Technology, Sweden
- Mitsuhisa Sato - Riken, Japan
- Martin Schulz - Technical University of Munich, Germany

Appendix 4 – Attendees

Due to ACM/SC policy, we are not allowed to publish attendee's names without their explicit consent. In the future, we will seek after this consent during the workshop. For this year, we are able to provide Table 3 with the attendance level by type of attendance.

Table 3 – Observed Attendance Level

Type of Attendance	Count
Hubb (virtual audience)	31
In-person	20
Zoom (speakers & organizers)	14
Total	65