FUTURE COMPUTING COMMUNITY OF INTEREST MEETING OF AUGUST 5-6, 2019

A report by the

FAST-TRACK ACTION COMMITTEE ON STRATEGIC COMPUTING

NETWORKING & INFORMATION TECHNOLOGY
RESEARCH & DEVELOPMENT SUBCOMMITTEE

COMMITTEE ON SCIENCE & TECHNOLOGY ENTERPRISE

of the

NATIONAL SCIENCE & TECHNOLOGY COUNCIL

SEPTEMBER, 2019

About the National Science and Technology Council

The National Science and Technology Council (NSTC) is the principal means by which the Executive Branch coordinates science and technology policy across the diverse entities that make up the Federal research and development enterprise. A primary objective of the NSTC is to ensure that science and technology policy decisions and programs are consistent with the President's stated goals. The NSTC prepares research and development strategies that are coordinated across Federal agencies aimed at accomplishing multiple national goals. The work of the NSTC is organized under committees that oversee subcommittees and working groups focused on different aspects of science and technology. More information is available at https://www.whitehouse.gov/ostp/nstc.

About the Office of Science and Technology Policy

The Office of Science and Technology Policy (OSTP) was established by the National Science and Technology Policy, Organization, and Priorities Act of 1976 to provide the President and others within the Executive Office of the President with advice on the scientific, engineering, and technological aspects of the economy, national security, homeland security, health, foreign relations, the environment, and the technological recovery and use of resources, among other topics. OSTP leads interagency science and technology policy coordination efforts, assists the Office of Management and Budget with an annual review and analysis of Federal research and development in budgets, and serves as a source of scientific and technological analysis and judgment for the President with respect to major policies, plans, and programs of the Federal Government. More information is available at https://www.whitehouse.gov/ostps.

About the Networking and Information Technology Research and Development Program

The Networking and Information Technology Research and Development (NITRD) Program is the Nation's primary source of federally funded coordination of pioneering information technology (IT) research and development (R&D) in computing, networking, and software. The multiagency NITRD Program, guided by the NITRD Subcommittee of the NSTC Committee on Science and Technology Enterprise, seeks to provide the R&D foundations for ensuring continued U.S. technological leadership and meeting the needs of the Nation for advanced IT. The National Coordination Office (NCO) supports the NITRD Subcommittee and the Interagency Working Groups (IWGs) that report to it. More information is available at https://www.nitrd.gov/about/.

About the Fast Track Action Committee on Strategic Computing

The NITRD Subcommittee's Fast-Track Action Committee (FTAC) on Strategic Computing was created to update and disseminate the goals and approaches necessary for sustaining and enhancing U.S. scientific, technological, and economic leadership in strategic computing research, development, and deployment. More information is available at https://www.nitrd.gov/news/Updating-US-Strategic-Computing-Objectives.aspx.

Copyright Information

Event Summary

The Future Computing (FC) Community of Interest (CoI) meeting on August 5–6, 2019, explored the evolving computing landscape to inform agencies about potential opportunities as well as gaps in the Nation's future computing objectives. The meeting focused on where computing will be in the next decade and beyond while also looking at emerging and future applications. It considered the need for new software concepts and approaches to effectively capitalize on new hardware architectures and paradigms. The long period of sustained growth in computing power over the last five decades, characterized by Moore's Law and Dennard Scaling, is expected to end over the next decade. The continued improvement in computing performance will now require moving to new modalities and new means of cooperation and partnership for the benefit of the Nation.

The FC-CoI meeting was held at the offices of the Federal Networking and Information Technology Research and Development Program in Washington, D.C. The meeting brought together key members of industry, academia, and the Federal Government over a two-day period to discuss the future of computing. The meeting had been advertised in the Federal Register to encourage broad participation from the advanced computing community.[1] More detail about the meeting is available at https://www.nitrd.gov/nitrdgroups/index.php?title=FC-COI-2019.

The meeting began with opening remarks by Dr. Jacob Taylor, Assistant Director for Quantum Information Science at the Office of Science and Technology Policy and Co-Chair of the Fast Track Action Committee (FTAC) on Strategic Computing. In his opening remarks, Dr. Taylor noted, "Computation, and associated infrastructure, are critical for our Nation's economy and society. So much of what we do at the bleeding edge of computing right now will, 20 years from now, be part of the fabric of our society. Getting it right matters. We have seen the transformation of our economy and our society driven by the advances made by the community represented here today. We are hopeful that those advances can continue, but they will require working together."

Keynote addresses and plenary talks were given by experts in industry and academia who laid out possible paths beyond Moore's law, new applications and opportunities, and challenges in the software and architecture domains. Three panels were held over the two-day event that presented and discussed (1) views on the current landscape of computing, (2) emerging computing paradigms, and (3) future computing models. During the FC-COI meeting, the participants broke out into small discussion groups to delve deeper into questions surrounding the future of computing in these three topic areas. Many participants agreed that advances in computing should last through the end of this decade, but that the path to continued performance gains will likely come through heterogeneous approaches to computing and considerable concurrent software development.

[1] https://www.federalregister.gov/documents/2019/07/09/2019-14481/notice-of-community-of-interest-meeting-on-future-computing

Key Takeaways

Session I Panel: Current Landscape, Use Cases, and Applications

The Session I panel of experts examined the current technology landscape and changes in applications and use cases brought on by the evolving technology landscape. Panelists noted that rapid changes in hardware—including heterogeneous systems, systems with reduced precision, and neuromorphic computing—are likely to fragment the architecture in the future computing space. This fragmentation will require partnerships to bridge the gap between disparate groups and a significant rethinking on how the United States funds, develops, and sustains software over the coming decades.

Panelists also noted that artificial intelligence (AI) and deep learning are expanding the problem space and broadening the base of applications, requiring hybrid simulation and AI approaches. They also pointed out the need for the community to consider how to "bring the best in technology" to researchers for their science, and to use computing to improve the competitiveness of applications. The Session I breakout discussions focused on areas considered vital to enabling emerging and future applications and use cases.

Breakout I: Software in Support of Applications. Participants saw a holistic approach to software development as an effective way to capitalize on new hardware architectures and paradigms and to accelerate advances in science, engineering, and other disciplines. For this to be effective, a software ecosystem that has well-defined architectures and processes for hardening and integration was considered important. The expectation of a diverse user pool led to discussions on developing models that address how to effectively enable users; this would involve more frequent user and developer interactions, better training to ensure technology is fully utilized, and increased quality control throughout the development lifecycle.

Breakout II: Reducing Barriers to Partnerships. Participants suggested that reducing barriers is best aided by the formation of a unifying vision or galvanizing force. It was suggested that long-term support and direction of software infrastructure and security could best be achieved by a cooperative approach using a consortium. Metrics and trust were highlighted as important aspects of such partnerships. Other discussions centered on building partnership choices that support long-term commitments into research and development business models, and on building discovery and learning partnerships to complement goal-driven partnerships.

Breakout III: Enabling Data. Participants noted that as workflows become more heterogeneous and data-centric, enabling data to be fully utilized to solve the complex problems of tomorrow is paramount. This includes examining standards for data governance, management, veracity/trust, and discoverability; researching data storage and movement to determine the best models; and examining requirements for publishing or publicizing data. Also, participants noted that key needs include interoperability of data and a converged system that is able to handle data analytics, AI, and simulation and modeling workloads.

Breakout IV: Emerging New Applications. Many participants suggested that the emergence of deep learning and the integration of traditional simulation and AI methods will have profound implications for the future of high-performance computing applications and architectures, potentially requiring a programming paradigm shift to collaboration between the programmer and AI for developing software and systems. To capitalize on these changes, challenge problems were suggested as a way to inspire transformative thinking about how people approach solving difficult problems.

Session II Panel: Emerging Computing Paradigms

A number of new computing paradigms for high-end computing are emerging, and many of these technologies could be disruptive in terms of programmability and widespread acceptance. The panel experts agreed that heterogeneity in computing is here to stay and will become much more complex and generation of vast amounts of data will push the scalability limits. They saw this resulting in more disaggregated data models and input/output (I/O) subsystems, where I/O bottlenecks may stand as a barrier to architectural innovations. A major challenge the speakers identified is seamlessly integrating the hardware and software, so that users are isolated from the complexity of the hardware. New partnerships between industry, academia, and the Federal Government will be needed to fully realize and take advantage of the new computing paradigms that future computing offers.

Breakout I: Seamless Integration and Software. Session participants discussed high-level programming systems being essential to integration and usability and requiring research to prepare for future computing paradigms. The issue of seamless integration was observed to involve not just technology but also people—both developers and users—and culture. Achieving this integration culture would require an R&D environment that promotes and follows best practices, improves workforce development, and recognizes the critical role of software engineers.

Breakout II: Potholes in the Hardware Roadmap. Participants recognized that a low-cost path to hardware fabrication is needed to enable exploration of a wide variety of designs and algorithms. Specific areas of interest that were discussed included (1) neuromorphic computing architectures that are both efficient (fast with low space requirements) and highly reconfigurable, and (2) cryogenic electronics for quantum computing. In the area of storage, discussions called out the need for a roadmap for developing scalable and resilient storage. Finally, there was a recognition of the importance of software and the need to evaluate the hardware with application metrics.

Breakout III: Architectural and Systems Considerations. Participants raised a variety of points as to the different architectural and system considerations that should be taken into account, including memory and file system research and the effects of machining learning and AI. Machining learning could be used to design dynamic mapping and scheduling schemes and auto-tuning for heterogeneous hardware. In addition, for research involving multiple communities, participants thought that a phased model where the funding periods of the partners are shifted could address the different stages of work.

Breakout IV: Enabling New Partnership Opportunities. Participants suggested a number of key points to enabling partnerships. Partnerships that are diverse and built on trust through sustained funding tend to work best, especially when based on clear goals with specific objectives. Partnerships also enable accelerated development through better alignment between government requirements and industry trends. Faster contracting processes, even for complex contracts, could also speed innovation.

Session III Panel: Future Computing Models

The final panel of experts addressed the future computing landscape. Panelists explored some future computing models and misconceptions surrounding these models, and they shared their expert opinions on where future computing may be headed. One panelist stated a belief that future computing developments will be based on three axes: classical computing bits, AI and neuromorphic bits, and quantum bits. Another panelist saw the emergence of approximate computing, analog AI cores for memory computation, and optimal materials as defining the next phase. Although there is much uncertainty for the future of computing beyond 20 years, the panelists felt that the community needs to rethink what it means to compute, as well as how workloads will change. They further agreed that a properly trained workforce needs to be in place to fully take advantage of the opportunities that future computing will afford. Panelists also identified the need to develop full software stacks to make future computing paradigms useful.

Breakout I: Workforce. Participants saw workforce development as beginning with sustaining and reinforcing computational thinking in undergraduate science and engineering education. They believed that early interdisciplinary partnerships between computer science departments and domain user communities in engineering and those developing new modes of computing will be key. They emphasized the importance of supporting a curious, motivated, and energetic workforce by establishing an environment that tolerates "constructive failure."

Breakout II: Academia to Industry to Government. A number of participants suggested the creation of grand challenge research problems for fostering interdisciplinary teams along with structuring calls for proposals to encourage collaborations across the government, academia, and industry. Other ideas mentioned included creating partnership councils (lab-initiated with industry and academia) to encourage flow of information, requirements, and trends, as well as covering models for temporary assignments or detailing of researchers across all three groups.

Breakout III: Applications and Users of Future Computing. Participants saw the future of computing as being wide open. The expectation was that a much larger user base would emerge due to greater availability of data. Increased computing usage by end-users who are not interested in the underlying hardware solutions necessitates a higher level of usability, requiring significant investment. Finally, participants acknowledged that the fraction of developers who actually use hardware has been decreasing; yet, given increased demand, the talent pool must be maintained.

Breakout IV: Quantification, Verification, and Validation of Novel Computing Architectures. Participants discussed how novel computing architectures will require extensive quantification, verification, and validation. They noted that many systems cannot be built and tested completely before deployment. As such, metrics and benchmarks will need to change. Participants mentioned building metrics and benchmarks around the applications and validating at a higher level, checking answers rather than bits.

Next Steps

The FTAC appreciates participants for taking the time out of their busy schedule to provide insight and guidance into how the Nation should address future computing needs. Outputs from the meeting will be used to inform the FTAC's efforts on updating the report, *Strategic Computing Update: Enabling the Future of Computing*.