

COMPUTING

edge

- High-Performance Computing
- Hardware
- Artificial Intelligence
- Ethics

OCTOBER 2021

www.computer.org

YEARS
75
IEEE
COMPUTER
SOCIETY

 **IEEE**



Evolving Career Opportunities Need Your Skills

Explore new options—upload your resume today

www.computer.org/jobs

Changes in the marketplace shift demands for vital skills and talent. The **IEEE Computer Society Jobs Board** is a valuable resource tool to keep job seekers up to date on the dynamic career opportunities offered by employers.

Take advantage of these special resources for job seekers:



JOB ALERTS



TEMPLATES



WEBINARS



CAREER
ADVICE



RESUMES VIEWED
BY TOP EMPLOYERS

No matter what your career level, the IEEE Computer Society Jobs Board keeps you connected to workplace trends and exciting career prospects.



75 YEARS
IEEE
COMPUTER
SOCIETY



STAFF

Editor

Cathy Martin

Publications Operations Project Specialist

Christine Anthony

Production & Design Artist

Carmen Flores-Garvey

Publications Portfolio Managers

Carrie Clark, Kimberly Sperka

Publisher

Robin Baldwin

Senior Advertising Coordinator

Debbie Sims

Circulation: *ComputingEdge* (ISSN 2469-7087) is published monthly by the IEEE Computer Society, IEEE Headquarters, Three Park Avenue, 17th Floor, New York, NY 10016-5997; IEEE Computer Society Publications Office, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; voice +1 714 821 8380; fax +1 714 821 4010; IEEE Computer Society Headquarters, 2001 L Street NW, Suite 700, Washington, DC 20036.

Postmaster: Send address changes to *ComputingEdge*-IEEE Membership Processing Dept., 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Printed in USA.

Editorial: Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *ComputingEdge* does not necessarily constitute endorsement by the IEEE or the Computer Society. All submissions are subject to editing for style, clarity, and space.

Reuse Rights and Reprint Permissions: Educational or personal use of this material is permitted without fee, provided such use: 1) is not made for profit; 2) includes this notice and a full citation to the original work on the first page of the copy; and 3) does not imply IEEE endorsement of any third-party products or services. Authors and their companies are permitted to post the accepted version of IEEE-copyrighted material on their own Web servers without permission, provided that the IEEE copyright notice and a full citation to the original work appear on the first screen of the posted copy. An accepted manuscript is a version which has been revised by the author to incorporate review suggestions, but not the published version with copy-editing, proofreading, and formatting added by IEEE. For more information, please go to: http://www.ieee.org/publications_standards/publications/rights/paperversionpolicy.html. Permission to reprint/republish this material for commercial, advertising, or promotional purposes or for creating new collective works for resale or redistribution must be obtained from IEEE by writing to the IEEE Intellectual Property Rights Office, 445 Hoes Lane, Piscataway, NJ 08854-4141 or pubs-permissions@ieee.org. Copyright © 2021 IEEE. All rights reserved.

Abstracting and Library Use: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy for private use of patrons, provided the per-copy fee indicated in the code at the bottom of the first page is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Unsubscribe: If you no longer wish to receive this *ComputingEdge* mailing, please email IEEE Computer Society Customer Service at help@computer.org and type "unsubscribe *ComputingEdge*" in your subject line.

IEEE prohibits discrimination, harassment, and bullying. For more information, visit www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

IEEE Computer Society Magazine Editors in Chief

Computer

Jeff Voas, *NIST*

Computing in Science & Engineering

Lorena A. Barba, *George Washington University*

IEEE Annals of the History of Computing

Gerardo Con Diaz, *University of California, Davis*

IEEE Computer Graphics and Applications

Torsten Möller, *Universität Wien*

IEEE Intelligent Systems

V.S. Subrahmanian, *Northwestern University*

IEEE Internet Computing

George Pallis, *University of Cyprus*

IEEE Micro

Lizy Kurian John, *University of Texas at Austin*

IEEE MultiMedia

Shu-Ching Chen, *Florida International University*

IEEE Pervasive Computing

Marc Langheinrich, *Università della Svizzera italiana*

IEEE Security & Privacy

Sean Peisert, *Lawrence Berkeley National Laboratory and University of California, Davis*

IEEE Software

Ipek Ozkaya, *Software Engineering Institute*

IT Professional

Irena Bojanova, *NIST*

OCTOBER 2021 • VOLUME 7 • NUMBER 10

COMPUTING
eedge



18

Making, Together,
Alone: Experiences
from Teaching a
Hardware-Oriented
Course Remotely

25

The Graphics
System for
the 80's

44

AI Ethics:
A Long History
and a Recent
Burst of
Attention

High-Performance Computing

8 Training Efforts in the Exascale Computing Project

OSNI MARQUES AND ASHLEY BARKER

13 The HPC Certification Forum: Toward a Globally Acknowledged HPC Certification

JULIAN KUNKEL, WERONIKA FILINGER, CHRISTIAN MEESTERS, AND ANJA GERBES

Hardware

18 Making, Together, Alone: Experiences from Teaching a Hardware-Oriented Course Remotely

SUSANNE BOLL, TJADO IHMELS, TOBIAS LUNTE, AND HEIKO MÜLLER

25 The Graphics System for the 80's

NICK ENGLAND

Artificial Intelligence

33 Toward Trustworthy and Responsible Artificial Intelligence Policy Development

RAJ MADHAVAN, JACLYN A. KERR, AMANDA R. CORCOS, AND BENJAMIN P. ISAACOFF

38 Artificial Intelligence in Developing Countries

NIR KSHETRI

Ethics

44 AI Ethics: A Long History and a Recent Burst of Attention

JASON BORENSTEIN, FRANCES S. GRODZINSKY, AYANNA HOWARD, KEITH W. MILLER, AND MARTY J. WOLF

Departments

4 Magazine Roundup

7 Editor's Note: Hone Your HPC Skills

55 Conference Calendar

Subscribe to *ComputingEdge* for free at www.computer.org/computingedge.

Magazine Roundup

The IEEE Computer Society's lineup of 12 peer-reviewed technical magazines covers cutting-edge topics ranging from software design and computer graphics to Internet computing and security, from scientific applications and machine intelligence to visualization and microchip design. Here are highlights from recent issues.

Computer

A Human-Machine Collaboration Model for Urban Planning in Smart Cities

Smart cities are emerging around the world. In response, government institutions have made efforts to implement technologies that promote citizen welfare. The authors of this article from the June 2021 issue of *Computer* explore the advances in inclusive and participatory urban planning processes and propose a conceptual model that helps citizens and governments in the decision-making process.

Computing

A Confidence-Guided Technique for Tracking Time-Varying Features

Application scientists often employ feature tracking algorithms to capture the temporal evolution of various features in their simulation data. However, as the complexity of the scientific features increases along with advanced simulation modeling

techniques, quantification of reliability of the feature tracking algorithms is becoming important. One of the desired requirements for any robust feature tracking algorithm is to estimate its confidence during each tracking step so that the results obtained can be interpreted without any ambiguity. To address this, the authors of this article from the March/April 2021 issue of *Computing in Science & Engineering* develop a confidence-guided feature tracking algorithm that allows reliable tracking of user-selected features and presents the tracking dynamics using a graph-based visualization along with the spatial visualization of the tracked feature.

IEEE Annals

Walking Instead of Working: Space Allocation, Automatic Architecture, and the Abstraction of Hospital Labor

"Space allocation" was a central pursuit in postwar research on computing and architecture. Researchers sought an algorithm that could automatically design the most efficient floor plan for a

set of activities. In this article from the April–June 2021 issue of *IEEE Annals of the History of Computing*, the authors connect an early algorithm for computing architectural floor plans to the postwar British hospital. They examine how researchers adapted algorithmic methods for floor layout design developed in industrial capitalist settings to the promotion of the British welfare state. This article situates the automation of hospital design in postwar UK at the intersection of building science and healthcare management, with the aim to contribute critical perspectives on algorithmic reifications of work in early computer-aided architectural design systems.

IEEE Computer Graphics

An End-to-End Shape-Preserving Point Completion Network

Shape completion for 3D point clouds is an important issue in the literature of computer graphics and computer vision. In this article from the May/June 2021 issue of *IEEE Computer Graphics and Applications*, the authors



propose an end-to-end shape-preserving point completion network through encoder–decoder architecture, which works directly on incomplete 3D point clouds and can restore their overall shapes and fine-scale structures. To achieve this task, they design a novel encoder that encodes information from neighboring points in different orientations and scales, as well as a decoder that outputs dense and uniform complete point clouds. They augment a 3D object dataset based on ModelNet40 and validate the effectiveness of the shape-preserving completion network. Experimental results demonstrate that the recovered point clouds lie close to ground truth points.

IEEE Intelligent Systems

A Basic Framework for Explanations in Argumentation

The authors of this article from the March/April 2021 issue of *IEEE Intelligent Systems* discuss explanations for formal (abstract and structured) argumentation: the question of whether and why a certain argument or claim can be accepted (or not) under various extension-based semantics. They introduce a flexible framework, which can act as the basis for many types of explanations. For example,

they can have simple or comprehensive explanations in terms of arguments for or against a claim, arguments that (indirectly) defend a claim, the evidence (knowledge base) that supports or is incompatible with a claim, and so on. They discuss a real-life application and formally compare their framework to existing work.

IEEE Internet Computing

Examining Machine Learning for 5G and Beyond Through an Adversarial Lens

Spurred by the recent advances in deep learning to harness rich information hidden in large volumes of data and to tackle problems that are hard to model, there is currently tremendous excitement in the mobile networks domain around the transformative potential of data-driven artificial intelligence and machine learning (AI/ML)-based network automation, control, and analytics for 5G and beyond. In this article from the March/April 2021 issue of *IEEE Internet Computing*, the authors present a cautionary perspective on the use of AI/ML in the 5G context by highlighting the adversarial dimension spanning multiple types of ML (supervised, unsupervised, and reinforcement learning) and support this through

three case studies. They also discuss approaches to mitigate this adversarial ML risk, offer guidelines for evaluating the robustness of ML models, and call attention to issues surrounding ML-oriented research in 5G more generally.

IEEE micro

BabelFish: Fusing Address Translations for Containers

In this article from the May/June 2021 issue of *IEEE Micro*, the authors identify that containerized environments create page translations that are extensively replicated across containers in the TLB and in page tables. The result is high TLB pressure and redundant kernel work during page table management. To remedy this situation, this article proposes BabelFish, a novel architecture to share page translations across containers in the TLB and in page tables.

IEEE MultiMedia

Instagram Use as a Multimedia Platform for Sharing Images and Videos: Links to Smartphone Addiction and Self-Esteem

This article from the January–March 2021 issue of *IEEE MultiMedia* aims to analyze and determine

the factors that have an impact on the intensive use of Instagram and its relationship with smartphone addiction and self-esteem. A total of 389 Instagram users aged between 18 and 57 ($M = 23.98$; $SD = 5.37$) completed an online survey based on three standardized scales. The findings of the study suggest that there are statistically significant differences in the intensive use of Instagram and smartphone addiction, according to the individual's employment status and the level of studies completed. Furthermore, the multiple linear regression analysis established age and time spent on the social network as predictors of the intensive use of Instagram and smartphone addiction. Finally, the structural equation model showed a positive correlation between intensive use of Instagram and smartphone addiction, and a negative correlation between smartphone addiction and self-esteem.



Teaching Tangible Interaction Remotely During COVID-19: Transcending Physical Boundaries

Teaching pervasive computing courses is challenging during "normal" times, let alone during a global pandemic. With the transition to remote learning due to the COVID-19 crisis, the authors of this article from the April–June 2021 issue of *IEEE Pervasive Computing* invested a lot of effort into a comprehensive redesign of their

tangible and embodied interaction (TEI) course for an online format. They share their experience and lessons learned from teaching this remote version of the undergraduate TEI course during the fall of 2020.



Understanding the Insecurity of Processor Caches Due to Cache Timing-Based Vulnerabilities

This article from the May/June 2021 issue of *IEEE Security & Privacy* discusses a recently developed test suite for checking timing-based vulnerabilities in processor caches, which has revealed the insecurity of today's processor caches. The susceptibility of caches to these vulnerabilities calls for more research on secure processor caches.



What People Focus on When Reviewing Your App—An Analysis Across App Categories

User-generated app store reviews of mobile apps provide valuable information for developers. However, making sense of this large pool of data across many apps remains a challenge today due to the large data volumes and because app stores lack systematic, built-in analysis tools for user reviews. Read more in this article from the May/June 2021 issue of *IEEE Software*.



Benefit From the Internet of Things Right Now by Accessing Dark Data

Even as companies become aware of the powerful benefits that Internet of Things (IoT) data can bring to their organization, the investment costs of deploying new sensors in the field keep IoT initiatives as distant goals for many firms. At the same time, firms are swimming in substantial amounts of "dark data" that are created by operational technology, transactional systems, and other "things" already deployed by the company. An estimated 90% of this data is currently dark, because most of these "things" were not designed to record data in ways that facilitate aggregation with information from other sources. Using relatively low-cost strategies, companies can begin accessing and exploiting this data gold mine today. Read more in this article from the March/April 2021 issue of *IT Professional*. 🌟

Join the IEEE
Computer
Society
computer.org/join



Editor's Note

Hone Your HPC Skills

Nowadays, scientists, professionals, and students in a wide variety of fields can benefit from learning about high-performance computing (HPC). HPC is helping advance many scientific disciplines, and industries like finance and manufacturing are using HPC for data modeling and analysis. Because of its importance, understanding HPC can be a career boost. In this *ComputingEdge* issue, two articles from *Computing in Science & Engineering* discuss training opportunities that help people develop and perfect their HPC knowledge and skills.

"Training Efforts in the Exascale Computing Project" describes resources available to computational scientists through the US Department of Energy, including webinars, classes, and events related to key supercomputing

hardware and software. "The HPC Certification Forum: Toward a Globally Acknowledged HPC Certification" details a certification program through which HPC practitioners can identify gaps in their knowledge and demonstrate their competencies.

Hardware-oriented education and training can pose special challenges. In *IEEE Pervasive Computing's* "Making, Together, Alone: Experiences from Teaching a Hardware-Oriented Course Remotely," the authors recount how they adapted their lab-based electronics course to a remote format. "The Graphics System for the 80's," from *IEEE Computer Graphics and Applications*, discusses graphics hardware education, research, and development in the 1970s and 80s.

Artificial intelligence (AI) is a powerful technology that can have significant impact in people's

lives. The authors of *IEEE Intelligent Systems' "Toward Trustworthy and Responsible Artificial Intelligence Policy Development"* argue that governments and NGOs can collaborate with industry and academia to adopt policies that cultivate AI's benefits and curtail any negative outcomes. In *IT Professional's "Artificial Intelligence in Developing Countries,"* the author posits that AI applications are helping to grow economies, improve safety, and protect environments in developing countries.

Modern AI often raises ethical concerns. The authors of *Computer's "AI Ethics: A Long History and a Recent Burst of Attention"* highlight writers who considered ethical aspects of AI before it was common to do so, and then present their opinions about ongoing issues in AI ethics. 🤖

DEPARTMENT: EDUCATION

Training Efforts in the Exascale Computing Project

Osni Marques, *Lawrence Berkeley National Laboratory*

Ashley Barker, *Oak Ridge National Laboratory*

This article describes the training activities carried out under the auspices of the U.S. Department of Energy's Exascale Computing Project (ECP). While some of these activities are specific to members of ECP, others can be beneficial to the community at large. We report on training opportunities and resources that the broad computational science community can tap into. We seek to increase awareness about these resources, which we expect to go beyond ECP's scope and life cycle.

The mission of the Exascale Computing Project (ECP) of the U.S. Department of Energy (DOE) is to put together all the pieces in the puzzle for the realization of the nation's first exascale systems. These pieces involve mission-critical applications and an integrated software stack, assembled in concert with the efforts of U.S. high-performance computing (HPC) hardware companies for the identification and development of advanced computer system engineering and hardware components. This amalgam is deemed essential for the deployment of fully functional, capable exascale computing environments. Also, because the architectures of exascale systems will be complex, building upon specialized multicore units and proprietary interconnects and offering unprecedented levels of parallelism.

ECP involves two DOE organizations, the Office of Science (DOE-SC) and the National Nuclear Security Administration, and it is structured on the principles of co-design and integration. Its three key focus areas are as follows.

1. Application Development (AD): Exascale-capable applications are a basic element of ECP and will be the vehicle for delivering

solutions and insights to crucial but up to now intractable challenges.

2. Software Technology (ST): STs play an essential supporting role in application efficiency, and span from low-level system software to high-level applications development tools and libraries.
3. Hardware and Integration (HI): This area ensures the integration of ECP applications, software, and hardware innovations within DOE's computing facilities.

HI also includes a training and productivity (T&P) effort to help address the challenges facing the ECP teams by providing training on key exascale hardware and software technologies. In addition, one of the primary goals of T&P is to help the ECP AD and ST teams become better software developers. As researchers and scientists, team members do not necessarily have backgrounds in computer science or specialize in writing codes. However, in order for science projects to fully exploit the potential of exascale computing, the software must be optimized and ready for the launch of these new exascale systems. Consequently, the T&P activity focuses on critical technologies to improve the productivity of software developers and to advance the sustainability of software products, with emphasis on the dissemination of best practices for software development, testing, and deployment.

The T&P effort facilitates selected application-driven training on topics, such as programming models, tools, libraries and frameworks, data management and workflows, data analysis and visualization tools, system software, on-node parallelism and vectorization, application portability techniques, and software engineering design. The training is conducted through a variety of activities, such as seminars, webinars, deep-dive workshops, lectures, hackathons, and tutorials. Often, these activities serve as an introduction to topics that are pursued in depth by ECP teams that specialize in those topics.

In the next sections, we elaborate on the mechanisms we use for setting up the ECP training agenda, give examples of training sessions (held and planned), and provide some statistics about our target audience. We finish by giving pointers to available resources.

TRAINING AGENDA

The agenda for the T&P effort is determined through feedback we receive from an annual survey of the ECP community, from an advisory council formed by representatives from DOE computing facilities, from interactions with the ECP community (e.g., at the ECP Annual Meeting, see Appendix 1, from participants in the Argonne Training Program on Extreme-Scale Computing (ATPESC),¹ and also the HPC community at large (e.g., at SC conferences).

The annual survey asks participants to rate topics in the following categories: programming languages; programming models and runtimes; compilers; package managers and software distribution tools; data analytics and visualization; data management and workflows; mathematical libraries, scientific libraries and frameworks; performance analysis tools; version control and Web-based hosting services; automated testing frameworks; collaboration tools; and “something else.” These categories are defined in coordination with the AD, ST, and HI areas, and also the advisory council. According to the last survey, the top most desirable topics for training (aggregate of “may have impact” plus “high impact” answers) are listed in Table 1. The table includes information about events that have already been offered (o), planned (p), or remain “to be determined” (u), meaning that as it is the topic is either too broad or training opportunities may be available elsewhere. In the next annual survey, we will attempt

TABLE 1. Top Areas for Training According to 2019 ECP Annual Survey.

Topic	Event
CMake	Webinar (o), tutorial (o/p)
Spack	Webinar (p), tutorial (o)
Building applications within containers	Webinar (o/p), tutorial (o)
Git, GitHub, and Gitlab	Webinar (o), tutorial (o/p)
Python	Webinar (o)
CUDA	Hackathon (o), tutorial (o)
HIP	Webinar (o)
SYCL	Webinar (o)
Advanced MPI	Classes (u)
C++ and Using C++14/17 effectively	Webinar (o), classes (u)
OpenMP 4.5+	Hackathon (o)
AMD GPU	hackathon (o)

to narrow down the needs for the topics considered to be too broad but that are nonetheless desirable.

Communication is an important aspect of the T&P activity. The training section of the ECP website² is the primary way that training opportunities are conveyed to the ECP community. This website provides a comprehensive list of upcoming and past training events (e.g., events related to the topics listed in Table 1). The website also contains archives for all the materials from previous training events, including slides and video recordings of the presentations. These videos are also available on the YouTube channel *ExascaleComputingProject*. This approach allows us to share these materials outside of the ECP. Reaching beyond the researchers involved in ECP today helps to develop the HPC workforce while also laying the groundwork for the second generation of exascale developers and users, after ECP ends and exascale systems become more common and accessible to a broader base of users.

The website also points to specific events hosted by the three computing facilities of DOE-SC's Advanced Scientific Computing Research (ASCR) program: Argonne Leadership Computing Facility, National Energy Computational Research Center, and the Oak Ridge Leadership Computing Facility. These facilities also conduct training on topics of great importance to ECP. Thus, we wish to make sure the ECP community is aware of these training opportunities.

In order to coordinate training events between the ECP and ASCR's computing facilities, the Training Advisory Group (TAG) was formed. TAG meets each month to discuss training plans at each of the facilities and inside ECP; if there is interest and benefit, the members initiate jointly sponsored events, from hackathons to tutorials. This cross-lab collaboration through regular communication allows the ECP to better identify cross-training opportunities as well as potential scheduling conflicts.

The T&P effort uses a training newsletter to communicate upcoming training events. Members of the computational community, not necessarily affiliated with ECP, can also subscribe to the (monthly) ECP Training Newsletter, which has been recently redesigned, and stay informed about training activities that are open to more than just members of ECP.

A notable component of ECP's training effort is the Best Practices for HPC Software Developers (HPC-BP) series of webinars, which is carried out in coordination with the IDEAS-ECP Productivity Project (see Appendix 2). HPC-BP is opened to the general public, all is needed is a free ticket, which serves as a registration and is used for planning and statistics purposes. More information about the HPC-BP series can be found in the work by Marques,³ including the process used for delivering each webinar. We aim at providing a reproducible workflow, not only to produce the webinar series but also to extract best practices that can be helpful to the community. This includes early interactions with the presenters to ascertain that the webinars contain "big-picture takeaways" and pointers to supplementary information. Subsequently, we request that presenters curate the answers to eventual questions from the audience (i.e., in a Q&A document), and add more information as needed. At the time of this writing, 33 webinars have been offered, see *Events* in IDEAS-ECP.⁴ The speakers are from a variety of institutions, including Europe. The average number of tickets we have issued (through Eventbrite) for each webinar is 152, which includes (on average) 47 "affiliated with ECP." Actual participation in the webinars has been (on average) 50% of the number of tickets issued. Many participants in the webinars are recurrent. Concerning participants' affiliations, and as an example, the attendance of the 2020 May webinar was comprised of 30% academia, 56% research laboratories, and 14%

industry. Currently, of particular interest to the series are topics that serve as motivation and preparation for the exascale machines that will be fielded at ASCR's computing facilities. The webinars about the program models Kokkos and SYCL are examples of those. While Kokkos and SYCL warrant more than a 1-h webinar, these presentations pave the way for further learning: We are currently interacting with the developers of Kokkos for the development of an online course on that topic.

Many of the architectural details of the future exascale systems remain under nondisclosure agreement and, thus, not publicly disclosed, but it is anticipated that these systems will have more powerful accelerators per node compared to current systems. As the hardware will be built upon technologies from different vendors, it would be risky for developers to embrace proprietary programming models. This warrants providing training on programming models that are hardware agnostic. Simultaneously, the expected systems will bring new dimensions for the effective development and building of applications, for increased use of reduced precision for improving performance, for the analysis of floating-point software (and debugging), for performance measurements, etc. All these are topics that we consider in the implementation of our agenda.

AVAILABLE RESOURCES AND FINAL REMARKS

At the time of this writing, we have organized or been involved in more than 130 training events (since the start of ECP in September 2016), which have been attended by more than 5000 people. (We keep track of the events in an internal Jira dashboard for the HI component of ECP.) Information about many of these events can be found under the training tab of the ECP website.² Importantly, slides and other documents from the HPC-BP webinar series are preserved under *events* in IDEAS-ECP;⁴ recordings are available at the *ExascaleComputingProject* and *IDEASProductivity* YouTube channels. Given the constraints posed by the COVID-19 pandemic, readers may be interested in the panel series *Strategies for Working Remotely*, found under *events* in IDEAS-ECP,⁴ in which the panelists discuss experiences in working remotely, lessons learned, unforeseen benefits, etc.

Since the focus of our effort is on the needs of ECP, in the near-term some of the events we organize, e.g., on programming models, may have a “restricted audience.” However, we think that the materials from those events can still be beneficial to HPC developers in general, as the building blocks of exascale systems become more widely accessible. Overall, the majority of the events on our agenda has been conceived for a larger audience (see, e.g., the HPC-BP series in IDEAS-ECP).⁴

As ECP moves to completion (in mid-2023), we will continue to analyze the needs of the ECP community and set up a matching T&P training agenda. In this agenda, we will seek to identify best practices that can be adopted or adapted by the computational community in general. Apart from the technical contents of the events we organize, we are interested in exploring innovative modes for delivering online content, which we think would be beneficial for training across a distributed community. This is also motivated by the observation (instructors’ and ours) that events that involve hands-on exercises or demos, for example, tend to be more effective in a face-to-face environment. 🤖

ACKNOWLEDGMENTS

This work was supported by the ECP under Grant ECP 17-SC-20-SC, a collaborative effort of the DOE Office of Science and the National Nuclear Security Administration. The authors are grateful to the presenters of webinars and tutorials mentioned in this article; to people who have provided suggestions for topics and feedback on draft material; and to people who have helped with the logistics of the various T&P events. The work of many has been essential for making a vast amount of information available to the CiSE community.

APPENDIX 1

The ECP Annual Meeting. The ECP Annual Meeting is paramount to the success of the ECP endeavor. The ECP community includes researchers, domain scientists, mathematicians, computer and computational scientists, U.S. HPC vendors, project management experts, an external advisory group, an industry council (formed by senior technology executives from prominent industrial organizations), and ECP’s sponsors and program managers. The meeting highlights

technical accomplishments of the ECP teams and provides a collaborative environment that includes featured speakers, workshops, tutorials, and numerous planning and co-design meetings. Participants are usually unable to attend all discussions and training opportunities that could be relevant to their work: the 2020 annual meeting⁵ included a diverse set of 53 birds of a feather sessions, breakout sessions and panels, and 35 tutorials of various lengths. During the year, we seek to offer training events that are motivated by the agenda of the ECP annual meeting. As an added value, these events are typically open to the participation of non-ECP attendees.

APPENDIX 2

IDEAS-ECP. The Interoperable Design of Extreme-scale Application Software (IDEAS) Project,⁴ also part of the HI area, comprises a set of activities addressing challenges in software development productivity and software sustainability in computational science on high-performance computers. These activities include advanced methodologies for application productivity (e.g., agile workflows for scientific software, supported by metrics and diagnostics to gauge progress); customizable resources for improving the development of scientific application codes (e.g., through software productivity and sustainability plans); engagement with ECP teams to gradually improve software development practices; and training, outreach and community building. As part of its outreach, IDEAS-ECP has offered the *Better Scientific Software* tutorial at various venues, including SC conferences.

REFERENCES

1. ATPESC, “Argonne Training Program on Extreme-Scale Computing.” [Online]. Available: <https://extremecomputingtraining.anl.gov>
2. The Exascale Computing Project. [Online]. Available: <https://www.exascaleproject.org>
3. O. Marques, “Producing a webinar series,” 2019. [Online]. Available: <https://bssw.io>
4. IDEAS-ECP. [Online]. Available: <https://ideas-productivity.org>
5. ECPAM20, “2020 Exascale Computing Project annual meeting,” 2020. [Online]. Available: <https://www.ecpannualmeeting.com/overview.php>

OSNI MARQUES is a member of the Scalable Solvers Group, Computational Research Division, at the Lawrence Berkeley National Laboratory. His research interests include algorithms for the solution of eigenvalue problems, software libraries and tools for the solution of computational problems, methodologies for improving the development of software libraries and applications, and parallel and scientific computing. He is currently co-PI in the DOE-funded projects Advancing Software Productivity for Exascale Applications (IDEAS-ECP), and Frameworks, Algorithms and Scalable Technologies for Mathematics (FASTMath) SciDAC Institute. He is also a collaborator in the project Extreme-scale Scientific Software Development Kit for the ECP (xSDK4ECP). He has contributed to projects funded by the NSF for the development of the LAPACK and ScaLAPACK libraries of algorithms for dense linear algebra calculations. Dr. Marques received the M.Sc. and Ph.D. degrees in structural engineering from the Federal University of Rio de Janeiro, Brazil. Contact him at OAMarques@lbl.gov.

ASHLEY BARKER is the Group Leader for the User Assistance

and Outreach (UAO) team at the Oak Ridge Leadership Computing Facility (OLCF) located at Oak Ridge National Lab (ORNL). UAO is responsible for facilitating access to OLCF resources, providing training, documentation, and technical support to users, collecting and reporting on user facility data, and acquainting the public with the work conducted at the OLCF through scientific highlights. The OLCF supports more than 1200 users and 250 projects annually from a wide spectrum of science domains. She served as the National Climate Research Center (NCRC) Project Director from 2014 to 2016. The NCRC project represents a partnership between NOAA and DOE and through this partnership, the NCRC team has delivered multiple computer systems to NOAA, allowing the agency to advance its climate modeling and improve our understanding of climate variability and change. Ashley is also currently involved in the Exascale Computing Project (ECP) as the Control Account Manager (CAM) for training and productivity. Ms. Barker received the Master of Science degree in information sciences from the University of Tennessee. Contact her at ashley@ornl.gov.



PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.

MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field. **OMBUDSMAN:** Email ombudsman@computer.org
COMPUTER SOCIETY WEBSITE: www.computer.org

EXECUTIVE COMMITTEE

President: Forrest Shull; **President-Elect:** William D. Gropp; **Past President:** Leila De Florian; **First VP:** Riccardo Mariani; **Second VP:** Fabrizio Lombardi; **Secretary:** Ramalatha Marimuthu; **Treasurer:** David Lomet; **VP, Membership & Geographic Activities:** Andre Oboler; **VP, Professional & Educational Activities:** Hironori Washizaki; **VP, Publications:** M. Brian Blake; **VP, Standards Activities:** Riccardo Mariani; **VP, Technical & Conference Activities:** Grace Lewis; **2021-2022 IEEE Division VIII Director:** Christina M. Schober; **2020-2021 IEEE Division V Director:** Thomas M. Conte; **2021 IEEE Division V Director-Elect:** Cecilia Metra

BOARD OF GOVERNORS

Term Expiring 2021: M. Brian Blake, Fred Douglass, Carlos E. Jimenez-Gomez, Ramalatha Marimuthu, Erik Jan Marinissen, Kunio Uchiyama
Term Expiring 2022: Nils Aschenbruck, Ernesto Cuadros-Vargas, David S. Ebert, Grace Lewis, Hironori Washizaki, Stefano Zanero
Term Expiring 2023: Jyotika Athavale, Terry Benzel, Takako Hashimoto, Irene Pazos Viana, Annette Reilly, Deborah Silver

revised 25 May 2021

BOARD OF GOVERNORS MEETING

TBD

EXECUTIVE STAFF

Executive Director: Melissa A. Russell; **Director, Governance & Associate Executive Director:** Anne Marie Kelly; **Director, Conference Operations:** Silvia Ceballos; **Director, Finance & Accounting:** Sunny Hwang; **Director, Information Technology & Services:** Sumit Kacker; **Director, Marketing & Sales:** Michelle Tubb; **Director, Membership & Education:** Eric Berkowitz

COMPUTER SOCIETY OFFICES

Washington, D.C.: 2001 L St., Ste. 700, Washington, D.C. 20036-4928; **Phone:** +1 202 371 0101; **Fax:** +1 202 728 9614; **Email:** help@computer.org

Los Alamitos: 10662 Los Vaqueros Cir., Los Alamitos, CA 90720; **Phone:** +1 714 821 8380; **Email:** help@computer.org

MEMBERSHIP & PUBLICATION ORDERS: **Phone:** +1 800 678 4333; **Fax:** +1 714 821 4641; **Email:** help@computer.org

IEEE BOARD OF DIRECTORS

President: Susan K. "Kathy" Land; **President-Elect:** K.J. Ray Liu; **Past President:** Toshio Fukuda; **Secretary:** Kathleen A. Kramer; **Treasurer:** Mary Ellen Randall; **Director & President, IEEE-USA:** Katherine J. Duncan; **Director & President, Standards Association:** James Matthews; **Director & VP, Educational Activities:** Stephen Phillips; **Director & VP, Membership and Geographic Activities:** Maike Luiken; **Director & VP, Publication Services & Products:** Lawrence Hall; **Director & VP, Technical Activities:** Roger U. Fujii



DEPARTMENT: EDUCATION

The HPC Certification Forum: Toward a Globally Acknowledged HPC Certification

Julian Kunkel, *University of Reading*

Weronika Filinger, *University of Edinburgh*

Christian Meesters, *Johannes Gutenberg-Universität Mainz*

Anja Gerbes, *Goethe-Universität Frankfurt am Main*

The goal of the HPC Certification Forum is to categorize, define, and examine competencies expected from proficient HPC practitioners. The community-led forum is working toward establishing a globally acknowledged HPC certification process, a process that engages with HPC centers to identify gaps in users' knowledge, and with users to identify the skills required to perform their tasks. In this article, we introduce the forum and summarize the progress made over the last two years. The release of the first officially supported certificate is planned for the second half of 2020.

The ever-changing nature of high-performance computing (HPC) has always compelled the HPC community to invest in continual efforts to train new and existing practitioners. Historically, these efforts were tailored around a typical group of users possessing, due to their background, a certain set of programming skills. However, as HPC has become more diverse in terms of hardware, software, and user backgrounds, the traditional training approaches have become insufficient to address the training needs of our community. An increasingly complicated HPC landscape makes the development and delivery of new training materials challenging. During training delivery, educators need to address the knowledge gaps resulting from the diverse backgrounds of their learners. It is not uncommon for an attendee to

come with a specific learning objective related to their work tasks and not be interested in the core HPC knowledge. In that sense, the term HPC practitioner describes anyone involved in providing or using HPC systems, e.g., a user that runs an application on an HPC-resource, a developer for HPC-systems, or an administrator. At the same time, we define the term "HPC" inclusively, capturing parallel computing and cluster computing, e.g., high-throughput computing or multitask computing, as these too suffer from a lack of knowledge with regard to performance issues.

How should we develop training for users, often coming from disciplines that have not traditionally used HPC resources, and are only interested in learning a particular set of skills? How can we satisfy their training needs if we do not really understand what these are? HPC centers struggle to identify and overcome the gaps in users' knowledge, whereas users struggle to identify the skills required to perform their tasks.

The goal of the HPC Certification Forum (HPCCF) is to clearly categorize, define, and examine competencies expected from proficient HPC practitioners. The HPCCF is the central authority, and curates and maintains the certification program. The program consists of three parts: competencies defined in a modular and easily expandable skill tree, an examination process to verify that practitioners possess those skills, and the certification demonstrating their knowledge. Although the forum is not involved in the development of any training materials or tools, it supports the ecosystem around the competencies.

The ultimate goal of the forum is to offer a free, globally acknowledged certification program that will make HPC education and training more transparent and quantifiable for training providers, and easier to navigate for practitioners. This article highlights relevant aspects of our activities, more details can be found in.¹

COMMUNITY-LED FORUM

The forum is organized around several key roles, which include: the general chair, a publicity chair, and curators for the skill-tree, topics, and examinations. While the board leads the effort, members of the community are expected to contribute to the effort, and anyone is free to benefit from it. Active members can gain nomination and voting rights via an annual steering board election. Decision making is lightweight at the moment: while we have defined roles for steering board members that include final authority in the event it is needed, thus far we have made decisions democratically without the need to rely on this formal mechanism. Basically, any contribution is either accepted or discussed and modified until it is accepted.

The forum uses Slack for its monthly meetings and organizes two face-to-face meetings per year (one at ISC-HPC and the other at the annual SC). GitHub and the Forum's webpage are used to coordinate the effort and publish information.

All software used by the forum is Open Source and freely available to allow everyone to participate. The forum aims to provide an ecosystem revolving around the certification specification (including the skill tree and the examination framework), which consists of tools that cover, for example, branding of

teaching materials, referring and cross-linking to the competency definitions, and compiling curricula. In particular, we hope to catalogue and reference the existing content of third-parties to allow practitioners to browse the skills and navigate to relevant open and commercial teaching material.

Note that there is currently no direct funding for the effort, but we support all proposals and efforts that members bring forward and associate their work with the forum. For example, in the ESIWACE project, some contributions regarding HPC IO are expected. Ultimately, we believe that the sustainability of the effort depends upon the recognition of its importance and the voluntary contribution of institutions and individuals.

CATEGORIZATION OF COMPETENCIES

The forum groups a well-defined set of competencies into a skill, and a skill is identified by a set of learning outcomes and relevant metadata that clearly specifies what a practitioner should be able to do to be said to possess that skill. The skills are organized in a tree structure from a coarse-grained representation (corresponding to the tree branches) to a fine-grained representation mapped onto the tree leaves. On the leaf level, a skill is orthogonal to other skills—their narrowed scope means they intentionally can be taught in sessions ranging from a 1.5-hour lecture up to a 4-hour workshop. Skills may cover technology-specific knowledge, such as the skill “USE1.1-B Command Line Interface” for Linux basics or the skill “K4.2-B SLURM Workload manager,” which describes how a cluster manages user jobs. We believe this granularity allows practitioners to select skills relevant to their circumstances, and allows lecturers to prepare modular training sessions with well-defined content while still achieving comparable training outcomes for a varied range of practitioners’ backgrounds. Cross-linking between skills belonging to different branches is allowed and provides for the reuse of the skill definitions and eases the navigation of the tree.

EXAMINATION

As the Forum aims to keep the certification free for practitioners, an online examination process has

been chosen. The summative assessment will be conducted primarily using multiple-choice questions. For each skill a pool of questions and answers will be created, drawing from both internal and external contributions, and an examination will consist of randomly selected questions. Future developments will include questions beyond the multiple-choice type.

We believe the incentive to deliberately cheat during the assessment, e.g., by having the exam filled by someone else, is low. Therefore, we address this issue in a lightweight and cost-effective fashion. Our process deploys several strategies to minimize the risk of cheating, such as raising the examinees' awareness, using a large pool of questions, setting time limits for each question, and a delay between registering for and taking the actual examination. Since knowledge can quickly become obsolete, each certificate needs to indicate when (month and year) the qualifying examination took place. Also, because the examination of a single fine-grained leaf-level skill would be too easy to pass with short-term memorization and more prone to cheating, the certificates bundle multiple skills together. To ultimately provide trust, the Forum hopes to provide the automatic generation of short tests for prospective employers that would allow validating the knowledge of applicants under their own supervision or in assessment centers.

We manage exam questions internally within the HPCCF consortium. To facilitate external contributions, we provide on our wiki-webpage a light-weight interface for suggesting a question for each skill-set which can be used by anyone—those skill-related questions are reviewed and ultimately managed by the exam curator who is part of the steering board. A well-defined process is being created for prospective contributors to verify the suitability of the contributed questions. For a question to be approved, it will need to undergo a review process where comprehensibility, logic (does it have at least one clear answer?) and rigor (does the question lead to the expected answer?) have to be met. Particularly, each exam will be designed considering a variety of aspects, such as competency assessment, learning goals, exam objectivity, and taxonomy of educational objectives.²

As HPCCF exams are potentially interspersed between courses for a particular examinee, examinees may profit from the forward learning effect, i.e.,

the outcome of such a lightweight test may influence how a practitioner will study the material further. Lecturers, who will be provided with “their” examinees summarized outcomes, may profit from insights to conceptual issues course participants may exhibit.

The questions are the only proprietary component of the HPCCF—using restrictive license terms for authors while giving them credit. By providing potential employers with our tests, too, will have the means to test how an exam is seen and complete it from a student's point of view; this is a measure with established relevance.³

RELATED EFFORTS

The certification program is a new community-wide effort and to the best of our knowledge nothing similar was attempted at such a scale, both with respect to the comprehensibility of the covered content and the international reach of the effort. The other related efforts mainly focus on either providing a comprehensive catalogue of existing training materials and opportunities, providing simple badges confirming participation in a specific training event, or establishing a branded well-defined content and teaching practices recognized by the community. A number of institutions and organizations attempted to catalogue the existing training materials, keep a list of training events and bring the HPC training community together, the ongoing efforts include: PRACE Training Portal, the HPC University and the ACM SIGHPC Education Chapter. An example of a badging effort is XSEDE Training Badges Program. Finally, the last example refers both to the Carpentries initiative, and HPC Carpentry which was developed in recent years. The HPC Certification Forum recognizes the importance of these efforts and is actively engaging with their contributors.

CONCLUSION

The program of the HPCCF allows the existing content to be reused but also makes it possible to create a new ecosystem in which HPC centers, research labs, academic institutions, and commercial companies could offer the best of their teaching material. The HPCCF aims to support existing activities and complements them by providing a unified and clear way of mapping out the relevant HPC competencies.

It should be emphasized that the HPCCF does not regulate the content of training material; we purposely separate the definition of skills, the examination, and the certification from the content delivery. The program does not prescribe a curriculum or any fixed order in which skills should be obtained, thus providing flexibility. It eases the navigation between different competencies without being overly restrictive. We are hoping that a majority of the existing and newly created teaching resources can be branded indicating the skills they cover.

We believe the program will bring multiple benefits to everyone involved in HPC teaching and training. Making clear what skills are required or recommended for a competent HPC user would be helpful to both the HPC service providers and practitioners. Training providers could bundle together skills that are most beneficial for specific user roles and scientific domains, which would allow practitioners to browse through skills to quickly identify and learn the skills required to perform their tasks. The variety of training offered within the HPC community makes finding the right resources more complicated than it should be. We hope that the certification program will eventually provide useful information on where the desired skills are taught. The examination confirming that a certain set of competencies has been acquired makes the learning process more complete and meaningful.

By participating in the program, HPC training providers can increase the visibility of their teaching opportunities and share their resources more effectively. The mapping of the skills defined by the program onto the existing training materials should also help to identify any potential gaps and improve the integrity of the offered training. Finally, the certificates recognized by the whole HPC community will simplify the intercomparison of independently offered courses and provide additional incentives for participation. Overall, the flexibility of the program allows for the construction of more personalized and just-in-time pathways to learn about HPC.

To achieve these goals, the forum welcomes contributions from volunteers. For this initiative to truly fulfill its role the involvement of the members of the HPC training community with diverse backgrounds and experiences is required. As the HPCCF community is managed collaboratively according to

self-managed evolving rules,⁴ it is welcoming and the expected contribution is not demanding, due to the informal nature of the forum. 🌍

REFERENCES

1. J. Kunkel, et al., "One year HPC certification forum in retrospective," *J. Comput. Sci. Educ.*, vol. 11, no. 1, pp. 29–35, 2020, doi: 10.22369/issn.2153-4136/11/1/6.
2. L. W. Anderson, "A taxonomy for learning, teaching, assessing. A revision of Bloom's taxonomy of educational objectives," London, U.K.: Pearson, 2Rev edition, Dec. 19, 2000.
3. R. Brink, "Electronic assessment in higher education," *Educ. Stud.*, vol. 37, no. 5, pp. 503–512, 2011, doi: 10.1080/03055698.2010.539733.
4. E. Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge, U.K.: Cambridge Univ. Press, 2015, doi: 10.1017/cbo9781316423936.
5. Webpage of the HPC Certification Forum, 2020. [Online]. Available: <https://www.hpc-certification.org>

JULIAN KUNKEL is currently a Lecturer with the Computer Science Department, University of Reading. He was a Postdoc with the Research Department, German Climate Computing Center (DKRZ), which partners with the Scientific Computing Group, Universität Hamburg. He manages several research projects revolving around high-performance computing and particularly high-performance storage and teaches parallel and distributed computing. Besides his main research focus, he is involved in several community-building activities and aims to improve the training situation for researchers. Contact him at julian.kunkel@gmail.com.

WERONIKA FILINGER is an HPC Applications Consultant with EPCC, The University of Edinburgh. She has been deeply involved in the design and development of the first Massive Open Online Course (MOOC) on Supercomputing, and facilitated all runs of the course. She is teaching the Practical Introduction to HPC online course, which is offered as a part of an online master's program at the University of Edinburgh. For several years, she has been a member of the EPCC outreach team, taking HPC-related activities to public events. She is also the Co-Chair of the Outreach Committee of the ACM SIGHPC Education Chapter and the Publicity Chair of the International HPC Certification Program. She is also

involved in running the International HPC Summer School. Over the years, she has worked on a number of collaborative projects such as CRESTA, ADEPT, APES, SARGASSO and DEEP-EST, and provided consultancy for the Software Sustainability Institute. Contact her at w.filingher@epcc.ed.ac.uk.

CHRISTIAN MEESTERS is a Computational Scientist with the HPC Center, Johannes Gutenberg-University, Mainz. He was a Postdoc with the University Hospital, Bonn, where he developed genetic epidemiology workflows and applications to leverage them to HPC environments. He coordinates local courses and teaching efforts. As the site's life science contact for HPC, he aims to develop generic life science workflows for HPC clusters. Contact him at meesters@uni-mainz.de.

ANJA GERBES is with the Center for Scientific Computing and is a member of Hessian Competence Center for High Performance Computing, Goethe University, Frankfurt/Main. A considerable part of her job role is to develop a range of courses and resources to enable users to work with the cluster. The goal of her Ph.D. is to study the compiler for deficits in terms of performance when translating HPC applications and to understand the limitations of compilers in making the necessary optimizations. These insights can then be incorporated into the compiler for future automatic compiler optimization. Automatic program transformation using source-to-source instrumentation of parallel programs will prepare HPC applications for future performance analysis. Contact her at gerbes@fias.uni-frankfurt.de.

ADVERTISER INFORMATION

Advertising Coordinator

Debbie Sims
Email: dsims@computer.org
Phone: +1 714-816-2138 | Fax: +1 714-821-4010

Advertising Sales Contacts

Mid-Atlantic US:
Dawn Scoda
Email: dscoda@computer.org
Phone: +1 732-772-0160
Cell: +1 732-685-6068 | Fax: +1 732-772-0164

Southwest US, California:
Mike Hughes
Email: mikehughes@computer.org
Cell: +1 805-208-5882

Northeast, Europe, the Middle East and Africa:
David Schissler
Email: d.schissler@computer.org
Phone: +1 508-394-4026

Central US, Northwest US, Southeast US, Asia/Pacific:
Eric Kincaid
Email: e.kincaid@computer.org
Phone: +1 214-553-8513 | Fax: +1 888-886-8599
Cell: +1 214-673-3742

Midwest US:
Dave Jones
Email: djones@computer.org
Phone: +1 708-442-5633 Fax: +1 888-886-8599
Cell: +1 708-624-9901

Jobs Board (West Coast and Asia), Classified Line Ads

Heather Bounadies
Email: hbuonadies@computer.org
Phone: +1 623-233-6575

Jobs Board (East Coast and Europe), SE Radio Podcast

Marie Thompson
Email: marie.thompson@computer.org
Phone: +1 714-813-5094

DEPARTMENT: EDUCATION & TRAINING

Making, Together, Alone: Experiences from Teaching a Hardware-Oriented Course Remotely

Susanne Boll, Tjado Ihmels, Tobias Lunte, and Heiko Müller, *University of Oldenburg*

Imagine you plan to teach a course with electronics and tinkering, with making and crafting in the lab, with students in teams—and then you must run it remotely while keeping physical distance. Our Makers' Lab course suddenly had to be transformed into a remote course due to COVID-19. In this article, we share our experiences on how we ran a practical hardware-oriented course—over a distance.

MAKERS' LAB

Makers' Lab—"Things that Think"—is a practical, hardware-oriented course in the Master's program of computer science at the University of Oldenburg. Over fourteen weeks, teams of students design and implement a physical, interactive system using a variety of prototyping techniques such as Arduino microcontrollers and electronics, 3D-printing, and laser cutting.

In the in-person course, students come in for weekly sessions, work together in teams, and regularly present their team's progress in class. The on-premise sessions are always very interactive, and the different levels of prototypes are brought to class such that everyone can try them out in different stages of development.

With COVID-19, we were faced with a problem: how can we run this course that relies so much on physical proximity, teamwork, and the rapid exchange of ideas between groups, without any meeting in person?

TASK

In previous years, students could choose an assignment from a set of different topics based on our lab's

research projects. This year, we decided on a single assignment for all groups. We anticipated that this would set a level ground for all students and leave enough time for remote supervision of the individual teams.

The assignment for all groups was to create an interactive puzzle box for entertainment purposes—a 20-cm wooden cube that keeps a secret (see Figure 1). For each of the four "sides" of the box, the students were asked to design a riddle using sensors and actuators and to integrate them under one coherent theme. Once the puzzles on all four sides are solved, a mechanism should open the box and reveal a treasure inside.

This assignment supported the remote teaching experience positively in several ways. Most importantly, the task could be broken into several smaller puzzles plus the opening mechanism, allowing individual students within one group to work in parallel, each from their own home. The box has finger joints, thus the box can easily be integrated into one object at the very end. As all groups worked on the same topic, we were able to provide a great number of default materials that were likely to prove useful right from the start.

The nature of the task afforded prototyping by the use of breadboard electronics, which students could do at home, and laser cutting for the wooden shell. The students were introduced to the tools for designing the box with additional engraving and cutouts. The actual cutting was done by one supervisor and then delivered to the students following the necessary safety rules. In this way, the assignment still offered experiences with laser cutting and tools even though the students could not experience them in action in our fab lab.




FIGURE 3. First set of electronics components for one of the students.

It allowed us to set up dedicated channels for meetings, resources, and private channels for each team. We scheduled bi-weekly video conference meetings in which students presented the project progress, challenges, and next steps, followed by discussion. In the off weeks, we offered regular check-ins and Q&A sessions in an online plenary. To facilitate a familiar and social course community, we asked the students to turn on their cameras and actively participate in each session. Additionally, we offered individual team meetings with the tutor whenever needed.

RUNNING THE CLASS

The Makers’ Lab concept follows an iterative design approach, with frequently revised prototypes of increasing fidelity and a final user evaluation at the end of the semester. Of course, user feedback and evaluation are greatly restricted in times of this pandemic. However, we were able to include basic feedback cycles throughout the project by encouraging students to rely on comments from roommates and friends whom they already chose to bubble together with, even if they do not exactly match the target user group. As all groups had chosen target users similar to their peers, this feedback often proved useful. Creative approaches for remote user tests were also sought-after and we pointed the students to a recent collection of literature on doing fieldwork in these times¹.

The first two weeks of the semester were spent

on requirements analysis and project specifications. Students had to define a target user group as well as a context of use and incorporate the arising requirements and capabilities into an overall concept or box theme. Furthermore, students familiarized themselves with the available hardware and gathered inspirations for their very first puzzle ideas from existing resources.

The following fourteen days were dedicated to QnD Prototyping. We provided a brief introduction to the method, followed by a closed intragroup prototyping session along with a short presentation of the intermediate results at the end of the meeting. We used video conferencing for this with students often holding up their prototypes to the camera and telling stories about how certain interactions would work. The remote nature of the session also encouraged groups to work on several ideas in parallel, rather than focus on a single prototype.

Afterwards, students were encouraged to continue working on their QnD prototypes. In this, they were able to collaboratively elaborate and finalize their overall concept and individual puzzle ideas. The prototype also supported them in communicating their ideas (see Figure 4, left) and gathering feedback. The exercise ended with groups dividing responsibilities for the puzzles among their members and transitioning into planning the electronic implementation (see Figure 4, right). All students received individually packed electronic components packages addressing their needs. From now on, every student worked on their own personal components at home, though always in communication with their group members and keeping later integration in mind. Progress, problems, and challenges were presented and discussed in the bi-weekly meetings and, alongside the technical implementation, puzzle concepts were refined based on test runs and new ideas being developed.

In the final three weeks, the students prepared SVG files for the laser cutting based on a common template. After the wooden parts had been cut, they could be picked up outside our institute and the individual sides of each box were implemented. Finally, the individual components were integrated into a complete box for the final prototype. With this, the students were able to run final user evaluations on a very small scale. Figure 5 shows one of the puzzles in




FIGURE 4. Results of QnD prototyping and early integration of first hardware sketches into the cardboard box.




FIGURE 5. Theremin puzzle. Beautiful illustration on the side of the wooden treasure box, integrated with electronics.

which the sound of a siren was to be reproduced by gestural movements.

A final presentation including a demonstration of the prototype marked the end of this successful course (see Figure 6). During the whole term, the students were explicitly instructed to collect material in all stages of their project in form of pictures, videos, sketches, or other artifacts to document all steps of their development.

FEEDBACK FROM STUDENTS

We conducted qualitative interviews with individual students to collect their personal experiences about remote project work.

All participants were extremely satisfied with the course concept, organization and results. Although most had little or no previous experience in working with hardware, all of them were able to work very well from home and became familiar with the technologies




FIGURE 6. Impressions of solving one side of the four languages puzzle box, and the 3-D printed vault riddle of the ancient treasure box.

without relying on human contact. The students appreciated working and iterating with early prototypes for the evolution of ideas, but many would have preferred the conceptual phase (5 weeks) to be somewhat shorter in favor of the technical implementation and integration phase (7 weeks).

Some indicated that the remote project work has enabled them to work together more productively due to less overhead, increased flexibility, location and time independence, and a high availability of both team members as well as lecturers. It has been suggested that some of the provisional solutions should also be considered in a “new normality” as they are more effective and do not depend on physical contact, such as remote Q&A sessions. Others, however, felt restricted by the situation, for example by the lack of physical access to a lab and because there was not as much of a joint development process as under normal circumstances in a collaborative project.

It is important to mention that most members of each group already knew each other from their studies, and that some of them had completed joint projects in the past. We assume that being acquainted was helpful for a course like this.

The main criticism shared by all students was the high integration effort required at the end of the course. Despite all efforts to anticipate integration aspects during early implementation, most teams were confronted with several challenging problems that required improvising. Cables that were too short were soldered together, dried hot glue was melted

with a soldering iron, or cardboard scaffolding was assembled for the multistage placement of components in the box (see Figure 7). All of this would have undoubtedly been noticed earlier in cosituated development.

IMPRESSIONS FROM SUPERVISORS

The first thing that all lecturers noted was how very impressed we were by the students’ projects, both in terms of quality and variety between the groups. While there were initial concerns that the results of the groups might be too similar due to the nature of the task, this semester has proven that students’ individual ideas will develop the same open-ended prompt into vastly different end products. It was exciting to watch how the teams very quickly developed their conceptual idea. One group developed a cursed treasure box in which solving riddles revealed coins which, when placed in the right spot, finally allowed the box to open. Another group dedicated their box to language, with each side of the box featuring a different alphabet and a puzzle related to its unique characteristics. Yet another group designed a wonderful booklet of the story of an abandoned house as a guide to solving the puzzles.

The second big surprise of the semester for us was how well QnD prototyping worked remotely. At this stage of the design process, where prototypes still function more as props in storytelling, rather than semi-functioning parts of a system, videoconferencing




FIGURE 7. Impression of cabling and integration effort for the “ghost eviction in an abandoned house” project.

worked very well in sharing ideas and feedback both within and between groups. In addition, because it was impossible for different people to actively work on the same prototype, students within a group more readily worked in parallel, developing several ideas at the same time. Because each student had materials at home, the QnD session served as an initial kick-off rather than a singular experience. The students then went on to explore different ideas and refine their prototypes. This led to a much higher quality and more creative prototype and was a better source for the next iterations than in the years before. Because of this, we are looking forward to also offering at-home QnD sessions in the future.

Although we only learned of this after the fact, students increasingly met up in their groups towards the end of the semester and considered it essential to finishing their projects. This implies that, while a practical course is possible in a reduced-contact setting and class-level meetings can be moved online, our approach was not fully suited to a strictly no-contact situation. It was fortunate that the situation in Germany safely allowed gatherings of small groups toward the end of the project.

While students could theoretically contact tutors whenever needed, in practice they rarely made use of this opportunity for questions about prototyping or fabrication methods. Thus, more time needs to be budgeted for learning new programs and technologies.

Preparing the files for laser cutting proved especially challenging, even with a template for a basic box provided. A precurated selection of online tutorials might prove useful to aid this process in the future.

Finally, the excellent results of this semester came at a price. While we were deeply impressed with the students’ performance, it became increasingly clear that this is also a result of the amount of effort put into this project. With all work being shifted to the students’ homes and no distinction of place or time between university and daily life, students often worked far beyond what was common in prior years. We believe that any future remote course must be careful to clearly communicate expectations and at what point they are fulfilled, even more than for an in-person class. We will, in the future, investigate sharing strategies for managing work-life balance in working from home with the students.

CONCLUSION

With only a few weeks of preparation, we were able to run a practical, hardware-oriented digital fabrication lab course in times of social distancing. The outcomes were very positive: we turned the class into a good remote learning experience for students under these demanding conditions. In the coming semester, we will be teaching a remote AR/VR lab. We will work on utilizing the lessons we learned in this course, including how to provide materials to students, and how to help

them achieve work-life balance. And, without a doubt, we will also encounter some new challenges. 🌍

ACKNOWLEDGMENTS

We thank the students of this course for their effort, their creativity, and their dedication to this course which was a new challenge for all of us: E. Furuno, T. Ihmels, S. Kühlewind, L.-M. Müller, F. Qaribpour, S. Schuirmann, S. Schultze, N. Trilck, S. Vöge, and A. Withöft.

REFERENCE

1. D. Lupton. Doing fieldwork in a pandemic (crowd-sourced document), 2020. [Online]. Available: <https://docs.google.com/document/d/1clGjGABB2h2qbduTgfqribHmog9B6P0NvMgVuiHZCl8/edit?ts=5e88ae0a#>

SUSANNE BOLL is a professor in media informatics and multimedia systems with the University of Oldenburg,

Oldenburg, Germany. Contact her at susanne.boll@informatik.uni-oldenburg.de.

TJADO IHMELS received the Bachelor's degree in computer science. He is currently working toward the master's degree in the master programme engineering socio-technical systems at the University of Oldenburg, Oldenburg, Germany. Contact him at frank.tjado.ihmels@uni-oldenburg.de.

TOBIAS LUNTE is a resesarch associate with the Media Informatics and Multimedia Systems Group, the University of Oldenburg, Oldenburg, Germany. He received the Ph.D. degree from the University of Oldenburg. Contact him at tobias.lunte@uni-oldenburg.de.

HEIKO MÜLLER is a postdoctoral fellow with the Media Informatics and Multimedia Systems Group, the University of Oldenburg, Oldenburg, Germany. Contact him at heiko.mueller@uni-oldenburg.de.

IEEE Computer Society Has You Covered!

WORLD-CLASS CONFERENCES — Over 215 globally recognized conferences.

DIGITAL LIBRARY — Over 800k articles covering world-class peer-reviewed content.


CALLS FOR PAPERS — Write and present your ground-breaking accomplishments.

EDUCATION — Strengthen your resume with the IEEE Computer Society Course Catalog.

ADVANCE YOUR CAREER — Search new positions in the IEEE Computer Society Jobs Board.

NETWORK — Make connections in local Region, Section, and Chapter activities.

Explore all of the member benefits at www.computer.org today!



DEPARTMENT: VISUAL COMPUTING: ORIGINS

The Graphics System for the 80's

Nick England

Forty years ago, the first General Purpose Raster Graphics Processor made the transition from research project to commercial product. This is the story of the creation of a new graphics system and the startup company that produced it in the early days of raster computer graphics.

Forty years ago, you could buy a Programmable Graphics Processor from Ikonas Graphics that could produce ray-traced, antialiased images of bicubic B-spline surface objects, volumetric rendering of 3-D data sets, and real-time display of constructive solid geometry operations (see Figure 1).

The processor could also perform fast image processing, real-time 3-D line and polygon drawing, and almost anything else that you could program. This recollection is about how this design and the company that made this product came to be.

Note: I have probably mischaracterized, omitted, embellished, miscredited, and otherwise made mistakes in telling this story from dim memories. For convenience, please insert an imaginary IIRC before each sentence.

I will try to convey the background of my idea for designing such a system, the environment in which the prototype took shape, the genesis of a company that turned a lab prototype into a product, some of the technical and business lessons learned along the way, the pleasure of seeing smart people use a tool that I helped create, and what legacy this effort might have left.

GRAPHICS EDUCATION

After graduating in Electrical Engineering and working for a few years, I returned to North Carolina State University (NCSU) in 1972 to learn something about computers and digital hardware. I had done some work as an undergraduate for Professor John Staudhammer, so I contacted him hoping to get a part-time job.



FIGURE 1. Ray-traced image computed entirely on an Ikonas graphics processor designed in 1979.

I walked into his lab and experienced my own Road to Damascus moment—for there I saw a computer screen displaying a moving 3-D object. I was immediately (and still am) hooked.

John had created a computer graphics research group as part of the department's new Signal Processing Lab—back in the days when university fiat required all computing resources to be centralized behind glass walls and fed with punched cards, the Lab's computer had to be called a Signal Processor.

Digital Object Identifier 10.1109/MCG.2020.2983816
Date of current version 28 April 2020.




FIGURE 2. Debbie Ogden's textile design system using Rand tablet, Varian 620, and color display.

It was indeed a very fine Signal Processor, but the Adage AGT-30 was also a great interactive 3-D graphics display embedded within a highly capable 30-bit computer. The AGT-30 had a hybrid digital-analog section for providing 3-D transformations of vector lists on the way from core memory to a calligraphic display. The interactive software tools were amazing in an era of punch cards and dial-up teletype. There was an on-screen editor that even let you use button switches for search/replace/insert, etc. Only a person who has edited code using a Model 33 Teletype can appreciate how wonderful this was. You could insert assembly code into a FORTRAN program to handle real time interrupts and the like, and the instruction set included indirect addressing so making linked lists for creating/editing 3-D object representations was a natural. It was a great machine on which to learn 3-D graphics.

John had brought in a smart gang of grad students and undergraduates to form an amazingly creative research group where hardware and software development were seamless partners. If you needed hardware, you designed and built it; if you needed software, you wrote it—it really was all the same process. Ed Tripp and Dave Wooten designed and built an interface from some additional generic core memory to the AGT-30. That let Jeff Eastman create an interactive 3-D modeling program on the Adage. John got a donation of the bits and pieces of an early educational computer

system that included a slo-mo magnetic video disk, a Varian 620 minicomputer, and a Rand tablet.

Glen Williamson redesigned the disk to act as a 640×480 full color video refresh buffer (one track each for R, G, and B) and Jeff built a FIFO interface from the Varian to write scan lines to the disk. Meanwhile Ed and Dave built a high-speed DMA link between the Adage and Varian. There it was: one of the earliest full color raster display systems tightly coupled to an interactive 3-D computer. As the next step, Jeff wrote a scan-line hidden surface display algorithm on the Varian (in assembly language) which received data from the 3-D modeling program running on the Adage and created a color shaded surface display via the video disk.¹

This hardware/software system started running not long after I showed up in the lab. My first task was designing and building an interface for the Rand tablet and writing a simple sketch program. Later, Debbie Ogden used these tools to help create a paint program for textile design (see Figure 2).

Hardware, software, whatever was needed, we built it. I think that our hardware design orientation (modular functions with well-defined interfaces) showed up in our software as well—I remember that it seemed quite easy to borrow code from others in the lab.

This fertile (and fun!) mixture of hardware and software innovation was where I learned graphics, hardware engineering, and software development. That set the stage for what was to come next in this tale.

MOTIVATION

Around 1977, Professor Ray Stroh got a small research grant from NASA Langley Research Center to investigate new ways of generating cockpit displays. NASA wanted to experiment with raster displays for the “glass cockpit” of the future—up until then everything had been calligraphic display based. I latched onto Ray's project (I honestly do not know what would have happened if he had tried to say no) and started thinking about how to create real-time raster display

of current cockpit info plus support NASA's future display concepts such as "highway in the sky."

By this time, I was experienced and proficient at designing/building pieces of hardware and writing graphics and image processing software. The conventional microprocessors of the time were way too slow to handle this task. At first I thought about dedicated hardware vector and character generators. I built a prototype hardware vector generator that ran about 100 ns/pixel, but it was not clear to me how I could turn that into a complete display generator. I had also designed (but not yet built) a fast processor to drive a calligraphic display—something like an updated AGT-30—but that design was not suitable for driving a raster display either. And then a "perfect storm" of new technology appeared.

BUILDING A MACHINE

Remember that this all happened pre-Internet. The way you found out about something was through magazines and conferences. And it was before gate arrays and ASICs. The only people who could make chips were really big companies like AMD, Texas Instruments, Signetics, etc. A series of articles appeared in one of the trade rags about AMD's new 2901 bit slice family – 4 bits of arithmetic unit, register files, etc., per chip. Put together four of these chips and you had the guts of a very fast 16-bit computer. Other articles appeared about TRW's new 16 × 16 multiplier and multiplier-accumulator chips (you would not believe what a huge pain it was to build a multiplier out of off-the-shelf chips). And another article appeared about Mostek's new 16K × 1 DRAM (wow – 512 × 512 × 1 display in only 16 chips!).

Well, you can guess the result—I realized I could now build a programmable graphics processor plus frame buffer that was fast enough to create the kind of 3-D displays NASA wanted—and give me a wonderful new toy with which to play. I started drafting a design on the backs of sheets of fanfold printout (no PC-based CAD back then) and ordering parts—only \$60 each for those new 16K × 1 DRAMs.

There weren't any PC-based printed circuit layout systems back then either, so I built the prototype the same way we had built other projects in the lab. I used wire-wrap IC sockets on 8" × 10" printed circuit boards laid out using the standard




FIGURE 3. Prototype graphics system at NCSU with logic analyzer probes attached.

technique of the time—black tape on acetate sheets plus Bishop Graphics stick-ons for DIP sockets and edge fingers. I was not absolutely certain of my design (no hardware modeling software then either), so I just laid out power and address lines in the memory array and left all the control logic to be wire-wrapped. I also laid out a board to take the 40 pin DIP packages for the 2901 (eight of them made up my 32-bit processor) and more edge connectors for ribbon cable to connect to the 64-bit wide instruction RAM. This machine had a separate instruction memory (composed of multiple fields to control multiple simultaneous functions) that ran in parallel with the data memory. The processor accessed a 512 × 512 × 2 frame buffer over a 32-bit bus that also had an interface from the host Varian 620 minicomputer. The frame buffer fed a video digital-to-analog converter so I could also run the system as a 256 × 256 × 8 frame buffer (see Figure 3).




FIGURE 4. Wire-wrapped Ikonas programmable matrix multiplier.

Once I had established the processor's instruction set, Ray Stroh wrote a cross-assembler that ran on the Adage. I built and debugged the system and wrote a Bresenham algorithm vector drawing routine for the 32-bit processor. Ray's program generated a display list on the Adage (where it was also displayed on the calligraphic display), passed the display list to the Varian which then sent it along to some static RAM on the display processor's data bus. And there we had it—a real-time raster display system that pleased NASA and looked like it would be a lot of fun to play with.

Naturally, I wanted to do more with it. I got inspired by attending SIGGRAPH in 1977 and started building an augmented-reality system for interactively modeling real-world objects. At this time, fellow grad student Turner Whitted was rendering B-spline surfaces, and I wanted to create models for his display program. Glen Williamson and I built a stereo display from two TV cameras pointed at an object and presented a combined image to the user via a polarization stereoscope (two monitors combined with cross-polarized images). It displayed B-spline patch meshes and cursor objects using my new graphics system combined with the TV camera images.

The mesh computation for my SIGGRAPH '78 paper presentation² actually happened on the Adage and was then displayed via my graphics processor. My plan was to use the TRW multipliers for 3-D transformation as well as for patch evaluation and subdivision. Mary Whitton was busy wire-wrapping the matrix multiplier boards, but I never got to use them for patch evaluation (or to finish my Ph.D. program).

[Aside: Mary and I had wed in 1974. After seeing how

much fun we were having in the graphics lab, she switched from teaching junior high school math to being an enthusiastic electrical engineering student.]

STARTUP

Then, other graphics researchers saw what I had built in the lab and wanted a toy/tool like that for themselves. John Staudhammer and some of his other students had built a couple of hardware systems in his basement, so Mary and

I figured we would follow his model. John was smart enough to get out of that business just as we decided to give it a try—our sophisticated business plan was to build four or five systems, take whatever money we made, and go on a vacation to Europe. When the money ran out, we would come home and get real jobs.

If that sounds like we didn't know what we were doing, it's because we didn't. Neither of us had any business experience. We probably should have named the company Clueless, Inc., but after toying around with the usual Tekagraphitronix names we settled on Ikonas Graphics Systems (fellow grad student Stavros Boinodiris contributed the name). Corporate headquarters was the back room of our rental house.

Our first corporate purchase was a power screwdriver to disassemble my model railroad in that back room. Our next purchase was a computer system. Mary assembled a Heathkit H-11 (DEC LSI-11) while I built an ADM-3K terminal kit (kits were cheaper). We also started designing/building a $512 \times 512 \times 8$ demo graphics system. Neither one of us had any electronics manufacturing experience so we went with what we knew—the same wire-wrap board construction I had used at the university (see Figure 4).

At SIGGRAPH '78, we distributed a one-page flyer promising to build most anything a customer wanted and picked up our first purchase order—a batch of frame buffer boards for Chuck Csuri's graphics lab at Ohio State. It was not a complete system because Marc Howard (NCSU lab alum) wanted to build his own system there.

The first order for a complete system, including programmable graphics processor, came from Don




FIGURE 5. Ikonas Graphics system block diagram.

Woodward at the University of Texas Health Science Center, Dallas, TX, USA. That first order really defined the system we'd produce (eventually making about 400 copies). As Don and I talked on the phone he'd ask questions like "Will it do 1024×1024 as well as 512×512 ?", "Does it have video input?", etc., and I'd reply "Sure!" if I had even a vague idea of how I would design such a feature. I certainly had no shortage of naïve optimism and self-confidence.

PRODUCT INTRODUCTION

I designed the system around a synchronous multi-master bus with 32 bits of data plus 24 bits of address. Via this bus, both the host computer and the embedded graphics processor had the same access to memory, control registers, etc. (see Figure 5). Some customers bought display systems with no processor, while a few others added their own custom processors. I added a few more functional units to the Ikonas version of my original processor. The result was a machine that was general purpose, yet optimized for graphics and imaging functions. In a single instruction you could add two numbers, multiply two numbers, increment a loop counter, write a pixel, and branch to a subroutine based on a condition code—all simultaneously in a single clock cycle.

The frame buffer memory could be accessed in multiple ways—one pixel at a time, multiple pixels at once, or as 32 bit general purpose RAM. The frame buffer memory could be used to store a Z buffer as well as multiple frames of color info. The video display stream carried 32 bits per pixel as RGBA or, via a bit-level crossbar switch, in most any other format. The bus and

processor clock was independent from the video pixel and frame rates which were completely programmable.

Actually almost everything was programmable, including the display resolution—the downside being that when you turned the system on, nothing worked until you programmed it.³ The upside of display programmability was that the system became a favorite of visual perception researchers and of graphics terminal developers like DEC, Tektronix, and HP

While I defined the system architecture and designed many of the plug-in boards, development now began including many additional people. We kept on with the "that's how we did it in the lab" philosophy, recruiting people we had known at the NCSU lab: Steve Holzworth to write software and Allan Sadowski part time to help build boards. Mary's brother Robert took on the financial management (not many bookkeepers had a Ph.D. in math). In high hopes, we moved the company out of our house to an office a few miles away. As additional orders started showing up, Allan roped in more NCSU students who turned up to solder and wire-wrap in the afternoons after class (see Figure 6).

Turner Whitted and John Jarvis at Bell Labs got system #2, and Henry Fuchs and Fred Brooks at UNC-CH got #3. Things mushroomed from there and soon we needed some full-time employees. One of Allan's classmates was Xuan Le who volunteered "I have a cousin." Actually, Xuan had several cousins, and eventually we ended up with an assembly area full of Vietnamese refugees. Pot luck lunches were great!

For SIGGRAPH '79, we had finished the new backplane and chassis, built a $512 \times 512 \times 24$ display (a full-color frame buffer was a pretty big deal back then),



FIGURE 6. Students building early Ikonas wire wrap boards.

interfaced it to our LSI-11's Q-Bus, and had some demo programs and images to display. I had not quite finished the new processor design, so I packed a suitcase full of chips and tools. My plan (honest) was to finish the design the first night, wire the boards the next day, and write running demo code before the exhibits closed. This was something far beyond naïve optimism and that plan did not survive past the first night.

But, boy, did we have images to display—Turner Whitted presented a paper with his first famous ray tracing images. So we loaded Turner's images from our 8" floppy disks (one disk each for red, green, and blue) and attracted an amazing amount of interest—we had the hit hardware of the trade show displaying the hit imagery of the technical conference. For SIGGRAPH '80, Turner said "You ought to contact Loren Carpenter at Boeing." Loren's fractal mountain images drew crowds to our 1980 booth...and we also could demonstrate the graphics processor this time. Over the years we continued to showcase customer images along with our own creations—SIGGRAPH and NCGA shows were always hectic, fun, inspiring, and rewarding.

GROWING THE COMPANY

Our system really did become a standard for graphics research labs and it was wonderful to see that our tools were being used in industry and academia (see Figure 7). We provided information on interfacing to the internal bus and several customers built hardware as well as software. For a couple of years over half of the SIGGRAPH technical papers were from researchers using our equipment.

The GRAPHICS SYSTEM for the 80's

RDS-3000 Graphics Processor and Raster Display System

If your graphics and imaging applications are demanding, the IKONAS RDS-3000 series is the system that can meet your need. The RDS-3000 offers:



POWER

- High Speed Architecture designed for computer graphics and image processing
- Fast 32 bit processor for graphics data generation
- Hardware Matrix Multiplier for 3-D transformations, vector products, and filtering operations
- Real Time Video Processing Module for image processing applications
- Video Input Module for real time "frame grabbing"



FLEXIBILITY

- Software selectable 512² or 1024² display format
- Variable frame and line rates: 200-2000 lines/frame
- Pan and scroll in pixel increments, zoom in integer ratios
- Full Window and Viewport Control

PROGRAMMABILITY

- Graphics Processor is completely user micro-programmable and executes the highly parallel code needed for real time and near real time applications
- IDL, the IKONAS DISPLAY LANGUAGE, is a high level command language which makes the IKONAS package of standard graphics routines easy to use.



EXPANDABILITY

- RDS-3000 components are modular allowing easy expansion of systems
- Small frame buffer systems can be upgraded at a later time by adding processing modules and image memories up to 1024² x 32

Photo credits: Terrain Model/R. Wilson, B. Marshall, Computer Graphics Research Group, Ohio State University Mountain/Loren Carpenter, Boeing Computer Services

IKONAS
IKONAS GRAPHICS SYSTEMS, INC.
403 Glenwood Ave
Raleigh, NC 27603, 919/833-5401

Reader Service Number 31

FIGURE 7. Advertisement in the first issue of IEEE Computer Graphics and Applications magazine.

Gary Bishop at UNC-CH and Preston Gurd at Waterloo both created C compilers for our graphics processor. Those compilers let our users be far more productive and creative than having to write 64-bit assembly code that looked like the following:

```
SOLNEG: B6 RA11 CARH0 RPS ALUMAR BD
NCCMEMAC JMPPDF SOLWT1
SOLWT1: RA6 PR ALUMAR CCMEMAC JMPDF.
SOLWR1: RA2 B6 REOS LRESWR MASHIKA
```

Along with the manufacturing expansion, we had recruited NCSU grads Henry Rich and Pete Evans to work on additional hardware design. Mary designed a programmable matrix multiplier and demonstrated B-spline patch evaluation with it to finish her MSEE degree.

We also added several more software developers to create tools and a graphics language (IDL) for

real-time display applications. NASA, Lockheed, Boeing, and many other aerospace groups used IDL for cockpit display prototyping; the last Ikonas systems in daily use were still generating displays in a cockpit training simulator over 25 years later. NASA definitely got their money's worth out of that initial university research grant.⁴

We had a wonderful team of people and it was a fun and exciting time to be on the cutting edge of graphics development. Besides working with academic and industrial research labs, we delivered systems to a number of firms in the entertainment industry—Marks & Marks, Robert Abel, Lucasfilm, Atari, and others. It was great fun to be able to cheer for “our pixels” on TV or at a movie—yes, of course it was all produced with their software, but we still took pride in supplying useful tools to creative people.

As the business expanded, we naturally ran into financing issues but we still didn't know all that much about running a business. We decided to merge into a local software company in a deal put together by a venture capitalist to raise money through a public offering. We realized that may not have been our smartest move when the SEC padlocked the doors of the investment bankers that were going to handle this public offering. Eventually, in 1982 Boston-based Adage, Inc. bought the entire mess in order to acquire Ikonas. This was the same Adage that had made the AGT-30 on which I learned graphics. Adage moved manufacturing to Massachusetts and we no longer had Vietnamese spring rolls at our pot luck lunches. Sigh.

Things definitely changed as part of a larger company (see Figure 8), and one particular benefit was the ability to hire a few new engineers. Tim Van Hook came from Ohio State where he had developed animation in Chuck Csuri's lab and built an extremely fast polygon-drawing hardware system in a startup that had folded. Tim was amazingly creative and productive. His academic training was in Fine Arts, but he taught himself software and hardware development. In 1985–1986, Tim developed the dixel buffer (linked-lists of front and back surfaces at each pixel) for his ray-tracing and real-time CSG code mentioned earlier.⁵

The hardware was the same that I had designed in 1979–1980, but it took time for software development that would exercise it to its full potential.



FIGURE 8. Final version—Ikonas/Adage RDS-3000 with 15" × 15" printed circuit boards.

LESSONS AND LEGACY

We learned a few lessons along the way.

- ▶ Building tools for smart people is fun—understanding your customers and their needs are key.
- ▶ Manufacturing products is tougher and far more expensive than making lab prototypes.
- ▶ Creating a viable business without any experience can be stressful (understatement!).
- ▶ It takes innovative software to show the benefits of flexible hardware.
- ▶ Smart, honest partners are the best investment you can make.
- ▶ When going through airport security, do not refer to your wire wrap tool as a wire wrap gun.

Tim was the major thread that carried on whatever legacy my early programmable graphics processor has had. Tim, Mary, and I left Adage and started Trancept Systems. There Tim designed a board-level programmable graphics + imaging + computation processor (with multiple processing elements and 200-bit-wide instructions) that plugged into Sun Microsystems workstations.⁶ After Sun bought Trancept in 1987, Tim developed volume rendering and other interesting software, added imaging instructions to Sun's SPARC processor, and eventually left for SGI. At SGI, he was architect for the Nintendo 64 graphics processor, then cofounded ArtX where he was architect for the Nintendo GameCube graphics processor.

ArtX was then bought by ATI and I would like to think that Tim's early experiences with the Ikonas system provided some DNA that still shows up in today's GPGPU concepts. Eventually ATI was acquired by AMD (the same company that had provided the chips for my first programmable graphics processor), completing a cosmic computer graphics hardware circle of some sort.

All-in-all it was a Grand Adventure thanks to all the wonderful coworkers and patient customers who made it possible. 🍷

REFERENCES

1. N. England, NCSU Graphics Lab. [Online]. Available: <http://www.graphics-history.org/ncsu/index.htm>
2. N. England, "A system for interactive modeling of physical curved surface objects," *ACM SIGGRAPH Comput. Graph.*, vol. 12, no. 3, pp. 336–340, Aug. 1978.
3. N. England, "A graphics system architecture for interactive application-specific display functions," *IEEE Comput. Graph. Appl.*, vol. 6, no. 1, pp. 60–70, Jan. 1986.
4. N. England, Ikonas Graphics Systems, Inc. [Online]. Available: <http://www.graphics-history.org/ikonas/index.htm>
5. T. Van Hook, "Real-time shaded NC milling display," *ACM SIGGRAPH Comput. Graph.*, vol. 20, no. 4, pp. 15–20, 1986.
6. N. England, Trancept Systems, Inc. [Online]. Available: <http://www.graphics-history.org/trancept/index.htm>

NICK ENGLAND is now retired after multiple startup companies and has fun restoring vintage Teletype and U.S. Navy radio equipment ("see <http://www.navy-radio.com>"). To contact him with comments, corrections, additional info, etc., you can try decvax!harpol!duke!unc!mcnc!ikonas!nick, or, better yet, nick@graphics-history.org.

Contact department editors David Kasik at dave.kasik@gmail.com, Chris Johnson at crj@sci.utah.edu, and Mary Whitton at mcwhitton@gmail.com.



IEEE COMPUTER SOCIETY
Call for Papers

Write for the IEEE Computer Society's authoritative computing publications and conferences.

GET PUBLISHED
www.computer.org/cfp

75 YEARS
 IEEE COMPUTER SOCIETY

IEEE

Toward Trustworthy and Responsible Artificial Intelligence Policy Development

Raj Madhavan, Jaclyn A. Kerr, Amanda R. Corcos, and Benjamin P. Isaacoff,
AAAS Science and Technology Policy Fellows, U.S. Department of State

Artificial intelligence (AI) technologies are already creating new sectors of the economy and revitalizing industries. Governments and societies around the globe are working to determine where and how AI fits into existing legal, ethical, and regulatory frameworks or where new ones are needed. Policy and governance are paramount to ensuring emerging technologies like AI are utilized to their fullest potential to benefit humanity while mitigating inherent risks and addressing societal concerns. This article provides the views of the four American Association for the Advancement of Science (AAAS) Science & Technology Policy Fellows working on a host of issues surrounding AI along with an overview of some of the U.S. Government efforts in these areas.

Advances in sensing and software, combined with the availability of large datasets and streaming data, algorithmic advances in machine learning, and unprecedented levels of funding in both the public and private sectors, are fueling the exponential growth of artificial intelligence (AI) and related emerging technologies that have the potential to transform and positively impact society. An expected outcome of this transformation will be a notable shift in the interaction between humans and machines that will likely lead to symbiotic relationships with significant implications for the global society as a whole. As more industries adopt AI's fundamental technologies, the field continues to drive profound economic changes and quality-of-life improvements worldwide. As with any emerging technology, AI systems* raise unique and complex legal, ethical, societal, and technological issues. In the

public imagination, they sometimes evoke either dystopian scenes of hapless humans fighting for survival against hostile superintelligences or exaggerated utopian hopes for AI as a panacea for many of humanity's problems. Such extreme narratives can obscure the more nuanced questions regarding specific implications in diverse areas, such as social and economic justice, democracy and authoritarianism, and military competition. The perceived lack of transparency and accountability of AI technologies, combined with limited public technical understanding, compounds both of these extreme views.

While public perception, cultural responses, and resulting societal acceptance have enormous implications, policy decisions often play important roles in determining to what extent societies are able to achieve the potential benefits of the new technologies and avoid the associated risks. Policymaking is often reactive, i.e., responding to situations after they occur. But there are tremendous benefits when policymaking is proactive—whether through fostering a

*Artificial General Intelligence (AGI) and Artificial Super Intelligence (ASI) remain aspirational areas of development and are likely many years away. The vast majority of AI research and applications today focus on *Narrow AI*, a system that can perform a specific and specialized task with or without human intervention.

safe environment for rapid innovation or preventing possible setbacks down the line. Thoughtful, proactive policy can be used to build trust, transparency, and ethical principles into the technological ecosystem, guarding against negative social impacts, building international collaboration, and promoting beneficial innovation. Crafting such policy requires technical understanding and discernment, which requires bringing technological expertise into the policy-making process.

The American Association for the Advancement of Science (AAAS) Science & Technology Policy Fellowship (STPF) offers one mechanism by which interested scientists and engineers can contribute to policy-making processes. The AAAS STPF is unique among policy fellowships because it provides scientists and engineers the experience of working in the U.S. Federal Government. Unlike other policy programs aimed solely at early career professionals, the AAAS STPF is open to candidates at any stage of their career, provided they have fulfilled the associated educational requirements.¹ The AAAS STPF places around 250 scientists and engineers across the three branches (executive, congressional, judicial) of the U.S. Government each year. Currently, there are trained professionals with backgrounds in diverse fields, including social sciences, physical sciences, mathematics, and engineering, serving as Fellows within the Federal workforce. This article provides an overview of some of the ongoing U.S. Government efforts and the viewpoints of four Fellows working on a host of issues surrounding AI.

U.S. GOVERNMENT AND INTERNATIONAL AI TECHNOLOGY POLICY EFFORTS

Many nations across the globe have developed and implemented national AI strategies and the United States is no exception. Executive Order (EO) 13859, signed in February 2019, established the American AI Initiative.² The EO lays out the Administration's high-level goals on AI and directs the Federal Government to pursue five pillars for advancing AI: 1) invest in AI research and development (R&D), 2) unleash AI resources, 3) remove barriers to AI innovation, 4) train an AI-ready workforce, and 5) promote an international environment that is supportive of American AI

innovation and its responsible use. The administration utilizes a series of policy, planning, and coordination bodies to implement the EO's high-level goals. The American AI initiative is now over a year and half old, and in this time the government has made significant progress on achieving the goals set out in the EO.³ Critically, the EO has empowered the entire administration to move forward on AI with each federal agency working to advance the American AI Initiative under their own remits.

The Office of Management and Budget (OMB) issued a draft memorandum in January 2020 that provides guidance to all U.S. Federal Departments and Agencies to inform the development of regulatory and non-regulatory approaches regarding technologies and industrial sectors that are empowered or enabled by AI, and it considers ways to reduce barriers to the development and adoption of AI technologies.⁴ Consistent with the EO, OMB guidance on these matters seeks to support the U.S. approach to free markets, federalism, and good regulatory practices, which has led to a robust innovation ecosystem. These regulatory principles are designed to achieve three goals: 1) ensure public engagement, 2) limit regulatory overreach, and 3) promote trustworthy technology.

The United States certainly is not the only nation or group charting a path forward on AI. Other nations have also developed their own AI strategies, including the EU, the United Kingdom, the Netherlands, Australia, Canada, Japan, India, China, and Russia, among others. Several private sector and civil society groups have put forward principles and guidelines regarding how AI systems should be developed and deployed. Discussions about AI issues have also taken place at multilateral and multistakeholder fora, such as the Group of Seven (G7), the Group of Twenty (G20), and the Organisation for Economic Cooperation and Development (OECD). The United States joined with a group of democratic, like-minded partners to lead the way in developing the first set of intergovernmental Principles on AI at the OECD.⁵ These principles reflect the shared fundamental values of the international community for the responsible development and use of AI. These principles were subsequently endorsed by the G7 and the G20 leaders, and several additional countries have endorsed these principles. In addition, the United States joined the G7 nations in launching

the Global Partnership on Artificial Intelligence (GPAI) in June 2020.⁶ GPAI is an international and multistakeholder initiative to guide the responsible development and use of AI consistent with human rights, fundamental freedoms, and shared democratic values.

IMPLICATIONS OF AI POLICY

AI policy decisions and their associated implementation will have complex and far-reaching societal implications, including profoundly influencing the economy, national security, and democracy.

Digital technologies have been a driving force in economic growth since the second half of the 20th century. With continued breakthroughs in machine learning and widespread applications of AI, these technologies are poised to drive the next transformative changes in the global digital economy. The economic impacts of AI are far-ranging, affecting growth, development, and competitiveness of countries around the world. Leadership in AI R&D and talent will be major determinants of technological and economic competitiveness for years to come, as will the policies that encourage this innovation, and that help cultivate and attract the requisite human capital.

Intellectual property (IP) policy, for example, plays an important role in fostering innovation and creativity. As AI technologies advance, it is important that IP policies with respect to AI be reviewed so as to keep pace with the new technologies. As such, the United States Patent and Trademark Office (USPTO) recently requested public input regarding the impacts of AI on the IP law and policy in order to evaluate “whether further examination guidance is needed to promote the reliability and predictability of patenting AI inventions,” which was subsequently expanded to include copyright, trademark, and other IP rights.^{7,8} The World Intellectual Property Organization (WIPO), a specialized agency of the United Nations (UN), also recently recognized the need to hold international conversations on IP in the context of AI since the technologies are developing quickly, intersect multiple disciplines, and cross international borders. WIPO recently released its “Revised Issues Paper on IP Policy and AI”⁹ and held its second “Conversation on AI and IP” to identify the main IP policy questions that need to be addressed in relation to AI, both of which are likely to continue to evolve with further stakeholder input.

Given the number of different IP rights and related interests to consider and that stakeholders may (often) have conflicting views as to whether current policies appropriately protect all of these rights and interests, it is not trivial to review IP policies at the national level, much less at the international level. Gathering and appropriately analyzing enough submissions from a fully representative spectrum of stakeholders will be key to ensuring that—if IP policies do in fact need to be changed—those policies would be updated so as to actually make a useful difference to stakeholders. If the results of these reviews find that IP policies are appropriate for existing and future AI technologies, and so do not require intervention, it will likely satisfy stakeholders to know they already have access to optimal IP policies.

In order to promote continuing beneficial growth from advances in AI technologies, early consideration should also be given to possible second- and third-order repercussions of these advances, such as impacts on inclusivity, well-being, development, and safety and security. Predicting and mitigating potential risks and harms before they can become sources of grievance or backlash is critical to preventing suboptimal reactive policies. Where jobs in some sectors could be lost as a result of automation, for example, governments could institute new policies to support workforce skill development and job transitions (as elaborated in the American AI Initiative). While international economic competition and growth around AI technologies could lead to further digital divides and reinforce patterns of economic inequality, AI and machine learning tools also offer vast potential to contribute to inclusive growth and development through applications to resource management, transportation, education, public health, and other sectors. Policies promoting accessible AI ecosystems, digital infrastructure, and shared data and knowledge can further advance these technologies to become tools of social entrepreneurship and community empowerment in rural and urban communities and developed and developing countries alike. Many of these key considerations at the international level are currently being addressed precisely through the types of multilateral engagements noted in the “U.S. Government and International AI Technology Policy Efforts” section, whereby countries can agree on key principles and norms that should guide

the responsible future development and use of AI technologies.

One particularly critical area for such norm development concerns the impact of AI technologies on digital freedom and democracy. While AI technologies have significant potential to advance democratic values by supporting innovation and novel forms of collaboration to solve societal challenges, they can also risk inadvertent harm or be used deliberately to bolster authoritarian abuses. If AI technologies reinforce biases, for example, their use may exacerbate existing inequalities of opportunity and treatment across communities, negatively impacting perceptions of fairness and justice. Even otherwise socially beneficial applications can sometimes risk undermining individual security and autonomy, or reducing transparency and accountability. In oppressive contexts, some of these technologies may also be used abusively in ways that directly violate human rights and democratic values by curtailing freedom of expression, association, and assembly, allowing arbitrary or unlawful interference with individual privacy, and restricting or manipulating the free flow of information and ideas. For example, while mass data collection and aggregation of the types possible today—through smart city sensors, CCTV cameras, health record databases, and social media platforms—could fuel advances in machine learning and AI technologies for social good, such as efficient energy usage, increased public health, and safety, they may also enable new forms of digital illiberalism. These include, for example: more pervasive forms of surveillance of individual activities in public and private settings; social ranking and predictive policing mechanisms that can be used to coerce individual behavior; and targeted mis- and dis-information campaigns, including the use of bots, sentiment analysis, and synthetic content to affect emotions and cognition, spread extremism or doubt in democratic institutions, influence public opinion, and undermine democratic public sphere discourse.

AI technologies pose a number of novel challenges for policymakers. Their dual-use and cross-sectoral nature can allow closely related applications of the same technologies to be used simultaneously for good or ill. The dominant role of the private sector in many areas of AI technology development might necessitate new models of public–private partnerships and

multistakeholder governance. The level of technical and/or legal expertise required of those engaged in AI policy development and governance may create unintended opaqueness and confusion, leading to a sense of democracy deficit and an environment of diminished trust. The transnational scope of the development and deployment of AI technologies requires simultaneous consideration of the interdependent impacts across many jurisdictions. A collaborative all-hands-on-deck approach will be needed to address these challenges, align interests, and bridge conventional silos of expertise—from the technical to the legal, across sectors and communities, and between stakeholders and countries. By working together, industry, NGOs, academia, governments, and international organizations can promote the adoption of policies and norms that maximize the benefits from AI and machine learning technologies while minimizing negative impacts. Such collaboration will likely require a renewed consideration of first principles of democratic values and governance to build common understandings across governments influenced by different histories, societal challenges, and public concerns. 🤝

ACKNOWLEDGMENTS

The authors are AAAS Science & Technology Policy Fellows at the U.S. Department of State, Washington, D.C. 20520 U.S.A. The views expressed in this article are those of the authors and do not necessarily reflect those of the U.S. Department of State or the U.S. Government.

REFERENCES

1. AAAS STPF. [Online]. Available: <https://www.aaas.org/programs/science-technology-policy-fellowships>
2. American AI Initiative. Feb. 2019. [Online]. Available: <https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence>
3. American AI Initiative Annual Report. Feb. 2020. [Online]. Available: <https://www.whitehouse.gov/wp-content/uploads/2020/02/American-AI-Initiative-One-Year-Annual-Report.pdf>
4. OMB Guidance for Regulation of Artificial Intelligence Applications. Jan. 2020. [Online]. Available: <https://www.whitehouse.gov/wp-content/uploads/2020/01/Draft-OMB-Memo-on-Regulation-of-AI-1-7-19.pdf>

5. OECD. [Online]. Available: <https://www.oecd.org/going-digital/ai/principles/> & <https://oecd.ai/>
6. GPAI. Jun. 2020. [Online]. Available: <https://www.state.gov/joint-statement-from-founding-members-of-the-global-partnership-on-artificial-intelligence/>
7. USPTO Federal Register Notice on AI and Patents. Aug. 2019. [Online]. Available: <https://www.federalregister.gov/documents/2019/08/27/2019-18443/request-for-comments-on-patenting-artificial-intelligence-inventions>
8. USPTO Federal Register Notice on AI and Other IP Rights. Oct. 2019. [Online]. Available: <https://www.federalregister.gov/documents/2019/10/30/2019-23638/request-for-comments-on-intellectual-property-protection-for-artificial-intelligence-innovation>
9. WIPO and AI, [Online]. Available: https://www.wipo.int/about-ip/en/artificial_intelligence/policy.html

RAJ MADHAVAN is currently an AAAS Science & Technology Policy Fellow and Program Analyst with the U.S. Department of State's Division for International Communications and Information Policy in the Bureau of Economic and Business Affairs, Washington, DC, USA, where he is actively involved in the formulation, coordination, and oversight of policy related to information and communication technologies and digital economy. He has previously held appointments with two U.S. Government agencies (DOE/ORNL & DOC/NIST) and two Universities as a faculty member (University of Maryland College Park & Amrita University, India). He has authored and coauthored 200+ papers in journals, conferences, and magazines and has co-edited two books and four journal special issues focusing on robotics, intelligent systems, performance evaluation, benchmarking, standardization, technology-policy, and ethics. He received the Ph.D. degree in field robotics from the University of Sydney, Camperdown, NSW, Australia, and the M.E. (research) in systems engineering from the Australian National University, Canberra, ACT, Australia. He is a Senior Member of the IEEE. Contact him at MadhavanR@state.gov.

JACLYN A. KERR is currently an AAAS Science & Technology Policy Fellow and Science and Technology Policy Adviser with the U.S. Department of State's Office of the Science and Technology Advisor to the Secretary (STAS), Washington, DC, USA, where she is involved in policy work concerning the roles of digital technologies in human rights, development,

and security. From 2016 to 2019, she was a Postdoctoral Research Fellow with the Center for Global Security Research (CGSR), Lawrence Livermore National Laboratory, where she conducted research on cyberstrategy, Internet governance, and the national security implications of emerging digital technologies. She was a cybersecurity Predoctoral Fellow with the Harvard University's Belfer Center for Science and International Affairs and Stanford University's Center for International Security and Cooperation, a Visiting Scholar with Harvard's Davis Center for Russian and Eurasian Studies, and has held research fellowships in Russia, Kazakhstan, and Qatar. She has previous professional experience as a Software Engineer with Comcast and Symantec. She received the Ph.D. and M.A. degrees in government from Georgetown University, Washington, DC, USA, and the M.A. degree in Russian, East European, and Eurasian Studies, and B.A.S. degree in mathematics and in Slavic Languages and Literatures from Stanford University, Stanford, CA, USA. Contact her at KerrJA@state.gov.

AMANDA R. CORCOS is currently an AAAS Science & Technology Policy Fellow and Foreign Affairs Officer with the U.S. Department of State's Office of Intellectual Property Enforcement in the Bureau of Economic and Business Affairs, Washington, DC, USA, where she is actively involved in developing, managing, and implementing intellectual property policy for digital rights management and emerging technologies. Prior to joining the State Department, she was a Postdoctoral Fellow with the Northwestern University Department of Chemistry, and she worked with Northwestern University's Innovation and New Ventures Office to support inventors and entrepreneurs with commercializing their discoveries and technologies. She received the Ph.D. degree in chemistry from the University of Wisconsin–Madison, Madison, WI, USA, in 2016. Contact her at CorcosA@state.gov.

BENJAMIN P. ISAACOFF is currently an AAAS Science & Technology Policy Fellow at the U.S. Department of State working as a Science and Technology Adviser in the Office of the Science & Technology Adviser to the Secretary of State, Washington, DC, USA, where he focuses on artificial intelligence policy. Prior to joining to the State Department, he was the OSA and SPIE 2018-2019 Arthur H. Guenther Congressional Fellow working in the U.S. Senate. He received the Ph.D. degree in applied physics from the University of Michigan, Ann Arbor, MI, USA, in 2018. Contact him at IsaacoffBI@state.gov.

Artificial Intelligence in Developing Countries

Nir Kshetri, *University of North Carolina at Greensboro*

Artificial intelligence (AI) applications are being used to improve economic, social, and environmental performance in developing economies. AI applications in these economies already do some tasks faster and better than humans. For instance, China's Ant Financial's AI-based chatbot system outperforms humans in delivering customer satisfaction.¹

Other applications perform activities that human beings or other technologies are not capable of doing. An autonomous four-wheeled robot developed by Peru's National Engineering University explores mines to detect dangerous substances such as methane, carbon dioxide, and ammonium. It collects information about their levels and trends. Using sensors, it detects location and generates actions to be taken such as the best routes for the workers to escape (<http://www.andina.com.pe/Agencia/noticia-robot-minero-tambien-tendria-utilidad-agricultura-y-desastres-naturales-621069.aspx>).

AI is arguably the fundamental technology of the Fourth Industrial Revolution (4IR).² Unsurprisingly many developing countries have made serious commitment to develop AI capabilities. They are accelerating the process of establishing new policies, regulations, and practices for utilizing AI. Countries such as China, India, Kenya, Mexico, Russia, and the UAE have released AI strategy documents (<https://www.strategy-business.com/blog/Is-AI-the-Next-Frontier-for-National-Competitive-Advantage?gko=931c6>). Others have developed broader visions to transform their economies using 4IR technologies (<https://www.iol.co.za/business-report/economy/wef-africa-botswana-tells-the>

[-world-it-wants-to-reduce-dependence-on-diamonds-31998283](#)).

There is, however, a wide variation among developing countries in the development of the AI industry. China is a clear global leader. Most other countries perform poorly in the main ingredients needed for building high-quality AI applications such as big data, computing power and manpower. Nonetheless, the AI ecosystem is gradually emerging in these countries.

SOME NOTABLE AI APPLICATIONS LAUNCHED IN DEVELOPING COUNTRIES

Table 1 provides illustrative examples of how AI is being used in diverse economic sectors.

Farming and Agriculture

AI has gained popularity in farming and agricultural sectors. Especially large-scale agriculture firms use AI to improve the speed and accuracy of planting and crop management techniques. Brazil's sugar and ethanol producer Raízen has teamed up with the start-up Space Time Analytics to develop an AI-based solution to forecast the sugarcane harvest's size for one year in advance (<https://www.weforum.org/agenda/2017/03/artificial-intelligence-could-help-reverse-latin-america-s-economic-slowdown/>).

AI-based apps are diffusing rapidly among smaller firms. The Brazilian startup Solinftec's AI assistant Alice integrates and processes data from machines, people, climate stations, and other sources (<https://www.nanalyze.com/2019/08/brazil-agriculture-technology/>). To use Alice, farmers embed smart black boxes in their machinery and deploy IoT devices in fields. Alice calculates farmers' needs and provides real-time recommendations (<https://www.croplife.com/precision/growmark-teams-up-with-solinftec-to-increase-farm-efficiency/>).

Digital Object Identifier 10.1109/MITP.2019.2951851

Date of current version 17 July 2020.

TABLE 1. AI Use in Diverse Industries in Developing Economies.

Sector/ industry/activities	Company/app	Use of AI	Outcome
Farming and agriculture	Brazil's Solinftec's AI assistant Alice	Integrates and processes data from machines, people, network of climate stations, and other sources	Provides real-time actionable recommendations
Mining and energy	Chile's Codelco	Monitor mining equipment to predict maintenance needs, detect anomalies and help prevent problems	Efficiency in operations.
Transportation	China's Didi	Predict traffic jams	Minimize the impact of traffic congestions
Healthcare and medical	Nigeria's RxAll	Assesses a drug's compound	Used in many foreign markets.
Finance, banking and insurance	China's Ant Financial	Uses deep-learning technology to detect fraud.	Losses related to fraud: one in 1 million.
HRM	Chile's AIRA	Publishes vacancy announcements, reads/ranks résumés, uses psychometric tests, conducts video interviews and measures performance of applicants.	Reduces time that human recruiters need to spend.

The initial solution was developed for the sugarcane industry. It also offers solutions for other crops. As of August 2019, Solinftec was being used on more than 6.5 million hectares. It monitored 20 000 pieces of equipment and managed 100 000 active daily users (<https://www.nanalyze.com/2019/08/brazil-agriculture-technology/>). Solinftec has opened North American offices in West Lafayette, Indiana (<https://www.precisionag.com/market-watch/brazils-solinftec-sets-up-shop-in-the-u-s/>). It is also expanding to other South American countries as well as in Russia and Ukraine (<https://agfundernews.com/the-road-to-automated-agriculture-begins-in-brazil.html>).

Mining and Energy

Chile's copper producer Codelco uses autonomous trucks at its mines. These vehicles reduce costs by operating for longer hours. They also reduce accidents (<https://hgomezgroup.com/2018/06/13/mining-technology-where-it-is-and-where-it-is-going-in-chile>). It uses AI to monitor mining equipment to predict maintenance needs and make sure that operations run efficiently (<https://www.mining-journal.com/innovation/news/1359598/codelco-to-deploy-ai-solution>).

Transportation

China's mobile-based ride-sharing service Didi uses predictive algorithms in order to help predict traffic jams to minimize the impact of traffic congestions to

drivers.³ The company claimed that by using past as well as real-time data, it can forecast demands 15 minutes in advance with an 85% accuracy level. This information is used to build predictive dispatching models and send vehicles earlier to areas with high degrees of congestion (<https://technode.com/2017/08/28/three-interesting-facts-just-learned-didis-big-data/>).

Healthcare and Medical

Nigeria's RxAll's handheld scanner fights fake drugs. It assesses a drug's compound by connecting the device to cloud-based databases, which contain information related to what the drugs should contain. The information is sent back to the phone. The database is updated using AI. The app also shows which other neighborhoods tested the drug. This gives information related to bad suppliers. It has also been used in Myanmar. The company plans to enter into Ghana, Cambodia, and Kenya (<https://www.theguardian.com/global-development/2019/jun/05/fake-medicine-makers-blockchain-artificial-intelligence>).

Finance, Banking, and Insurance

China's Ant Financial uses big data and AI to manage credit risk and lower loan delinquency rates³ and detect fraud.⁴ Banks in Latin America also use chatbots to provide customer services. A survey found that 83% of consumers of Brazilian financial service would trust banking advice entirely generated

by a computer (<https://www.accenture.com/us-en/insight-financial-services-distribution-marketing-consumer-study>).

Kenya's Kenindia Assurance plans to use AI to detect fraudulent motor insurance claims (<https://www.kbc.co.ke/kenyan-insurers-utilizing-artificial-intelligence-to-curb-fraud-cases/>). In addition to the industry's Integrated Motor Insurance Data System (IMIDS) through the Association of Kenya Insurers, it plans to establish a data center of customers' insurance history.

Human Resources Management (HRM)

The Chilean company AI recruitment assistant's (AIRA) system publishes vacancy announcements in recruitment websites. It reads and ranks the résumés and uses psychometric tests. It also conducts video interviews with applicants. An applicant's performance is measured with indicators related to emotion analytics. Factors such as attention levels and facial expressions are converted into numbers. After all these processes are completed, human recruiters conduct in-depth interviews with the highest-ranked candidates (https://www.accenture.com/_acnmedia/pdf-48/accenture-ai-south-america.pdf?la=es-la).

ECONOMIC EFFECTS OF AI

Some have argued that AI will have a negative impact on developing nations' export-led growth model. AI and automation in developed countries will arguably lead to a decline in the outsourcing of manufacturing and other jobs (<https://www.atlanticcouncil.org/wp-content/uploads/2018/06/The-Global-Innovation-Sweepstakes.pdf>).

However, by using AI in diverse industries, developing nations may compensate for such losses. Even more impressive, technology firms from developing countries such as Brazil, China, Nigeria, and Russia are selling AI solutions in foreign markets. Especially the entry of young AI firms such as Brazil's Solinftec into developed countries is an important trend that can shape AI's global competitive landscape.

Regarding the economic growth mechanisms, traditionally economists viewed new technologies as economic growth drivers through effects on total factor productivity (TFP). The factors of production are capital (e.g., machines and buildings) and labor

that drive economic growth. Broadly TFP is a function of technological, economic, and other factors. TFP measures how efficiently the factors of production are being used. Technologies of the past such as electricity and IT boosted productivity. For example, the World Bank Group's World Development Report 2016 found that the Internet and e-commerce increased TFP in Vietnam by 1.9% and 3.6%, respectively, during 2007–2012.

The global professional services company Accenture argues that in addition to its role in driving TFP, AI also acts as a new factor of production. Various uses of AI are discussed in the previous section to illustrate this possibility. First, AI functions as a new workforce by replicating labor activities at a higher scale and speed. AI performs tasks that humans cannot and learns faster. AI-enabled robots and intelligent machines can also function as a physical capital. Unlike conventional forms of capital such as machines and buildings, AI's self-learning capabilities help improve over time.

Accenture has used case studies of Latin American countries to illustrate AI as a new factor of production. The company forecasts that Brazil's gross value added (GVA) in 2035 will be US\$3452 billion without AI. If AI's impact is limited to TFP, its projected GVA will be US\$3526 billion. When AI is viewed as a factor of production, the value will further increase to US\$3884 billion (https://www.accenture.com/_acnmedia/pdf-48/accenture-ai-south-america.pdf?la=es-la).

OPPORTUNITIES TO DEVELOP THE AI INDUSTRY

There are a number of opportunities to develop AI applications in developing countries. First, local technology hubs are rapidly evolving in many developing countries, which are developing solutions to solve local problems. As of October 2019, Africa had 618 tech hubs, which provide foundations for the AI industry.

Countries such as Ethiopia have launched high-profile AI initiatives. Most of Ethiopia's more than 30 official universities and 130 polytechnics emphasize on technology. In 2012, the Ministry of Science and Technology established its own university and developed a US\$250 million technology park.⁵ The country's AI and robotics research company iCog has produced several apps (<https://nationalinterest.org/blog/buzz/will-ai-cripple-or-leapfrog-developing>

-nations-growth-8771). One such project involves developing software for AI tablets to distribute to children, which would help teach themselves coding, mathematics, and English.⁵

Second, some multinationals are taking advantage of various resources in developing countries to develop cutting edge AI solutions. Such activities are likely to create positive externalities and spillover effects of AI-related knowledge to the local economy. Unilever developed an autonomous forklift in Brazil and launched first in its manufacturing units in the country before any other markets (<https://www.weforum.org/agenda/2017/03/artificial-intelligence-could-help-reverse-latin-america-s-economic-slowdown/>).

A related point is that many developing countries are increasingly participating in the value chain of the global AI industry. One such area is data labeling. According to a 2018 McKinsey report, data labeling is the biggest obstacle to AI adoption.⁶ According to analyst firm Cognilytica, the third-party data labeling solutions market was US\$150 million in 2018, which will increase to more than US\$1 billion by 2023.⁷

Big Western companies are taking major initiatives in order to perfect the ML training. The data-labeling industry employs hundreds of thousands of workers in developing countries such as Kenya, India, and the Philippines. Some examples of data-labeling include teaching self-driving cars, the meanings of road signs, or the difference between a child and a fox. The India- and U.S.-based data annotation company iMerit had 2200 employees in India to label data generated by manufacturing, medical imaging, autonomous driving, retail, insurance agriculture, and other industries. Its Kolkata operation employs 460 women to train computer vision algorithms used in autonomous vehicles and augmented reality systems for companies such as Amazon, Microsoft, eBay, and TripAdvisor.⁷

Nairobi, Kenya-based Samasource labels data for Walmart, Google, Microsoft, Glassdoor, Continental, and General Motors. It employs more than 2800 people.⁷

Some economic sectors are highly digitized, which makes it easier to develop AI-based solutions. For instance, industries such as financial services, telecommunications, and retail in Africa have relatively large amounts of data (<https://www.cio.com/article>

[/3431656/what-africas-approach-to-ai-can-teach-the-world.html](https://www.cio.com/article/3431656/what-africas-approach-to-ai-can-teach-the-world.html)). Likewise, banking and retail industries are expected to dominate AI expenditure in the Middle East (<https://www.thenational.ae/business/technology/ai-spending-to-grow-43-in-middle-east-and-africa-in-2019-1.917352>).

The availability of data at more granular levels is important to develop better algorithms, which is especially critical in personalizing experiences for users. However, doing so might compromise users' privacy or data confidentiality (https://www.accenture.com/_acnmedia/pdf-48/accenture-ai-south-america.pdf?es-la). Whereas ethical and data privacy issues act as a hindrance to the development of the AI industry, such issues are less of a concern in developing countries. For instance, due to China's almost non-existent privacy controls, Chinese companies have easy access to the data of over one billion users. The large datasets help algorithms produce more accurate results and predictions (<https://carnegieendowment.org/2019/01/22/we-need-to-get-smart-about-how-governments-use-ai-pub-78179>).

BARRIERS FACING DEVELOPING COUNTRIES IN THE DEVELOPMENT OF AI

Developing countries experience various barriers and challenges in the development of the AI industry and markets. First, limited datasets are available for AI projects and the available data are of questionable quality. Due to wrong, low quality and irrelevant data, many ICT projects targeted at the poor fail to deliver the benefits that are promised by the initiators of such projects. One example is the "Index-based" crop insurance programs that have been promoted widely among small-scale farmers. The payment to a policyholder relies on satellite images to detect if extreme weather has affected a given area. The area may cover up to 1000 farmers. However, Canada's International Development Research Centre found that the promised benefits have not materialized because the technology gathers data on wide areas. It provides general views about the effects of drought or floods but fails to accurately measure rainfall at a local level. There have been cases in which satellite data indicated that an area had sufficient rainfall but some farmers experienced crop loss due to microclimates.

They were not offered insurance payouts. Some discontinued their insurance schemes (<https://www.reuters.com/article/us-climate-change-kenya-insurance/kenyan-farmers-snap-crops-with-phones-to-improve-insurance-payouts-idUSKBN1WQ0Q7>).

Second, in a discussion of AI's impacts on developing countries, it is also important to discuss noneconomic costs such as the loss of privacy. Such concerns exist at various stages from data gathering to develop AI applications to the actual use of AI. For instance, AI in China is being developed without giving sufficient considerations to ethical issues.⁸

Some countries are using AI to build surveillance tools that might violate privacy laws and human rights. These concerns have been especially strongly voiced against countries that have used China-developed solutions. China has sold AI and facial recognition software in developing countries such as Serbia, Turkey, Russia, Ukraine, Azerbaijan, Angola, Laos, Kazakhstan, Kenya, Uganda (<https://www.scmp.com/news/world/europe/article/3033267/huaweis-facial-recognition-technology-causes-anxiety-serbi>), Ecuador, Bolivian, and Peru (<https://foreignpolicy.com/2018/08/09/ecuadors-all-seeing-eye-is-made-in-china/>).

The Chinese company CloudWalk works with the Zimbabwean government to develop a facial recognition program. China described this as a "win-win" deal. Chinese AI companies can train ML algorithms on Africans to diversify their datasets and Zimbabwe gets access to use the latest technology to monitor its population.⁹ Especially, deployment of China-developed solutions in countries with poor track records on human rights is a concern.

Third, some developing countries are using AI solutions developed by foreign companies, which often perform poorly and have limited usability. A Russian company is reported to be selling face recognition technology in South Africa and Kenya. These systems have a low level of accuracy and precision in Africa.¹⁰

Other impediments include unsupportive regulatory and policy environments. Some countries have a history of cutting Internet connections to tackle problems as varied as cheating by students in exams and political unrest by activists. In June 2019, Ethiopia shut down the Internet for three days during national examinations to prevent students from cheating (<https://www.bloomberg.com/news/articles/2019-06-13/exam-cheats-cited-in-three-day-internet-shutdown-in-ethiopia>).

In January 2019, Zimbabwe disrupted Internet connectivity for a week due to concerns related to unrest in disputed elections (<https://www.bloomberg.com/news/articles/2019-01-21/zimbabwe-lawyers-sue-mobile-operators-over-internet-shutdown>).

SUMMARY

AI applications are tackling economic and social challenges facing developing countries. Economically speaking, AI possesses unique mechanisms that allow it to have significant impacts on economic productivity. While developing countries may experience a decline in outsourcing jobs from developed countries, the potential negative impact of such decline can be minimized by appropriate policy to deploy AI solutions. The true potential of AI comes from the ability to complement as well as enhance traditional factors of production.

AI development in developing countries depends on many factors, inter alia, technology entrepreneurship, knowledge and expertise, data availability, and government policy. The success of an AI project is heavily dependent on the quality of data and algorithms. Especially the lack of data availability is a big obstacle to develop good quality of AI systems. AI performance is influenced by the lower level of training in AI algorithms.

There are major differences between developing and developed countries in AI uses. Due to the lack of huge investments AI-enabled robots and intelligent machines may not be as relevant and important in most developing countries. The most valuable uses of AI would be in performing tasks that humans are not currently capable of doing such as detecting dangerous gas leaks in mines, dealing with traffic congestion, and providing real-time actionable recommendations to farmers. 🌍

REFERENCES

1. W. Knight, "Meet the Chinese finance giant that's secretly an AI company," *MIT Tech. Rev.*, Jun. 2017. [Online]. Available: <https://www.technologyreview.com/s/608103/ant-financial-chinas-giant-of-mobile-payments-is-rethinking-finance-with-ai/>
2. C. Herweijer and D. Waughray, "Fourth industrial

- revolution for the earth harnessing artificial intelligence for the earth," *PWC*, Jan. 2018. [Online]. Available: <https://www.pwc.com/gx/en/sustainability/assets/ai-for-the-earth-jan-2018.pdf>
3. S. Zoo, "What Africa can learn from China about data privacy," *World Econ. Forum*, Jun. 2019. [Online]. Available: <https://www.weforum.org/agenda/2019/06/what-africa-can-learn-from-china-about-data-privacy/>
 4. B. Perez and Z. Soo, "China a fast learner when it comes to artificial intelligence-powered fintech, experts say," *South China Morning Post*, Oct. 2017. [Online]. Available: <https://www.scmp.com/tech/innovation/article/2117298/china-fast-learner-when-it-comes-artificial-intelligence-powered>
 5. C. Galbraith, "Artificial intelligence catches fire in Ethiopia," *HuffPost*, Dec. 2017.
 6. M. Chui, J. Manyika, and M. Miremadi, "What AI can and can't do (yet) for your business," McKinsey & Company, New York, NY, USA, Jan. 2018. [Online]. Available: <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/what-ai-can-and-cant-do-yet-for-your-business>
 7. M. Murgia, "AI's new workforce: The data-labelling industry spreads globally," *Financial Times*, Jul. 2019. [Online]. Available: <https://www.ft.com/content/56dde36c-aa40-11e9-984c-fac8325aaa04>
 8. D. Meyer, "A.I. regulation is coming soon. Here's what the future may hold," *Fortune*, Oct. 2019. [Online]. Available: <https://fortune.com/2019/10/24/german-eu-data-ethics-ai-regulation/>
 9. "Exporting repression? China's artificial intelligence push into Africa," *Council Foreign Relations*, Dec. 2018. [Online]. Available: <https://www.cfr.org/blog/exporting-repression-chinas-artificial-intelligence-push-africa>
 10. J. Snow, "How Africa is seizing an AI opportunity," *Fast Company*, Mar. 2019. [Online]. Available: <https://www.fastcompany.com/90308114/how-africa-is-seizing-an-ai-opportunity>

NIR KSHETRI is a professor of management in the Bryan School of Business and Economics, University of North Carolina at Greensboro, Greensboro, NC, USA. Contact him at nbkshetr@uncg.edu.

Call for Articles

IEEE Pervasive Computing seeks accessible, useful papers on the latest peer-reviewed developments in pervasive, mobile, and ubiquitous computing. Topics include hardware technology, software infrastructure, real-world sensing and interaction, human-computer interaction, and systems considerations, including deployment, scalability, security, and privacy.

Author guidelines:
www.computer.org/mc/pervasive/author.htm

Further details:
pervasive@computer.org
www.computer.org/pervasive

IEEE pervasive COMPUTING
 MOBILE AND UBIQUITOUS SYSTEMS

AI Ethics: A Long History and a Recent Burst of Attention

Jason Borenstein, *Georgia Institute of Technology*

Frances S. Grodzinsky, *Sacred Heart University*

Ayanna Howard, *Georgia Institute of Technology*

Keith W. Miller, *University of Missouri*

Marty J. Wolf, *Bemidji State University*

Artificial intelligence (AI) ethics has become a hot topic in the popular press and in scholarly writing. In this column, five noted scholars give their opinions on what AI issues will become important in the foreseeable future.

During World War II, a Massachusetts Institute of Technology professor named Norbert Wiener worked on the automatic control of a cannon. In 1948, Wiener¹ coined the term *cybernetics* and wrote about computers:

... we are already in a position to construct artificial machines of almost any degree of elaborateness of performance. Long before Nagasaki and the public awareness of the atomic bomb, it had occurred to me that we were here in the presence of another social potentiality of unheard-of importance for good and for evil.

Terry Bynum² cites Wiener’s work in 1948 and in his later book in 1950³ as the start of computer ethics as a scholarly field. Even in this early work, we see the importance of artificial intelligence (AI) issues inside computer ethics, although the formal study of AI is often traced back later to 1955.^{4,5}

Lately, the idea of exploring ethical issues in AI seems commonplace, but it was not always so. We searched Google Scholar for articles or books with a title that includes (“ethics” or “ethical”) and (“AI” or “artificial intelligence”). We got the counts shown in Table 1 and Figure 1 for the years 1985 (when the first such article

TABLE 1. The counts of Google Scholar citations with (“AI” or “artificial intelligence”) and (“ethics” or “ethical”) in the title.

Year	Count	Year	Count	Year	Count	Year	Count
1985	1	1994	0	2003	4	2012	7
1986	1	1995	0	2004	3	2013	5
1987	0	1996	0	2005	6	2014	12
1988	0	1997	0	2006	2	2015	10
1989	0	1998	3	2007	2	2016	21
1990	0	1999	1	2008	6	2017	45
1991	1	2000	4	2009	2	2018	128
1992	0	2001	0	2010	2	2019	334
1993	0	2002	1	2011	8	2020	342

arrived⁶) through 2020. The count for 2020 is only a partial count, up to when this article was written.

Even though the scholarly literature on AI ethics was limited until the last few years, popular culture was far more engaged in issues related to what we now call AI. The term *robot* is often traced back to a 1920 play by Karel Capek⁷ called *R.U.R.* about automated beings revolting against the human race. Isaac Asimov’s Three Laws of Robotics,⁸ later expanded to four laws, have generated debate for decades. Even a short list of films involving AI⁹ is impressive for their treatment of human interactions with AI: *Metropolis*, released in 1927; *The Day the Earth Stood Still*, 1951; *2001: A Space Odyssey*, 1968; *Westworld*, 1973; *Star Wars*, 1977; *War*

Digital Object Identifier 10.1109/MC.2020.3034950

Date of current version: 14 January 2021



COPYRIGHT ISTOCKPHOTO, CREDIT/LVANDY

Games, 1983; *The Terminator*, 1984; *Short Circuit*, 1986; *Star Trek Generations*, 1994; *The Matrix*, 1999; *AI: Artificial Intelligence*, 2001; *I, Robot*, 2004; *WALL-E*, 2008; *Robot and Frank*, 2012; *Ex Machina*, 2015; *Blade Runner 2049*, 2017; and many others. Similarly, television and steaming movies have taken up these themes with a vengeance.¹⁰ In some sense, scholarly interest is merely catching up to popular culture in its focus on ethical issues and AI.




FIGURE 1. The counts of Google Scholar citations with (“AI” or “artificial intelligence”) and (“ethics” or “ethical”) in the title.

LOOKING FORWARD: WHAT ARE SOME IMPORTANT ISSUES TO EXPLORE?

When there is a sudden burst of interest (and publications) in a field, it is important to pay attention to the most significant issues and trends. Otherwise, we can be overwhelmed by minutia and spurious, overhyped speculations.¹¹

Many issues in AI ethics are by no means closed questions that have already been authoritatively resolved. Instead, there is a wide array of contentious arguments which we expect to continue for decades. To give *Computer* readers a closer look at three of these contentious issues, we describe three sets of questions that we think are, and will continue to be, significant AI ethics issues in the foreseeable future.

Issue 1. The ethics of exclusion: Deciding who should have a seat at the table when AI systems are being designed (Jason Borenstein and Ayanna Howard)

Many serious ethical challenges are emerging in relation to AI systems. Among them is how to identify and mitigate different types of biases embedded in the

technology.¹² There are also ethical concerns about the disproportionately harmful impacts AI systems are having on the poor, individuals who reflect gender diversity, and people of color.¹³ Moreover, the “dual use” potential of AI systems is a real worry in the sense that a presumably beneficial AI could be maliciously twisted to cause deliberate harm to the public.¹⁴ Along these lines, an AI system that can recommend medical services based on positively maximizing an individual’s health outcomes can also be used to turn away those same individuals from receiving comparable benefits to maximize a hospital’s revenue stream. Yet in this piece, we want to draw attention to a different issue: one that underlies many of these and other concerns in the realm of AI ethics. It is the overarching question of who should be involved in the process of designing AI systems. Given the widespread impacts that AI systems are having on our lives (both personally and professionally) and how powerful they are in terms of shaping society, drawing attention to who has a seat at the table during the design process is crucial. We contend that it is one of the most pressing ethical issues pertaining to AI of our time. Deployment and use decisions are, of course, crucial to examine as well, but for our purposes here, we limit the scope to design.

Those involved in designing computing devices are typically trained in computer science, engineering, or a related discipline. Of course, the expertise those disciplines provide is necessary, but not sufficient, to the task of designing AI systems. These specialized realms do not consistently expose students to societal implications interconnected with their future professional work. Moreover, computer scientists, engineers, and others may carry with them assumptions about the value neutrality of technology. The commonly held idea that technology is “value-free” can contribute, as a by-product, to the notion that AI will be better than human decision makers. Yet what such beliefs can obscure is that values (from the designer) are being embedded in technology during the design process. For example, a designer might train and validate an algorithm’s output for results biased against one group,

THOSE INVOLVED IN DESIGNING COMPUTING DEVICES ARE TYPICALLY TRAINED IN COMPUTER SCIENCE, ENGINEERING, OR A RELATED DISCIPLINE.

but that doesn’t mean the designer has tested for bias against another type of group or even thinks that it is necessary to do so. This is also true when a designer does not consider the intersectionality of identity attributes. A value-free tool in terms of gender may seem fairly accurate with respect to women, but that does not necessarily mean it works well with Black or Asian women.

Like the rest of us, a designer is a fallible, biased human being. Thus, a strategy is needed for sincerely identifying and acknowledging personal values and overcoming biases (or other shortcomings) when conducting research or designing new technologies. AI cannot achieve the lofty goal of making “better” decisions than humans, assuming that it is even possible, unless diverse voices come together to contribute to the designed solution. Identifying one’s own personal biases is difficult to achieve when everyone in the room has a similar background. And it is hard to unwrap one’s biases if everyone has similar experiences and similar blind spots. For one, groupthink is likely to result. We (the authors) are not necessarily claiming that

designers have bad intent; they probably do not. Many are seeking to promote “responsible computer science,” computing for good, or other worthwhile initiatives. But AI and its applications are so complex and reaching into so many facets of our lives that no singular person, discipline, or field is equipped to understand and represent the perspectives that should be included in the design process. It is only when we are confronted with different voices, backgrounds, and perspectives that, together, we can make a concerted, collective effort to improve.

So, given this state of affairs, the question thus becomes: “Who does and who should have a seat at the table when AI systems are being designed?” We unpack this question into its two component parts: 1) who is currently involved in designing AI systems and 2) who should be involved. In terms of who is currently involved, it is largely computer scientists and engineers. A typical profile of an AI designer is a person situated in a corporate setting in a relatively wealthy region of the world. Present-day AI designers are predominantly White or Asian males.¹⁵

This leads to the harder question, “Who should be involved?” Recognizing the ways in which AI is profoundly changing the world (often in troubling ways), we make the case that the AI design process should be more diverse and inclusive in several senses. This includes striving for more disciplinary diversity among the people taking part in the design process. Sociology, economics, philosophy, the law, gender and race studies, and public policy (among other fields) have valuable insights to share. Diversity in terms of the gender and race of designers is also crucially important; the examples resulting from AI design failures (in part due to the lack of diversity in this sense) are distressing and unfortunately too common, ranging, for instance, from facial recognition failures¹⁶ to health-care treatment errors.¹⁷ Aiming for regional diversity is essential as well. Too often, for example, the Global South finds itself excluded from decisions about emerging technologies generally, and more recently, AI.¹⁸ Finally, it is not just about the designers themselves; it is about who they are interacting with during the design process and when. Potential users and, more broadly, the public must authentically be a part of the picture. A seismic philosophical shift should occur from “What can we design for you?” to “What can we design with you?” so that AI is more authentically aligned with what is good for humanity.

And now comes another, even more difficult, question: “How do we effectively manage the process so that those who should be involved have a seat at the table?” Companies and organizations have tried to move the needle over the last few years, but the needle has only twitched a little. Sending the message that “we welcome you” does not resolve the issue if participation is mere tokenism and has no teeth to produce change. There are no simple solutions to be had for this problem. Yet what may help is a cultural shift in thinking, which acknowledges that finding solutions to what originally seemed like a “technical” design challenge requires engaging with individuals and communities who are not traditionally represented.

Issue 2. The ethics of research and development: The training and deployment of AI systems (Marty J. Wolf and Frances S. Grodzinsky)

Tay was an AI chatbot developed by Microsoft with a goal of learning human speech patterns. Within 24 h, Microsoft researchers had to shut it down when malicious users “taught” Tay to produce and publish anti-Semitic hate speech. In Wolf et al.,¹⁹ we argue that AI, or any software that learns, creates additional risk and places the burden of additional responsibility on not only the software developers who write the AI but also on those who oversee the training of the AI as well. Much of this stems from our concern that it is human subjects who interact with AI systems. This was the case with Tay. Designers did not sufficiently assess the risks to people who came from unleashing it on the open Internet instead of in a closed environment, where it could be closely monitored.

In the United States and many other countries, research on human subjects has a storied past. The notorious 1936 Tuskegee Syphilis Study serves as a marker of how not to conduct research involving human subjects. It also serves as a reminder that not all experiments should be conducted.

Since that time, a rich set of standards has been developed for research involving human subjects. Policies and procedures exist in those institutions where they can be enforced. U.S. universities that get any sort of funding from the U.S. federal government are obligated to ensure that all research that takes place at the university and involves human subjects meets those standards. In response, most universities around the

world have established an institutional review board (IRB) or a similar board that must be consulted at the beginning of a project involving human subjects. Prior to beginning any data collection, the research plan for the project must meet those minimum standards.

The job of the IRB is to ensure that the research procedures are designed in such a way that subjects are not exposed to any risk beyond that encountered in normal daily life. The IRB is responsible for considering physical, psychological, and social risks. Since the IRB must give its approval for the project to go forward, it is in the best interest of those proposing the experiment to consider and address these risks.

THE JOB OF THE IRB IS TO ENSURE THAT THE RESEARCH PROCEDURES ARE DESIGNED IN SUCH A WAY THAT SUBJECTS ARE NOT EXPOSED TO ANY RISK BEYOND THAT ENCOUNTERED IN NORMAL DAILY LIFE.

Privacy and confidentiality are two additional pertinent considerations of the IRB. The protocols and researchers themselves are responsible for ensuring that confidential information, such as names and salient identifying data, are not disclosed outside of the research team. There must also be provisions that protect the confidentiality of subjects whose information is to be retained over an extended period of time. The case of Tay raises two problems for us: First, what happens when the project comes out of private industry? Is there and, if not, should there be the same kind of oversight? And, secondly, what happens, as in the case of Tay, when the human subjects involved are not specifically defined, but rather general Internet users?

There are two other considerations worthy of note in the context of the development of AI on university campuses. Traditionally, computer science faculty have not engaged in research involving human subjects. Thus, within computer science departments there is no culture of considering the impact of “technical” computer science research on people, as would be common in social science or biology research, for example. The second problem arises, at least potentially, when a computer science project is presented to

an IRB. Those on the panel may likely have insufficient experience with or understanding about the complexities of AI, which is necessary to evaluate the proposal with respect to the risks the subjects and society would be subjected to by the research.

Thankfully, many universities have begun to address these and other shortcomings (for example, understanding confidentiality risks arising from “anonymous” data being combined with publicly available data) in the IRB approval process. It is clearly a work in progress. Yet it serves as a model to address an even bigger problem in the development of AI.

Many of those now working in industry who have come out of computer science programs are still subject to this technical view of computing. Their view

*BUILDING WITH SILICON, ELECTRICITY,
AND CLEVER ENGINEERING, AI HAS
ALREADY DELIVERED ARTIFACTS THAT
CAN CONVERSE WITH US, LEARN
WITH US, AND WALK WITH US.*

lines up with the definition of weak AI (as opposed to strong AI). In particular, they understand developing AI as functional, that is, performing a task. In that light, they set a goal and write software using standard, well-understood software development methodologies. While these methodologies may be appropriate for standard software, it is not clear that they work well or are considered to be best practice when the software being developed is designed as self-learning, either partially or fully modifiable.²⁰ Microsoft developers who deployed Tay and, subsequently, removed it from the Web did not consider the risks to those merely viewing the posts when Tay unleashed its hateful speech. Should we ascribe responsibility to Microsoft developers for this incident? Tay was quite specific as to its mission to change its behavior in real time according to what it learned from user responses and other information on Twitter and to publicly display its Twitter responses. Risk assessment did not address the larger sociotechnical context of Twitter and its users; nor did it simulate in a closed environment what would happen if learning produced offensive speech. There were not adequate controls for the

downstream users as, apparently, Tay learned from all users after deployment.¹⁹ This analysis points strongly in favor of looking at a broader context when designing AI applications. At this point, software developers may not even realize that they need to consider the impact the AI may have on those whose data are used to train it, on those who are subject to its output, or on society as a whole. As previously mentioned by Bornstein and Howard, the “dual use” of AI needs serious consideration in the ethical analysis. There is a shared ethical responsibility among those who develop the AI, those who train it, and those who deploy it as they position it for use in a global, sociotechnical environment.

At the beginning of this section, we were reminded that the Tuskegee Syphilis Study was a study that should not have been conducted. Given the experimental nature of AI research, development, training, and deployment, it is safe to suppose that not all AI ought to be developed. Each project requires that those involved give serious consideration to the question of whether the project even ought to be. Any project must not move forward until there is a clear and convincing argument that humanity will be better off with the AI.

Issue 3. In the future, should we be willing to consider some AI artifacts to be persons? (Keith Miller)

In Greek mythology, Pygmalion sculpts an ivory statue depicting a woman. He falls in love with the statue, and Aphrodite grants his prayer to bring the statue to life. The Greeks did not corner the market on the idea of animating nonliving matter. Jewish folklore includes the golem, and Mary Shelley wrote of Frankenstein’s experiment with reanimation. The subtitle of Shelley’s book is “The Modern Prometheus.” That brings us full circle to the Greeks since Prometheus is the Titan god of fire credited with creating the human race from clay.

Modern AI promises us a transformation similar to these stories. Building with silicon, electricity, and clever engineering, AI has already delivered artifacts that can converse with us, learn with us, and walk with us. After decades of overhyped speculation, AI is now presenting us with an amazing array of functioning machines and tantalizing suggestions of amazing advances in the near future. Furthermore, unlike the older fictional stories, we can purchase many of these creations right now, and we are promised increasingly

human-like devices in the foreseeable future.

Goaded by recent developments in AI, I think society as a whole and scholars from many disciplines are starting to engage seriously in an increasingly important question, which I will label “Question 0.”

- › Question 0: Should we consider some AI artifacts, either now or in the future, as persons?

I do not think questions about this issue will be settled quickly or easily; I do think it is vital that we focus on this issue with urgency. I want to distinguish this question from two related, but importantly different, questions:

- › Alternative Question 1: Can a future AI become sufficiently person-like in its behavior and appearance so that we will not be able to easily distinguish between it and a human being?
- › Alternative Question 2: Will society consider some AI artifacts, either now or in the future, as persons?

Both of these alternative questions are interesting, timely, and already being discussed in the literature. Importantly, both of these alternative questions might be answered empirically; if machines exist that routinely pass as human beings (that is, they often pass a physical Turing test), then the answer to Alternative Question 1 is, in my opinion, “yes.” If when we look about and no such machines exist, then the cautious answer to Alternative Question 1 is “at least not yet.”

Similarly, if we look about and observe that either many or most of our fellow humans treat sufficiently sophisticated AI artifacts as if they were persons, then the answer to Alternative Question 2 is “yes.” If we do not see that happening, then the cautious answer to Alternative Question 2 is “at least not yet.”

Both the alternative questions are questions of observation and description. I think the more important Question 0 is normative. Should we consider some AI artifacts, either now or in the future, as persons? At the heart of this question is a recognition that the designation of personhood is a societal choice, not a scientific classification. That societal choice could be expressed as law, custom, or regulations. Without some societal decision, we cannot devise a definitive physical or

behavioral test for personhood: we who are already considered persons have to come to an agreement (probably not a universal consensus) on what other entities, if any, we will allow into our “personhood club.”

Except in the case of cloning humans (for example, the replicants in *Blade Runner*), our “candidates” for personhood have no claim to be of our species. Humans are carbon based, and at least all common examples of AI today are silicon based. For some humans, that may make the answer to Question 0 easy. Only humans can be persons; AI artifacts are not humans, so they should not be considered persons.

But many people think that the easy rejection of AI personhood is too glib. And once you reject the idea that only humans can be persons, the question becomes more nuanced. There are existing potential answers to this question in the literature. In a provocatively titled article, Joanna Bryson²¹ stakes out a strong position: “robots should be slaves.” Bryson’s abstract begins: “Robots should not be described as persons, nor given legal nor moral responsibility for their actions.” Bryson goes on to argue that we should not create machines that would seriously compete for that designation.

It seems a reasonable question to ask a corollary to Question 0: Why should we build machines that appear to be persons? Is it merely because we can? Is it because such machines give us functionality that is difficult or impossible to get in other ways? These questions are, I think, important. They have been asked for years, but they are not as often answered, especially by the people and organizations developing and funding the development of these AI artifacts.

In the 10 years since Bryson’s publication, AI researchers have not appreciably slowed their work to make machines that increasingly closely resemble humans both in their behavior and their appearance. I anticipate that Question 0 and its corollaries will become far more contentious in the next few years. 🤖

REFERENCES

1. N. Wiener, *Cybernetics or Control and Communication in the Animal and the Machine*. Cambridge, MA: MIT Press, 1948.
2. T. W. Bynum, “A very short history of computer ethics,” *APA Newsltt. Philos. Comput.*, vol. 99, no. 2, p. 2, 2000.
3. N. Wiener, *The Human Use of Human Beings, Cybernetics*

- and Society*. Boston: Houghton Mifflin, 1950.
4. A. Kaplan and M. Haenlein, "Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence," *Bus. Horiz.*, vol. 62, no. 1, pp. 15–25, 2019. doi: 10.1016/j.bushor.2018.08.004.
 5. J. McCarthy, M. Minsky, N. Rochester, and C. E. Shannon, "A proposal for the Dartmouth summer research project on artificial intelligence," Stanford Univ., CA, 1955. [Online]. Available: <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>
 6. T. W. Bynum, "Artificial intelligence, biology, and intentional states in computers and ethics," *Metaphilosophy*, vol. 16, no. 4, pp. 355–377, 1985. doi: 10.1111/j.1467-9973.1985.tb00183.x
 7. K. Čapek, *R.U.R.: Rossum's Universal Robots (P. Selver and N. Playfair, Transl.)* 1920. [Online]. Available: preprints.readingroo.ms/RUR/rur.pdf
 8. I. Asimov, "Runaround," *Astounding Sci. Fiction*, vol. 29, no. 1, pp. 94–103, 1942.
 9. "List of artificial intelligence films," Wikipedia.org, 2020. [Online]. Available: https://en.wikipedia.org/wiki/List_of_artificial_intelligence_films
 10. "List of fictional robots and androids: Television films and series," Wikipedia.org, 2020. [Online]. Available: https://en.wikipedia.org/wiki/List_of_fictional_robots_and_androids#Television_films_and_series
 11. K. Miller, M. Wolf, and F. Grodzinsky, "This 'ethical trap' is for roboticists, not robots: On the issue of artificial agent ethical decision-making," *Sci. Eng. Ethics*, vol. 23, no. 2, pp. 389–401, 2016. doi: 10.1007/s11948-016-9785-y.
 12. A. Howard and J. Borenstein, "The ugly truth about ourselves and our robot creations: The problem of bias and social inequity," *Sci. Eng. Ethics*, vol. 24, no. 5, pp. 1521–1536, 2018. doi: 10.1007/s11948-017-9975-2.
 13. R. Benjamin, *Race After Technology*. Polity, 2019.
 14. M. Brundage et al., "The malicious use of artificial intelligence: Forecasting, prevention, and mitigation," Future of Humanity Institute, Centre for the Study of Existential Risk, Center for a New American Security, Electronic Frontier Foundation, OpenAI. [Online]. Available: <http://arxiv.org/abs/1802.07228>
 15. S. M. West, M. Whittaker, and K. Crawford, "Discriminating systems: Gender, race and power in AI," AI Now Institute, New York, 2019. [Online]. Available: <https://ainowinstitute.org/discriminatingystems.html>
 16. J. Buolamwini and T. Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," *Proc. Mach. Learn. Res.* vol. 81, pp. 77–91, 2018. [Online]. Available: <http://proceedings.mlr.press/v81/buolamwini18a.html>
 17. Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, "Dissecting racial bias in an algorithm used to manage the health of populations," *Science*, vol. 366, no. 6464, pp. 447–453, 2019. doi: 10.1126/science.aax2342.
 18. A. Kak, "The Global South is everywhere, but also always somewhere: National Policy Narratives and AI Justice," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc. (AIES '20)*, New York, 2000, pp. 307–312. doi: 10.1145/3375627.3375859.
 19. M. Wolf, K. Miller, and F. Grodzinsky, "Why we should have seen that coming: Comments on Microsoft's Tay 'experiment,' and wider implications," *ORBIT J.*, vol. 1, no. 2, pp. 1–12, 2017. doi: 10.29297/orbit.v1i2.49.
 20. F. S. Grodzinsky, K. Miller, and M. J. Wolf, "The ethics of designing artificial agents," *Ethics Inform. Technol.*, vol. 10, nos. 2–3, pp. 2–3, 2008. doi: 10.1007/s10676-008-9163-9.
 21. J. Bryson, "Robots should be slaves," in *Close Engagements With Artificial Companions: Key Social, Psychological, Ethical and Design Issues*, Y. Wilks, Ed. Philadelphia: John Benjamins, pp. 63–74, 2010.
- JASON BORENSTEIN** is a director at the Graduate Research Ethics Programs, School of Public Policy and Office of Graduate Studies, Georgia Institute of Technology, Atlanta, Georgia, USA. Contact him at borenstein@gatech.edu.
- FRANCES S. GRODZINSKY** is a professor emerita at the School of Computer Science and Engineering, Sacred Heart University, Fairfield, Connecticut, USA. Contact her at grodzinskyf@yahoo.com.
- AYANNA HOWARD** is a Linda J. and Mark C. Smith professor and chair at the School of Interactive Computing, Georgia Institute of Technology, Atlanta, Georgia, USA. Contact her at ah260@gatech.edu.
- KEITH W. MILLER** is a professor with the Department of Computer Science and the College of Education, University of Missouri–St. Louis, St. Louis, Missouri, USA. Contact him at millerkei@umsl.edu.
- MARTY J. WOLF** is a professor of computer science at Bemidji State University, Bemidji, Minnesota, USA. Contact him at Marty.Wolf@bemidjistate.edu.



Call for Articles

IEEE Software seeks practical, readable articles that will appeal to experts and nonexperts alike. The magazine aims to deliver reliable, useful, leading-edge information to software developers, engineers, and managers to help them stay on top of rapid technology change. Topics include requirements, design, construction, tools, project management, process improvement, maintenance, testing, education and training, quality, standards, and more.

Author guidelines:
www.computer.org/software/author.htm
 Further details: software@computer.org
www.computer.org/software



VANDERBILT UNIVERSITY



School of Engineering

TENURE-TRACK FACULTY POSITIONS IN COMPUTER SCIENCE (Destination Vanderbilt)

The Department of Computer Science launched in 2020 a multi-year faculty recruitment and hiring process for 20+ tenure-track positions at all ranks above normal hiring patterns, with preference at early-career appointments. In Year 1, the department welcomed eight new faculty members. In Year 2, the initiative will support at least eight new faculty positions starting in AY 2022-2023. Destination-CS is part of the university's recently launched Destination Vanderbilt, a \$100 million initiative to recruit new faculty.

The hiring initiative aspires to propel the Vanderbilt CS department to one of the leading academic programs nationally and beyond. Successful candidates are expected to teach at the undergraduate and graduate levels and to develop and grow vigorous programs of externally funded research. We seek exceptional candidates in broadly defined areas of computer science that enhance our research strengths in areas that align with the following investment and growth priorities of the Vanderbilt University School of Engineering:

1. **Cybersecurity and Resilience**
2. **Autonomous and Intelligent Human-AI-Machine Systems and Urban Environments**
3. **Computing and AI for Health, Medicine, and Surgery**
4. **Design of Next Generation Systems, Structures, Materials, and Manufacturing**

Ranked #14 nationally, Vanderbilt University is a private, internationally recognized research university located on 330 park-like acres 1.5 miles from downtown Nashville. Its 10 distinct schools share a single cohesive campus that values collaboration. The university enrolls over 13,500 undergraduate, graduate, and professional students, including 36% minority students and over 1,100 students from 84 countries. The School of Engineering is on a strong upward trajectory in national and international prominence, and has built infrastructure to support a significant expansion in faculty size. In the rankings of graduate engineering programs by U.S. News, the school ranks in the top 20 private, research-extensive engineering schools. Five-year average T/Tk faculty funding in the formerly combined EECS department is above \$800k per year per person. Nearly all junior faculty members hired during the past 15 years have received prestigious young investigator awards, such as NSF CAREER and DARPA CSSG.

With a metro population of over two million people, Nashville's top industries are trade, transportation and utilities; education and health services; professional and business services; government; and leisure and hospitality. Other industries include manufacturing, financial activities, construction, and information. Long known as a hub for health care and music, Nashville is a technology center with a considerable pool of health care, AI, and defense-related jobs available. The city recently has experienced an influx of major office openings by some of the largest global tech companies and prime Silicon Valley startups.

Vanderbilt University has a strong institutional commitment to recruiting and retaining an academically and culturally diverse community of faculty. Minorities, women, individuals with disabilities, and members of other underrepresented groups, in particular, are encouraged to apply. Vanderbilt is an Equal Opportunity/Affirmative Action employer.

Vanderbilt University has made the safety of our students, faculty and staff, and our surrounding communities a top priority. As part of that commitment, the University recently announced that students, faculty, and staff, are required to be vaccinated against COVID. As a prospective and/or a new employee at Vanderbilt, you will be required to comply with the University's vaccination protocol. Effective, August 1, 2021, proof of full vaccination or an approved accommodation will be required in order to work at Vanderbilt University.

Applications should be submitted at: <http://apply.interfolio.com/94225>. For more information, please visit: <http://vu.edu/destination-cs>. Applications will be reviewed on a rolling basis beginning December 1, 2021. Interviews begin around January 1, 2022. For full consideration, application materials must be received by January 31, 2022.

IEEE COMPUTER SOCIETY D&I FUND

Drive Diversity & Inclusion in Computing



Supporting projects and programs that positively impact diversity, equity, and inclusion throughout the computing community.



Do you have a great idea for new programs that will positively impact diversity, equity, and inclusion throughout the computing community?

The IEEE Computer Society Diversity & Inclusion Committee seeks proposals for projects, programs, and events that further its mission. New programs that deliver education, outreach, and support, including, but not limited to, mentoring programs at conferences, panel discussions, and webinars, are welcomed.

Help propel the Computer Society's D&I programs—submit a proposal today!

<https://bit.ly/CS-Diversity-CFP>



Donations to the IEEE Computer Society D&I Fund are welcome!



IEEE Foundation

IEEE Internet Computing


IEEE Internet Computing delivers novel content from academic and industry experts on the latest developments and key trends in Internet technologies and applications.

Written by and for both users and developers, the bimonthly magazine covers a wide range of topics, including:

- Applications
- Architectures
- Big data analytics
- Cloud and edge computing
- Information management
- Middleware
- Security and privacy
- Standards
- And much more

In addition to peer-reviewed articles, *IEEE Internet Computing* features industry reports, surveys, tutorials, columns, and news.

www.computer.org/internet



Join the IEEE Computer Society
for subscription discounts today!

www.computer.org/product/magazines/internet-computing

75 YEARS
IEEE
COMPUTER
SOCIETY



IEEE

SECURITY & PRIVACY

IEEE Security & Privacy is a bimonthly magazine communicating advances in security, privacy, and dependability in a way that is useful to a broad section of the professional community.

The magazine provides articles with both a practical and research bent by the top thinkers in the field of security and privacy, along with case studies, surveys, tutorials, columns, and in-depth interviews. Topics include:

- Internet, software, hardware, and systems security
- Legal and ethical issues and privacy concerns
- Privacy-enhancing technologies
- Data analytics for security and privacy
- Usable security
- Integrated security design methods
- Security of critical infrastructures
- Pedagogical and curricular issues in security education
- Security issues in wireless and mobile networks
- Real-world cryptography
- Emerging technologies, operational resilience, and edge computing
- Cybercrime and forensics, and much more


www.computer.org/security



Join the IEEE Computer Society for subscription discounts today!

www.computer.org/product/magazines/security-and-privacy

75 YEARS
IEEE COMPUTER SOCIETY



Conference Calendar

IEEE Computer Society conferences are valuable forums for learning on broad and dynamically shifting topics from within the computing profession. With over 200 conferences featuring leading experts and thought leaders, we have an event that is right for you. Questions? Contact conferences@computer.org.

NOVEMBER

2 November

- ICNP (IEEE Int'l Conf. on Network Protocols), Dallas, USA

4 November

- SLIP (ACM/IEEE Int'l Workshop on System-Level Interconnect Prediction), virtual

6 November

- SmartCloud (IEEE Int'l Conf. on Smart Cloud), Newark, USA

7 November

- IISWC (IEEE Int'l Symposium on Workload Characterization), virtual

12 November

- ICEBE (IEEE Int'l Conf. on e-Business Eng.), Guangzhou, China

14 November

- SC (Int'l Conf. for High-Performance Computing, Networking, Storage, and Analysis), St. Louis, USA

15 November

- ASE (IEEE/ACM Int'l Conf. on Automated Software Eng.), Melbourne, Australia
- BigMM (IEEE Int'l Conf. on Multimedia Big Data), Taichung, Taiwan

16 November

- AVSS (IEEE Int'l Conf. on Advanced Video and

Signal-Based Surveillance), virtual

17 November

- ICMU (Int'l Conf. on Mobile Computing and Ubiquitous Networking), Tokyo, Japan

19 November

- ITME (Int'l Conf. on IT in Medicine and Education), Wuyishan, China

22 November

- ATS (IEEE Asian Test Symposium), virtual

23 November

- IMITEC (Int'l Multidisciplinary Information Technology and Eng. Conf.), Windhoek, Namibia

24 November

- SNPD (IEEE/ACIS Int'l Conf. on Software Eng., Artificial Intelligence, Networking, and Parallel/Distributed Computing), Taichung, Taiwan

29 November

- ISM (IEEE Int'l Symposium on Multimedia), virtual

30 November

- ICRC (IEEE Int'l Conf. on Rebooting Computing), virtual

DECEMBER

1 December

- PRDC (IEEE Pacific Rim Int'l Symposium on Dependable

Computing), virtual

5 December

- HPBD&IS (Int'l Conf. on High-Performance Big Data and Intelligent Systems), Macau, China

6 December

- BDCAT (IEEE Int'l Conf. on Big Data Computing, Applications, and Technologies), Leicester, UK
- Blockchain (IEEE Int'l Conf. on Blockchain), virtual
- UCC (IEEE Int'l Conf. on Utility and Cloud Computing), Leicester, UK

7 December

- CSASE (Int'l Conf. on Computer Science and Software Eng.), Duhok, Iraq
- ICDM (IEEE Int'l Conf. on Data Mining), virtual
- RTSS (IEEE Real-Time Systems Symposium), Dortmund, Germany

9 December

- BIBM (IEEE Int'l Conf. on Bioinformatics and Biomedicine), virtual

12 December

- HOST (IEEE Int'l Symposium on Hardware-Oriented Security and Trust), Washington, DC, USA



13 December

- CIC (IEEE Int'l Conf. on Collaboration and Internet Computing), virtual
- CogMI (IEEE Int'l Conf. on Cognitive Machine Intelligence), virtual
- MSN (Int'l Conf. on Mobility, Sensing, and Networking), Exeter, UK
- PST (Int'l Conf. on Privacy, Security, and Trust), virtual
- TPS-ISA (IEEE Int'l Conf. on Trust, Privacy, and Security in Intelligent Systems and Applications), virtual

14 December

- SEC (ACM/IEEE Symposium on Edge Computing), San Jose, USA

15 December

- BigData (IEEE Int'l Conf. on Big Data), virtual
- FG (IEEE Int'l Conf. on Automatic Face and Gesture Recognition), virtual

17 December

- HiPC (IEEE Int'l Conf. on High-Performance Computing, Data, and Analytics), Bengaluru, India

18 December

- STI (Int'l Conf. on Sustainable Technologies for Industry 4.0), virtual

20 December

- MCSoc (IEEE Int'l Symposium on Embedded Multicore/Many-Core Systems-on-Chip), Singapore

2022

JANUARY

12 January

- ICOIN (Int'l Conf. on Information Networking), Bangkok, Thailand

26 January

- ICSC (IEEE Int'l Conf. on Semantic Computing), Laguna Hills, USA

FEBRUARY

7 February

- FOCS (IEEE Symposium on Foundations of Computer Science), Denver, USA

12 February

- CGO (Int'l Symposium on Code Generation and Optimization), Seoul, South Korea
- HPCA (IEEE Int'l Symposium on High-Performance Computer Architecture), Seoul, South Korea

MARCH

12 March

- ICSA (IEEE Int'l Conf. on Software Architecture), Honolulu, USA
- VR (IEEE Conf. on Virtual Reality and 3D User Interfaces), Christchurch, New Zealand

15 March

- SANER (IEEE Int'l Conf. on Software Analysis, Evolution, and Reengineering), Honolulu, USA

21 March

- PerCom (IEEE Int'l Conf. on Pervasive Computing and Communications), Pisa, Italy

APRIL

11 April

- PacificVis (IEEE Pacific Visualization Symposium), Tsukuba, Japan

MAY

9 May

- ICDE (IEEE Int'l Conf. on Data Eng.), virtual

21 May

- ICSE (IEEE/ACM Int'l Conf. on Software Eng.), Pittsburgh, USA

30 May

- IPDPS (IEEE Int'l Parallel & Distributed Processing Symposium), Lyon, France

Learn more
about IEEE
Computer Society
conferences

computer.org/conferences

Get Published in the *IEEE Open Journal of the Computer Society*

Submit a paper today to the premier open access journal in computing and information technology.

Your research will benefit from the IEEE marketing launch and 5 million unique monthly users of the IEEE *Xplore*® Digital Library. Plus, this journal is fully open and compliant with funder mandates, including Plan S.

Submit your paper today!

Visit www.computer.org/oj to learn more.

