COMPUTING edge

- Internet of Things
- Ethics
- Machine Learning
- Quantum Computing

MAY 2020

www.computer.org

IEEE COMPUTER SOCIETY

IEEE

COMPUTING edge

IEEE COMPUTER SOCIETY computer.org

STAFF

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.

IEEE Computer Society Magazine Editors in Chief

MAY 2020 • VOLUME 6 • NUMBER 5

COMPUTING edge

8 The Internet of Things Grows Artificial Intelligence and Data Sciences

30 Mining Insights from Visual Assets

45 The U.S. National Quantum Initiative

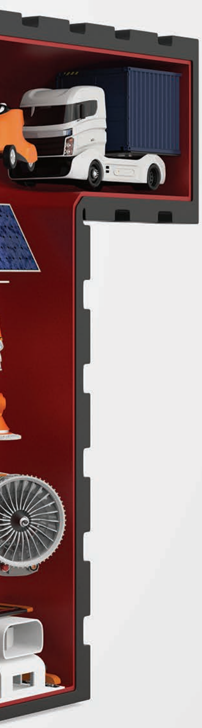

Magazine Roundup

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

The Promise of Interactive Shared Augmented Reality

Augmented reality (AR) is a game-changing technology that lets users see things they cannot otherwise see. Among other applications, shared reality could be used to improve the safety of traffic systems. Despite current limitations, the future is bright for interactive shared AR. Read more in the January 2020 issue of *Computer*.

Computing in Science & Engineering

Assessing the Impact of Outreach through Software Citation for Community Software in Geodynamics

The Computational Infrastructure for Geodynamics is a community of software users and user-developers who model physical processes in the Earth and planetary interiors. From 2010 to 2018, the community of researchers published upward of 638 peer-reviewed papers in more than 124 venues. The authors of this article from the January/February 2020 issue of *Computing in Science & Engineering* analyzed this corpus of publications to understand the impact of CIG workshops and tutorials, measured through software citation.

IEEE Annals of the History of Computing

Coping with the "American Giants"

At the beginning of the 1960s, several Western European computer companies faced financial issues and pressure from US competitors. A series of negotiations attempted to create a consortium of European manufacturers, and while some of these had a positive outcome, in general, they did not succeed; by 1964, IBM had more than a 60% share of the European computer market and General Electric acquired two of the most prominent companies in Europe. This article from the October–December 2019 issue of *IEEE Annals of the History of Computing* examines negotiations during the years 1962–1964, focusing on contacts between UK manufacturers English Electric, LEO Computers, and ICT, and Italian Olivetti, French Bull, and German Siemens.

IEEE Computer Graphics and Applications

Towards Placental Surface Vasculature Exploration in Virtual Reality

The authors of this article from the January/February 2020 issue of *IEEE Computer Graphics and Applications* present a case study evaluating the potential for interactively identifying placental surface blood vessels using magnetic resonance imaging (MRI) scans in virtual-reality (VR) environments. They visualized the MRI data using direct volume rendering in a high-fidelity CAVE-like VR system, allowing medical professionals to identify relevant placental vessels directly from volume visualizations in the VR system, without prior vessel segmentation. Participants were able to trace most of the observable vascular structure, and consistently identified blood vessels down to diameters of 1 mm, an important requirement in diagnosing vascular diseases.

IEEE Intelligent Systems

Keyword Generation for Sponsored Search Advertising: Balancing Coverage and Relevance

Automatically generating a pool of keywords used by potential consumers is a challenging issue for advertisers in sponsored search advertising (SSA). Such a keyword pool serves as the base for market research and determines the feasible space of consequent keyword-related decisions. This article from the September/October 2019 issue of *IEEE Intelligent Systems* presents a novel method for keyword generation with Wikipedia as a corpus of the source text (WIKG). Starting with a few seed keywords, the WIKG supports flexible keywords generation by taking advantage of Wikipedia's rich link structure to construct a graph of entry articles in an iterative way.

IEEE Internet Computing

Contextual, Behavioral, and Biometric Signatures for Continuous Authentication

Continuous authentication in the Mobile Internet of Things should be based as broadly as possible, since a wide range of factors continuously reveal unexpected correlations. Such factors may include captured events, continuous time series, and derived behavioral features. All these factors have been shown to correlate with the actual user identity, often in surprising combinations. More and more sensors are being deployed in autonomous devices, smart environments, and vehicles, enabling even further behavioral and contextual data to be analyzed. The pegs of this continuous authentication "big tent" are moving out further than ever before, bringing it closer to practical uses in our everyday lives. Read more in the September/October 2019 issue of *IEEE Internet Computing*.

IEEE micro

A Bunch-of-Wires (BoW) Interface for Interchiplet Communication

Multichiplet system-in-package designs have recently received a lot of attention as a mechanism to combat high SoC design costs and to economically manufacture large ASICs. These designs require low-power area-efficient off-die on-package die-to-die communication. Current technologies either extend on-die high-wire count buses using silicon interposers or off-package serial buses. The former approach leads to expensive packaging. The latter leads to complex and high-power designs. The authors of this article from the January/February 2020 issue of *IEEE Micro* propose a simple bunch-of-wires interface that combines ease of development with low-cost packaging techniques.

IEEE MultiMedia

Residual-Based Post-Processing for HEVC

The authors of this article from the October–December 2019 issue of *IEEE MultiMedia* propose a residual-based post-processing for high-efficiency video coding (HEVC). Based on the proposed network, residual-based video restoration network (residual-VRN), the decoded image quality has improved significantly. The authors experimentally verified the superiority of their method.

IEEE pervasive COMPUTING
MOBILE SYSTEMS | UBIQUITOUS COMPUTING | INTERNET OF THINGS

Printing Wearable Devices in 2D and 3D: An Overview of Mechanical and Electronic Digital Co-design

Multi-process additive manufacturing (AM) offers system designers new, exciting computational tools to rapidly realize smart wearable sensing devices in 2D and 3D shapes. The authors of this article from the October–December 2019 issue of *IEEE Pervasive Computing* guide readers through the novel development and fabrication process based on a digital co-design framework and highlight AM techniques, functional materials, and assembly procedures for designing wearables as flexible and stretchable on-skin patches, e-textiles, and smart accessories for everyday use.

IEEE SECURITY & PRIVACY

Toward a Data-Driven Society: A Technological Perspective on the Development of Cybersecurity and Data-Protection Policies

A data-driven society requires a common regulatory umbrella to provide a harmonized vision of cybersecurity. The authors of this article from the January/February 2020 issue of *IEEE Security & Privacy* describe cybersecurity policies and joint initiatives in the European Union and give insights about the need to align ongoing technological advances with such regulatory efforts.

IEEE Software

Taking the Middle Path: Learning about Security through Online Social Interaction

Integrating security into software development involves more than learning principles or applying techniques. Security can be integrated into software development practice by following a middle path, through which developers draw together knowledge received through training and software development techniques. Read more in the January/February 2020 issue of *IEEE Software*.

IT Professional

The Adoption of Cognitive Computing Technology in China

China is becoming a global leader in the adoption of cognitive technology. One of the major benefits that cognitive technology promises to bring to China is the ability to advance the social, economic, and environmental development of the large rural population. A controversial application of the technology is security surveillance; while the Chinese government claims that cognitive-powered surveillance helps reduce crime and create a safe environment for its citizens, others view the practice as technocratic governance and an abuse of civil liberties. This article from the November/December 2019 issue of *IT Professional* reviews the adoption and development of cognitive technology in China and discusses the impacts, concerns, and challenges of the technology for the country.

Leveraging the Internet of Things

Internet of Things (IoT) devices are common in today's homes and include everything from thermostats and light bulbs to smart speakers and fitness trackers. However, the IoT has enormous potential beyond these domestic applications. Organizations in most industries could leverage the IoT to improve their processes and inform their decision-making. This issue of *ComputingEdge* examines how two industries—manufacturing and education—can benefit from implementing IoT systems.

IT Professional's "The Internet of Things Grows Artificial Intelligence and Data Sciences" recommends that manufacturers utilize IoT data to improve operational efficiency, safety, and flexibility—and, ultimately, to stay competitive. The authors emphasize the importance of using artificial intelligence-based predictive analytics on IoT data. In "'Smarter' Education," also from *IT Professional*, the authors discuss the ways IoT could be used to support schools, including attendance management and equipment tracking.

Ethics is important in IoT system design, especially as more and more industries adopt IoT technology. *IEEE Pervasive Computing*'s "Carousel Kittens: The Case for a Value-Based IoT" argues for requirements engineering that takes into account the dignity, privacy, and freedom of users and data subjects. Similarly, *IEEE Software*'s "Ethics Is a Software Design Concern" focuses on ethical design in software development.

Machine learning is another technology that touches many industries and has many applications. "Mining Insights from Visual Assets: A Case Study," from *IEEE Computer Graphics and Applications*, examines one company's machine learning-based visual analytics platform. "Machine Learning for Internet Congestion Control: Techniques and Challenges," from *IEEE Internet Computing*, highlights machine-learning methods for mitigating network congestion.

This *ComputingEdge* issue concludes with two articles from *Computer* about organizations at the forefront of quantum computing: the IEEE and the US government. "A Role for IEEE in Quantum Computing" explains how IEEE initiatives—such as Rebooting Computing and the International Roadmap for Devices and Systems—will help improve quantum technology. "The US National Quantum Initiative" describes efforts by the US government to accelerate quantum research and development.

EDITORS: Phillip A. Laplante, plaplante@psu.edu
Ben Amaba, baamaba@us.ibm.com

This article originally appeared in IT Professional vol. 21, no. 3, 2019

DEPARTMENT: INTERNET OF THINGS

The Internet of Things Grows Artificial Intelligence and Data Sciences

Charla Stracener, Quentin Samelson, Joe Mackie, and Mitsuko Ihaza, *IBM*

The combination of growing competition and inexpensive connectivity has made the Internet of things (IoT) an ongoing topic in manufacturing. Sensors, devices, and machines all connected via the Internet are the "things" in IoT. The flood of IoT data can provide the information needed to stay competitive, by applying analytics and artificial intelligence—with the goal to improve operations' efficiency, safety, and flexibility.

The Internet of things (IoT) is at times referenced in part as *cyber-physical systems, Industry 4.0, the fourth industrial revolution, the industrial Internet, big analog data solutions, smarter planet, intelligent systems, and Digital Twin*. Regardless of how it is described, IoT has captured the world's imagination. Connecting computers, smart phones, sensors, appliances, machinery, vehicles, utilities, and a host of other elements into a reliable efficient "system of systems" could be the greatest engineering accomplishment of our generation—as well as the largest engineering challenge we have ever faced.

By the end of 2013, IoT and RPA deployments already included an estimated 20 billion connected devices (out of about 187 billion connectable devices). That number is predicted grow to 30 billion by 2020. These IoT devices come together as sensors, intelligence, actuators, and power into robotic process automation. IoT and RPA technology applications are growing at a rapid rate from the chat bots with which we interact in our everyday lives to manufacturing lines that are not only fully automated but also self-report performance and automatically request assistance when an anomaly is detected.

When we evaluate IoT in the industrial, transportation, and utility sectors, worldwide infrastructure is undergoing a user-centered, software-driven, digital transformation—a change from the material and mechanical innovations that have driven innovation in the past. For the United States, this transformation is urgently needed. Ripe for change, the United States infrastructure earned a barely passing grade of D+ from the American Society of Civil Engineers for its "aging electrical grid and pipeline distribution systems, some of which originated in the 1880s."

Despite systems and technologies spanning over multiple decades, we expect these heterogeneous systems to be stable and safe and to provide services in a secure manner. Ultimately, the IoT devices deployed in these environments will couple real-time analyses with machine-to-machine, machine-to-infrastructure, and user-to-machine communication so that they can adapt continually to changing circumstances. The scale of connectivity is unprecedented. Consider that a "mere" 20 billion devices yield 400 quintillion—4×10^{20}—potential communicating pairs.

Although even companies with rather rudimentary manufacturing operations often have IoT sensors in at least some of their machines, many (if not most) companies have not yet capitalized on the information that can be gleaned from the data produced by those sensors. We will especially call attention to the opportunity to monitor overall equipment effectiveness (OEE) using data from IoT sensors, to use that data to first become better at reacting to production slowdowns and stoppages, and then to develop the ability to predict and

DOI No. 10.1109/MITP.2019.2912729
Date of current version 21 May 2019.

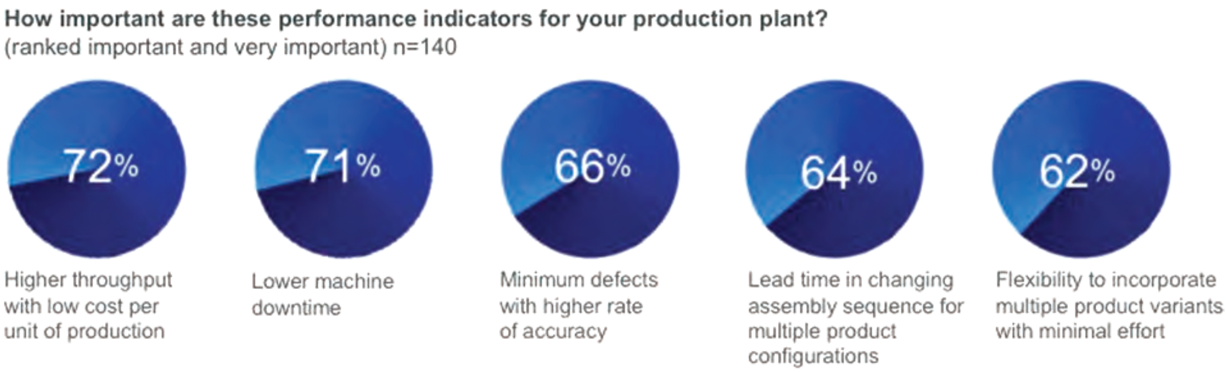

FIGURE 1. Electronics KPIs from the IBV's cognitive manufacturing study: *Why cognitive manufacturing matters.*

prevent issues on the manufacturing line—ultimately helping companies reach high OEE values.

With sensors everywhere, manufacturers have taken to instrumenting their plants with IoT to find the fast path to KPI improvements and delivery of more productive plant floor. When IBM's Institute of Business Value (IBV) asked electronics manufacturers which performance indicators were most important for their plants, COOs and manufacturing leaders answered with *throughput, uptime, and defect reduction*, in nearly equal measures (see Figure 1).

ARTIFICIAL INTELLIGENCE (AI) FUELED BY IOT

AI is making one of the biggest impacts on the digital community and our world. The field of AI research was first recognized at Dartmouth College in 1956. Engineers, scientists, and innovators led by Arthur Samuel of IBM, Allen Newell and Herbert Simon of Carnegie Mellon University, and John McCarthy and Marvin Minsky from the Massachusetts Institute of Technology, are widely regarded as pioneers in the field for their ground-breaking research. AI covers a large spectrum of machine intelligence from simple text-to-speech applications to autonomous vehicles and aircraft. Early efforts at AI struggled with the rudimentary computing technology available, but in the 1980ws, AI saw a resurgence with researchers attempting commercial applications with approaches based on deep learning and neural networks. Today, AI is appearing in our homes, cars, businesses, buildings, mobile phones, and appliances, from operator assistants to mapping software, to provide the most effective routes to your desired locations. In recent years, the infrastructure, network, and algorithms have become so robust that average Internet users, not just large technology institutions, have access to AI platforms. With hardware advancing and software becoming more user-friendly, AI is now available to a larger population for use and development, even to elementary school students' participating in hackathons over a weekend. AI is playing a major role in providing prescriptive analysis and recommendations to institutions, governments, and researchers on options and patterns that are not as apparent on the surface. AI has the ability to dig through mountains of data effectively and efficiently to provide answers and services in days that used to take years.

DATA PREPARATION

Before you can fully leverage data with AI and data science, the raw data must be cleansed and transformed into standardized higher quality data. The data preparation step is known to traditionally be the most time consuming. After surveying data scientists and engineers currently working in manufacturing in March of 2019 (From a survey of IBM clients working in the O&G industry.), not much seems to have changed in industry. When asked if they agreed that "I use 40% or more of my time in collecting, cleaning, and preparing data to model for data sciences and AI," 61% of those surveyed agreed. Additionally, when asked if they agreed that "three or more people work and share data and/or models in our group," 68% of those surveyed agreed. Automating this step with machine learning allows for less time spent on data preparation and more on the actual data analysis.

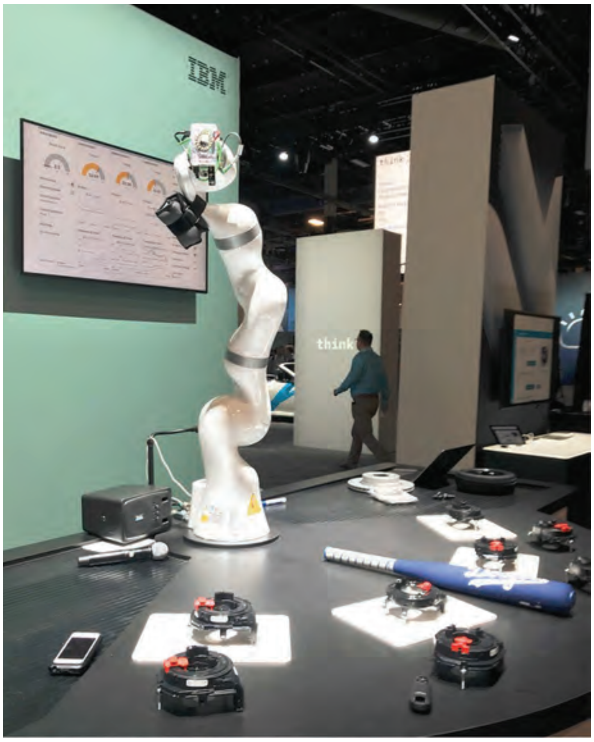

FIGURE 2. IBM Dallas Global Solution Center—Robotics, AI, IoT, and blockchain-integrated solution on an open platform combining multiple technologies.

IOT-ENABLING AI

The layered value of IoT and AI must meet those specific needs. IoT generates massive amounts of information, and AI helps to make sense of all this data, turning it into predictive findings and prescriptive recommendations. It is about helping industrial companies *reduce cost and downtime, maximize asset productivity, reduce risk of business disruption, improve quality, and streamline operations* to maximize asset operating efficiency by analyzing machines, predicting outages, handling equipment repairs, and automating equipment maintenance in real time to streamline global operations and keep critical assets operating at maximum efficiency. By surfacing dark data, AI allows you to work with new patterns of data and automate the equipment maintenance and quality management processes, delivering results instantly (see Figure 2).

Solutions need to be increasingly flexible. Many companies have found that collecting IoT data is easier than the challenging task of transitioning it into streamlined and productive decision support. That is what makes AI so crucial to the next generation of improved processes. The ability to combine structured and unstructured data along with other information sources—such as visual and audio files, operator and technician notes, and machine guidance manuals—requires not just machine learning but deeper intelligence. To that point, 49% of electronics executives in IBM's recent C-Suite study saw AI significantly impacting manufacturing processes in the next 2–3 years (see Figure 3).

Finally, to effectively manage resources, an organization must optimize multiple assets across systems and processes for greater asset availability, reliability, and performance. Embedding AI allows organizations to support planned and unplanned maintenance activities. This ranges from initial service requests to work order generation, through planning, to completion, and recording of actual results. It provides an integrated approach to manage discrete or complex assets, to help organizations overcome challenges of aging infrastructures, and in their siloed or disconnected systems. It is mission critical for the future of manufacturing to move beyond connecting systems for collaborating and optimizing. The knowledge an organization builds through its IoT assets and

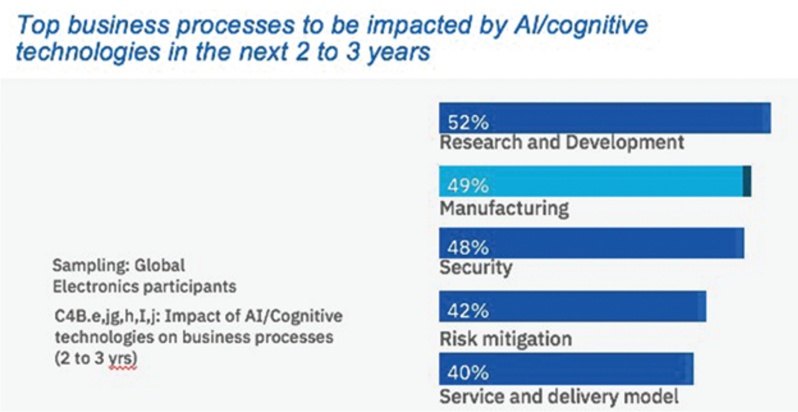

FIGURE 3. Top business processes to be impacted by AI/cognitive technologies: IBM's 19th Global C-Suite Study, The IBM Institute for Business Value ibm.biz /csuiteelectronics.

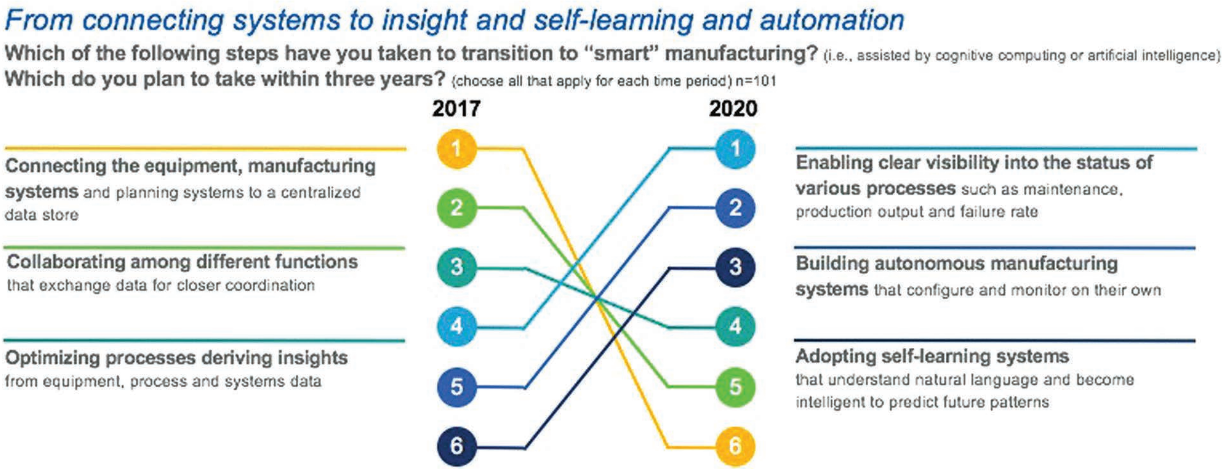

FIGURE 4. Why cognitive manufacturing matters in electronics: Activating the next generation of production: ibm.biz /cogmanufacture.

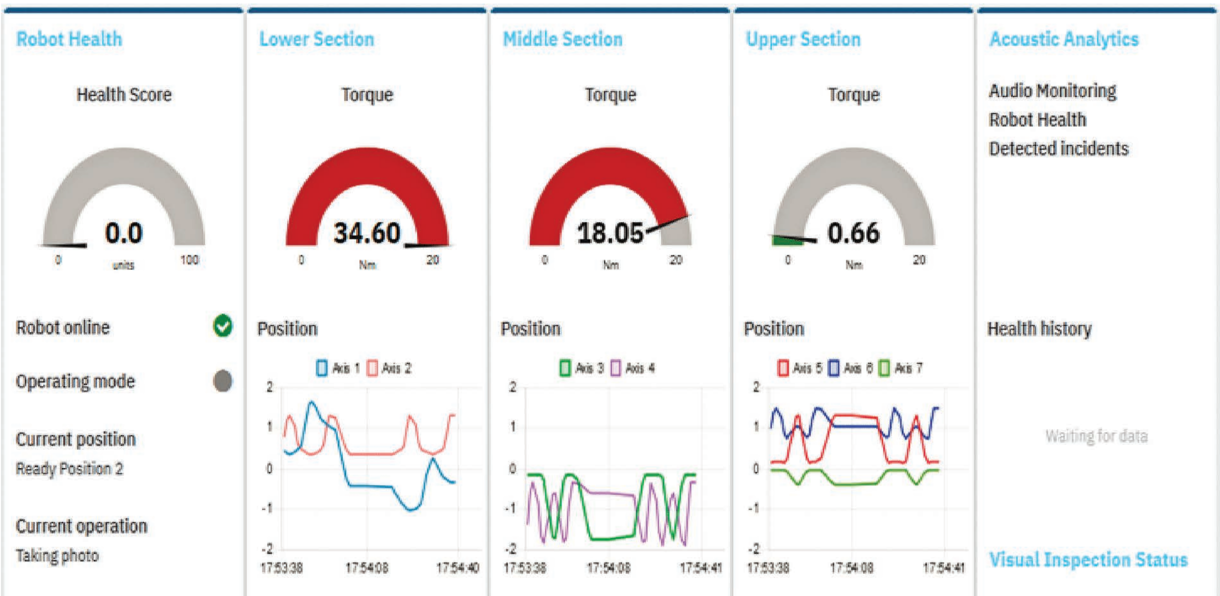

FIGURE 5. Robot health on torque and temperature (upper, middle, and lower elements), acoustical analysis, and visual recognition.

platforms will be the gateway for optimizing multiple processes simultaneously and heading for autonomous self-learning systems (see Figure 4).

During the manufacturing process, IoT data are collected from equipment asset sensors to manage and optimize the use of all assets to achieve greater asset availability, reliability, and performance (see Figure 5). By combining a very broad range of AI offerings, we can create a holistic view of the manufacturing process. This allows us to go beyond basic predictive maintenance and asset management, to assist with determinations around component and product defects and quality, which can affect the manufacturing process as well as a company's brand. We develop deeper understanding of interconnected quality issues that are not typically obvious and enable better interactions between business functions and more consistent outcomes across the entire organization. The system monitors production and asset health, improving manufacturing yield and asset optimization.

Follow this link to acquire additional information about *AI-powered manufacturing* and how the layered value of IoT and AI are helping industrial organizations get to cognitive manufacturing.

COMMON DATA SCIENCE AND ANALYTICS PLATFORM TO ENABLE AI SUCCESS WITH IOT

The oil and gas (O&G) industry has been working with IoT since before IoT was a thing. O&G organizations have had system collection and data acquisition systems since the 1960s to monitor and control remote equipment. These systems have evolved over time, and most O&G companies have multiple generations of systems connected to their networks. They also often have millions of devices of all different types and generations spread around the world. Often, these systems, and the other controls and IoT devices like them, process data in widely varying standards. This makes the role of a data administrator or engineer even more challenging. O&G organizations are not unlike other industrial businesses. They need to keep their workers safe while keeping costs down. This requires them to monitor the devices in real time. This often makes addressing alarms and alerts that come off these IoT devices challenging and even costly.

A data science team at IBM built a demonstration illustrating a real-life use case of IoT data and machine learning that got its inspiration from the Piper Alpha oil rig explosion. An alarm on a critical pump triggered a backup pump to start up. The damaged pump coming online caused an explosion in the rig. A total of 100 people died in the explosion that also destroyed the plant, costing its owners millions of dollars to repair it. One of the team's clients was struggling with a similar problem. The O&G client's control room operators were getting overwhelmed by alarms and alerts. The quantity of alerts and alarms made it challenging for them to understand which ones were the most critical to address. This problem was costly and time consuming, as staff needed to address each alarm and alert while looking for root causes. This organization's executives, like most in the industry, pressured their operations teams to increase the production in their plants and reduce all unnecessary down time. In this use case, the plant and control room operators needed to identify critical alarms that affect the production of petroleum products from pumps in their facilities. Potential for loss of life and the high cost of downtime means that operators need to predict when a pump might fail and where that pump is in the plant. Armed with this knowledge, plant operators can work to reduce their plant's down time and increase their production while avoiding unsafe situations. As anyone in the O&G industry can attest, downtime of a single-pump system can cost operators hundreds of thousands of dollars per hour. Preventing unnecessary costs is critical to the profitability of an operation and the company overall.

Leveraging predictive failure engines and a common data and analytics platform allows organizations to avoid emergency rush situations, which account for 4.2% of all reported injuries and fatalities in the O&G industry. The team built their demonstration to show their O&G client how to leverage their IoT data with machine-learning models to predict the pump failures in their plants and allow operators to take preventive action to avoid unnecessary downtime and rush situations. They built the demonstration in one week with an IoT dataset. Building off a common data science and analytics platform that is easily and quickly scaled ensured that this demonstration could be leveraged as a "minimum viable product"—the critical core of an actual application—and scaled into production with ease as needed. The unified platform that this team used is IBM Cloud Private for Data (ICPD). The ICPD platform aided the team in abstracting the data from the many devices around the plant allowing them to leave the data where it lived, thereby saving time in collecting, organizing, and preparing the data. By focusing less time on data preparation, the team was able to quickly move into the analytics phase and within a few days have a full-scale working machine-learning model built and tested. During this engagement, the team quickly connected devices, uploaded data, visualized that data, and infused machine-learning models into their systems. With this common foundation, the team quickly deployed the models at full scale and predicted equipment failure, even displaying where in the plant the failure was to occur. This outcome would allow the organization to schedule maintenance before a critical failure that would affect production.

If you would like to see this use case in action, the team shared with us a synopsis of the demonstration from their AI journey with their O&G client. This demonstration shows how industrial organizations of all kinds can tap into their IoT data leveraging a single data and analytics platform making it easier for them to connect to millions of data sources, virtualize critical data tables, visualize an initial analysis of the data, build initial regression models in Jupyter notebooks (A Jupyter notebook is an open document format based on JSON. See https://jupyter.org.), leverage machine learning to build and compare multiple models, and then deploy them. The demonstration (See https://ibm.ent.box.com/s/utnq8z1g3cbpdsk05otqevdjmyz960ac. (If readers experience any issues with this link, contact Joe Mackie at the e-mail address shown at the end of this paper.)) also shows how they used a web application, hosted in the cloud, to allow users to consume elements such as scored data, which are predictions of where and when failures would occur.

With the velocity and veracity of data facing IoT businesses today, it is critical that they leverage common platform architectures to help reduce friction and increase efficiency in tackling real challenges using their device data. A common platform, such as IBM Cloud Private for Data, makes it easy to collect, organize, analyze data, build machine-learning models, and collaborate among peers across the business to solve real business challenges. Building on a common platform, like this one, enables teams to leverage agile, deploy a minimally viable product, and continue to iterate in sprints over time to perfect and tweak their assumptions and models until the models can run themselves in a less supervised or even fully unsupervised way. When models can be let loose into the wild, with set limits, data scientists can rest easy knowing that the common platform will alert them when their models misbehave. Data scientists can thus focus on building new models to solve other business challenges. To adopt AI successfully, organizations need a solid information architecture built on seamless well-integrated common platform for lasting support.

PRACTICAL USES OF IOT

Does every implementation of IoT need to be as sophisticated as autonomous vehicles or semiconductor manufacturing? No—the data available via IoT devices can be used for mundane but meaningful purposes such as monitoring assembly lines or the conditions of a trailer carrying a load of fresh fruit.

Monitoring the stages of production on a manufacturing line can help companies improve their OEE by helping them react quickly to production line slowdowns and stoppages. Additionally, advanced analytics can recommend preventive action to avoid issues from ever occurring.

(OEE is a way to assess the overall performance of a manufacturing operation. A simple way to understand it is that OEE is the ratio of good products *actually completed* to the number of good products *planned*. OEE becomes vitally important to a factory when it begins to run out of useful capacity. Improving OEE from, say, 40% to 60% is like getting a 50% increase in capacity for free (A more exhaustive treatment of OEE and how IoT can be used to improve it can be found at https://medium.com/the-future-of-electronics/you-might-have-an-extra-production-line-s-hidden-in-your-factory-994491371430)).

Systems leveraging IoT data are inherently superior to manual monitoring mechanisms. Manual methods can be slow to react, losing minutes or hours of production time. Further, manual methods do not produce the data that will help you prevent future downtime, let alone predict it. We need data to be able to automatically detect and alert production staff to anomalies, and we need to collect and organize that data so that we can move past mere detection to prediction and optimization.

However, many companies are missing the opportunity to use the IoT data already being produced in their factories. Although IoT once may have looked like "too much money for too little benefit," the ratio has reversed itself. In fact, the cost/benefit ratio now suggests that building (and using) an IoT infrastructure may be one of the best investments a company can make, especially if OEE is well below world-class levels (which are in the 80+% range). The fact is that one of the best use cases for the IoT is to improve factory performance.

Utilizing IoT in a factory essentially involves the following three steps.

1. Use the IoT devices already in your production equipment, and add wired or wireless sensors where needed, to create the IoT Infrastructure that *produces* the necessary data.

2. Set up an IoT gateway (with built-in connectors for industrial equipment and sensors) to *consume, organize*, and *store* IoT data in a local server or a public cloud platform. (Note that to make the data useful, it has to be standardized and organized—partly so you do not have to worry about which supplier's sensors are being used in which locations and partly to turn the flood of data into useable information. The easiest way to do this is to store it in an existing cloud-based IoT platform.
3. Once the data have been standardized and stored, it can be accessed by analytical tools on the IoT platform. Those tools can be descriptive, predictive, or prescriptive; we can also build a "digital twin" for each line so that plant managers can tell exactly what is going on down on the plant floor whether they are in their office or halfway around the world. The key to this step is to begin *using* the information provided by the data.

Creating the IoT infrastructure allows a company to use all that IoT data for one or several of these purposes.

- *Detection* of an event or a "limit breach" (an "event" could be that a line has stopped; a "limit breach" might be that a temperature sensor is reading too high or low).
- *Monitoring* of key metrics (this is really another sort of detection—but of a condition rather than an event. Examples might be that a production line is only producing 80% of expected volume, or that the temperature has been gradually increasing, even though it has not gone past a limit yet).
- *Analytics* built using the data from a sensor or a combination of sensors—a variety of advanced analytical tools go past *reacting* to an event, or taking *preventive* actions due to monitoring, and provide *predictive* analytics. Advanced manufacturing analytics can tell you things like "your machine is going to shut down in 3 h unless you perform maintenance on it."
- *Cognitive* tools that can go beyond even the algorithms used in analytics, and learn from watching the data over time. Cognitive requires that IoT data be available for an extended period of time so that it is possible to learn from it.

CONCLUSIONS AND IMPLICATIONS

While 65% of enterprises will adopt IoT products by 2020 (according to Gartner), close to 64% claims to have failures with IoT (says a recent study by Cisco involving 1845 organizations). The greatest challenge for a successful Industrial IoT deployment is IT-OT convergence. The key to a successful IoT implementation lies in the alignment of IT and business goals. Identifying the operational technology (OT) goals is crucial for tangible IoT benefits. IT-OT convergence means utilizing the data-centric IT computing for data collected from the OT systems. IoT enables centralized data monitoring and analytics for generating actionable insights. Hence, for a successful IoT implementation, IT and OT skillsets need to work in tandem. The real benefits for a successful IoT implementation are derived from analyzing the large amounts of data collected from the OT systems, correlating that to the data from the IT systems, using AI to make predictions, and devising strategy based on the patterns and insights generated from this large pool of data.

REFERENCES

1. D. McCormick, "Power and the industrial Internet of things (IIoT) IBM software whitepaper executive summary smart plants," *Smart Grids*, Jan. 2015. [Online]. Available: https://www.slideshare.net/BenAmabaPhDPEngCPIML/wpsmartplantssmartgridsv3-44208020.
2. B. Marr, "Artificial intelligence and blockchain: 3 major benefits of combining these two mega-trends," *Forbes*, Mar. 2018. [Online]. Available: https://www.forbes.com/sites/bernardmarr/2018/03/02/artificialintelligence-and-blockchain-3-major-benefits-of-combining-thesetwo-mega-trends/
3. B. Amaba, T. Ahram, A. Sargolzaei, S. Sargolzaei, and J. Daniels, "Blockchain technology innovations," in *Proc. IEEE Technol. Eng. Manage. Conf.*, 2017, pp. 137–141.
4. Amer. Soc. Civil Eng., "2013 report card for America's infrastructure," 2014. [Online]. Available: http://www.infrastructurereportcard.org/
5. B. Amaba, "Industrial and business systems for smart cities," in *Proc. 1st Int. Workshop Emerg. Multimedia Appl. Services Smart Cities*, Orlando, FL, USA, Nov. 7–14, 2014, pp. 21–22.

6. G. Gilder, *Life After Google: The Fall of Big Data and the Rise of the Blockchain Economy*, Gateway ed. New York, NY, USA: Simon and Schuster, Jul. 17, 2018.
7. M. Hung, "Leading the IoT: Gartner insights on how to lead in a connected world," 2017. [Online] Available: https://www.gartner.com/imagesrv/books/iot/iotEbook_digital.pdf.
8. J. Dumas, "IoT failures plague 64% of users worldwide," IT Brief New Zealand, Aug. 23. 2018. [Online] Available: https://itbrief.co.nz/story/iot-failures-plague-64-users-worldwide.
9. "Research-driven insights to boost your business IQ," IBM Inst. Bus. Value. 2018. [Online]. Available: www.ibm.com/thought-leadership/institute-business-value.

CHARLA STRACENER is the GSC Electronics Industry Lead and Watson IoT/AI Technical Specialist. Contact her at charla@us.ibm.com.

QUENTIN SAMELSON is a Senior Managing Consultant within IBM's Electronics Center of Competence. Contact him at qbsamels@us.ibm.com.

JOE MACKIE is the Data Science and AI Hybrid Cloud Sales Leader for the U.S. Industrial Market, U.S. Western Region. Contact him at joe.mackie@ibm.com.

MITSUKO IHAZA is a Data Science and AI Technical Specialist at IBM. Contact her at mihaza@ibm.com.

EDITORS: Phillip A. Laplante, Penn, State University, plaplante@psu.edu
Ben Amaba, IBM, baamaba@us.ibm.com

DEPARTMENT: INTERNET OF THINGS

This article originally appeared in IT Professional vol. 20, no. 5, 2018

"Smarter" Education

Mohamad Kassab and Joanna F. DeFranco, *Penn State University*
Jeffrey Voas, *IEEE Fellow*

Today's "old-school" educational infrastructures create challenges for IT architects who want to infuse the Internet of Things (IoT) into education. Five-year old children already consider smart devices to be passé. Our next generation is already addicted to information on demand, almost from birth. The question is whether the current educational infrastructure can seize the opportunity to build and oversee the requirements necessary to allow IoT to enhance the educational experience for all at any age.

These challenges include ethical constraints (for example, students' privacy), technical constraints (such as big data captured from a wide range of heterogeneous sources), economic constraints (for example, the added cost of technology in education), and physical constraints (such as available technology and communication channels within schools). Regardless of these challenges, and based on our systematic literature review, we contend that IoT—with its distinguished features such as sensing and intelligence (artificial and regular)—can support and significantly benefit the pedagogical processes for all interrelated actors (faculty, students, and staff) as well as all assets involved (libraries, classrooms, labs, etc.).

According to Lenhart, 73 percent of US teenagers had access to a smartphone as of 2015 and 75 percent of high school students use laptops for educational purposes.[1] Nearly all US public schools have Internet access. While 69 percent of students are reported to want to use their mobile devices more frequently in the classroom, most of those students also want to automate more tasks such as note-taking, schedule checking, and research. Educational jurisdictions and institutions should embrace (not prohibit) personal devices that learners bring into the classroom and should allow students to use them as learning tools to capture intelligence more quickly and accelerate learning. It should be noted that there are security issues related to the idea of "bring your own device" (BYOD) that would need to be addressed if personal student devices are allowed to connect to an institutional IT infrastructure. That aside, it is apparent that IoT is poised to radically transform education.

Currently, there are more than 6.4 billion devices connected to the Internet excluding computers, cellphones, and tablets.[2] This number is projected to reach 20.8 billion devices by 2020, with some estimates suggesting as many as 100+ billion connected devices by 2020. Regardless of the exact number, spending in this market is expected to increase substantially, with the International Data Corporation (IDC) calculating that the worldwide market for IoT solutions will reach $7.1 trillion in four years.[3] IoT applications are already being deployed in diverse service domains such as medical, retail, consumer, smart home, environmental monitoring, and industrial Internet. Because of this ubiquitous nature, schools and academic institutions are now looking to incorporate IoT in educational activities. Why? Put simply, because they must adapt to this new world that their students already embrace. However, there is a unique twist: the students might understand this new technology better than their educators.

EMERGING SCENARIOS OF IOT IN EDUCATION

Emerging scenarios for IoT in education can be classified across three dimensions: delivery mode, perception, and learning principles (see Figure 1). Education can be delivered in one of three broad-based modes: face-to-face, remote, and hybrid. Scenarios applying IoT already exist for each mode. These scenarios suggest that IoT can actively complement and enhance

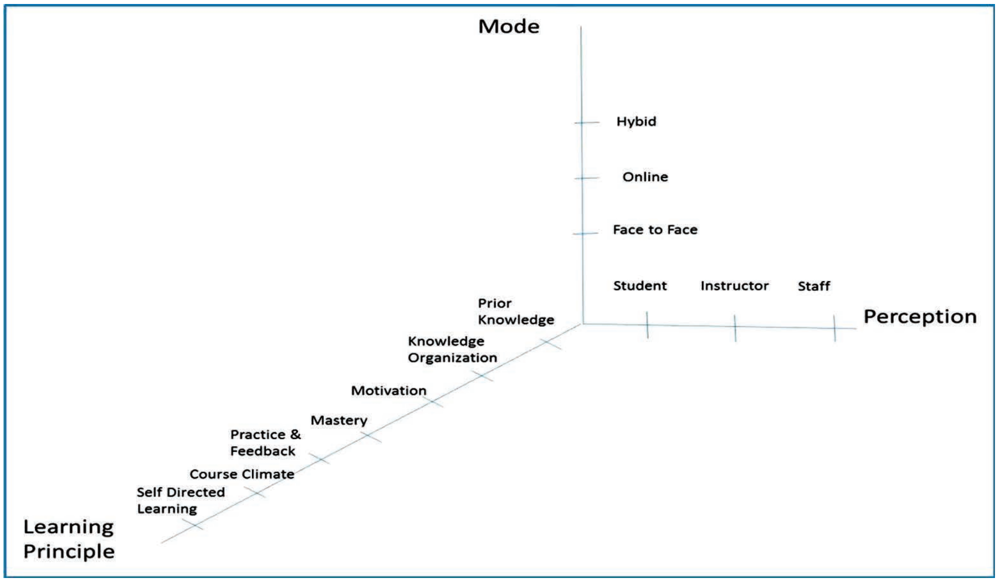

FIGURE 1. The three-dimensional scheme for IoT in education.

certain pedagogical activities with relevance to three perspectives: instructors, students, and staff.

From an instructor perspective, IoT can help manage class attendance and the availability of required equipment/devices: "Installing RFID readers at the school gate entrances, library, cafeteria, dormitory and teaching buildings, and other places to identify students' RFID electronic tags can obtain the students' activities trajectory."[4] Using IoT, instructors can initiate and manage class sessions with voice/facial/gesture commands, communicate with remote students at different locations, collect immediate feedback from students in terms of interests and likeability of an activity or a session, automatically collect data from sensors, and run analytics relevant to students in terms of behavior, performance, interest, and participation. IoT devices can help instructors confirm the identity of students and assist students with special needs.

From a student perspective, IoT can help students communicate with classmates (local or re-mote), share project data, discuss and annotate learning materials in real time, and access learning resources remotely (such as remote labs). In addition, IoT can provide students with adapted learning resources by integrating context (based on location, time, date, student interaction, knowledge level, etc.) and reasoning into a smart school system architecture.

From a staff perspective, IoT can play roles such as monitoring student emotional states and classroom environments. For example, an IoT scenario is reported in Wang[5] on monitoring and maintaining students' psychological health. Other reported scenarios relate to the potential assistance for staff members in managing and tracking fixed and portable academic resources: "Using a noise sensor, one classroom can communicate automatically to a neighbor classroom and inform them if the noise level exceeds a certain level. A warning message could be displayed on the LCD screen in the noisy room."[4] Public portable equipment (portable projectors, lab and sports equipment) can be tagged and tracked using RFID. The collected data from tracking portable equipment can be further utilized to automatically calculate patterns and trends and find inefficiencies. IoT can also assist staff members in managing events (like registration and sports events) and in managing safety and security. In addition, it can play a role in institutional energy management.

Of more concern is how IoT can assist in student learning. In Ambrose et al.,[6] the authors discussed seven principles that underlie effective learning: (1) student prior knowledge, (2) knowledge organization, (3) motivation, (4) mastery, (5) practice and feedback, (6) course climate, and (7) self-directed learning.

IOT HAS THE POTENTIAL TO MAKE A POSITIVE IMPACT ON EACH OF THE SEVEN PRINCIPLES OF EFFECTIVE LEARNING.

We believe that research by others has correctly suggested that IoT has the potential to make a positive impact on each of these seven principles. For example, in Sula et al.,[7] the authors propose an assistive smart environment system to support the learning process of an autistic student. The system identifies math and creativity ability using the Heuristic Diagnostic Teaching (HDT) process as well as sensors, a RFID tag reader, and a SmartBox device to provide personalized practice and feedback to each student (effective learning principle #5). In another example, the authors proposed an innovative system based on IoT to analyze the impact of several parameters of the physical classroom climate (effective learning principle #6).[8] The climate refers to a student's focus as it relates to their feelings and concentration during any given moment in a lecture. And in Antle et al.,[9] the Story of Things (SoT) system is proposed to enable children to learn the story behind every object they touch in a typical day. This system is inspired by Living Media and IoT, where the information is overlaid on the world through an augmented-reality contact lens to enhance the knowledge structure (effective learning principle #2).

CONSIDERATION FOR QUALITY REQUIREMENTS FOR IOT IN EDUCATION SCENARIOS

When specifying the functionality of IoT educational applications, attention is often focused on concerns such as fitness for purpose, big data, interoperability, etc. Conventional requirements elicitation techniques such as Quality Function Deployment (QFD), Joint Application Development (JAD), and domain analysis are usually adequate for these types of requirements. But in IoT educational applications, some nonfunctional quality requirements of concern include security and privacy, scalability, and humanization.

Security and Privacy

Security requirements are critical. It has become increasingly clear that educational systems are vulnerable to cyberattacks and that the number of attacks will increase. Students can stage cyberattacks that will disable a school. The Higher Education Information Security Council (HEISC) was established in July 2000 to provide coordination for the higher education sector. HEISC's mission is to support higher educational institutions as they improve information security governance, compliance, data protection, and privacy programs. To help better understand the nuance of information security issues in higher education, members of HEISC drilled down into the topic of information security and identified their top three strategic information security issues, "planning for and implementing next-generation security technologies" with increasing concern that IoT is one of the three strategic issues. IoT-employing surveillance can also address physical security, such as identifying active shooters.

Privacy is also a concern. Many of the devices used in provisioned, specialized IoT will collect various data whether that surveillance is known or not. But why is this data being collected? Who owns it? And where does it go? These are questions that need to be answered by legal professionals and governance entities that oversee education and educational standards. Education, and particularly higher education, is often identified as having a larger number of reported data breaches, and the Privacy Rights Clearinghouse (PRC) (https://www.privacyrights.org) database appears to confirm this view. In the US, there were 727 reported breaches in educational institutions between 2005 and 2014. This number is the second highest among seven investigated sectors (the first is healthcare). About 7 percent of all academic institutions in the US have had a least one breach. From 2005 to 2014, 66 percent of academic institutions listed in the PRC database experienced only one reported breach. However, about one-third of institutions with breaches have had more than one. Six percent of the listed institutions have experienced five or more reported breaches. Hacking and malware where an outside party accessed records via direct entry, malware, or spyware was the largest proportion of the reported breaches at 36 percent.

There is also a theme in academic research considering IoT functionality versus privacy. One line of research argues against restrictive regulations in response to security concerns that could prevent innovation in IoT applications. A recently conducted survey found that consumers were willing to trade off some privacy for the perceived benefits of information sharing.[10]

Scalability

Scalability creates concerns regarding the costs of implementing IoT in education. By embedding sensors into operational field environments and terminal devices, IoT networks can collect vast amounts of sensor data that reflects real-time environmental conditions of the operational field and events and activities that are occurring. The main question that arises is whether IoT devices and big data analysis will increase the existing divide into a two-class learning system: (1) those who can disburse this technology, and (2) those who cannot. At the same time, if school should be affordable for everyone, how will schools pay and service these devices? The financial obligations created by moving toward an IoT ecosystem focused on education cannot be discounted.

Humanization

There are ethical questions about the role IoT plays in human life. IoT applications involve computers interacting with computers; however, the success of IoT will depend less on this machine-to-machine connectivity and more on the humanization of such technologies. IoT clearly has the potential to reduce human autonomy. It might shift people toward particular habits and then shift power to corporations, political parties, or other organizations. For educational systems, this means that the controlling agents are those that control the tools used by academic professionals, but not the academic professionals themselves. Increased digitalization in education has a potential to create a loss of the established social benefits of attending school. There has been a long-standing, implicit social contract between educators and students—technology can erode that. Online universities are one example of the erosion of face-to-face and hand-to-hand interaction.

Conversely, IoT employed in virtual learning environments can support students with special needs (dyslexic and dyscalculic needs, for example).[7] IoT can offer the opportunity to repeat experiments without major cost or damage to property. Students with special needs are likely to face less frustration or shame in front of others since they can work and practice independently.

CONCLUSION

Advances in sensors, nanoelectronics, smart objects, cloud computing, big data, and communications on a large scale will provide continuous innovations in IoT and will impact many domains. The educational domain is no exception. While IoT education is a new conceptual paradigm in its initial phase, IoT is poised to transform the educational domain. Though there are advantages of injecting IoT technologies into education, we will likely have to compromise privacy, security, and the human-touch-associated long-standing disciplines such as K–12 education. That is a simple truth, not a myth. New thought leadership needs to be introduced to address this concern. It is ironic that today we are discussing “smarter” schools and “smarter” education, when the schools are supposed to teach us to be smarter. Now, we are trying to teach the schools. Alas, times have changed.

REFERENCES

1. A. Lenhart, *Teens, Social Media, And Technology Overview 2015*, Pew Research Center, 2015; http://www.pewinternet.org/2015/04/09/teens-social-media-technology-2015.
2. *Control Your IoT Destiny: Focus On The Business Impact Of IoT*, Gartner, 2017; https://www.gartner.com/en/information-technology/insights/internet-of-things.
3. I. Asseo et al., “The Internet Of Things: Riding The Wave

In Higher Education," *Educause Review*, June 2016; https://er.educause.edu/articles/2016/6/the-internet-of-things-riding-the-wave-in-higher-education.

4. M. Cata, "Smart University, A New Concept In The Internet Of Things," *14th RoEduNet International Conference—Networking In Education And Research* (RoEduNet NER), 2015, pp. 20–24.
5. Y. Wang, "The Construction Of The Psychological Health Education Platform Based On Internet Of Things," *Applied Mechanics And Materials*, vol. 556-562, 2014, pp. 20–24.
6. S.A. Ambroseet al., *How Learning Works: Seven Research-Based Principles For Smart Teaching*, John Wiley & Sons, 2010.
7. A. Sula et al., "A Smart Environment And Heuristic Diagnostic Teaching Principle-based System For Supporting Children With Autism During Learning," *28th International Conference On Advanced Information Networking And Applications Workshops* (WAINA), 2014, pp. 20–24.
8. H.F. Elyamany and A.H. AlKhairi, "IoT-Academia Architecture: A Profound Approach," *16th IEEE/ACIS International Conference On Software Engineering, Artificial Intelligence, Networking And Parallel/Distributed Computing*, 2015; doi.org/10.1109/SNPD.2015.7176275.
9. A.N. Antle, B. Matkin, and J. Warren, "The Story Of Things: Awareness Through Happenstance Interaction," *Proceedings Of The 15th International Conference On Interaction Design And Children*, 2016, pp. 20–24.
10. J.S. Winter, "Surveillance In Ubiquitous Network Societies: Normative Conflicts Related To The Consumer In-Store Supermarket Experience In The Context Of The Internet Of Things," *Ethics And Information Technology*, vol. 16, no. 1, 2014, pp. 20–24.

MOHAMAD KASSAB is an associate research professor of software engineering at Penn State University. Contact him at muk36@psu.edu.

JOANNA F. DEFRANCO is an assistant professor of software engineering at Penn State University. Contact her at jfd104@psu.edu.

JEFFREY VOAS is an IEEE Fellow. Contact him at j.voas@ieee.org.

IEEE COMPUTER SOCIETY

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.

COMPUTER SOCIETY WEBSITE: www.computer.org

OMBUDSMAN: Direct unresolved complaints to ombudsman@computer.org.

CHAPTERS: Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.

AVAILABLE INFORMATION: To check membership status, report an address change, or obtain more information on any of the following, email Customer Service at help@computer.org or call +1 714 821 8380 (international) or our toll-free number, +1 800 272 6657 (US):

- Membership applications
- Publications catalog
- Draft standards and order forms
- Technical committee list
- Technical committee application
- Chapter start-up procedures
- Student scholarship information
- Volunteer leaders/staff directory
- IEEE senior member grade application (requires 10 years practice and significant performance in five of those 10)

PUBLICATIONS AND ACTIVITIES

***Computer*:** The flagship publication of the IEEE Computer Society, *Computer* publishes peer-reviewed technical content that covers all aspects of computer science, computer engineering, technology, and applications.

Periodicals: The society publishes 12 magazines and 18 journals. Refer to membership application or request information as noted above.

Conference Proceedings & Books: Conference Publishing Services publishes more than 275 titles every year.

Standards Working Groups: More than 150 groups produce IEEE standards used throughout the world.

Technical Committees: TCs provide professional interaction in more than 30 technical areas and directly influence computer engineering conferences and publications.

Conferences/Education: The society holds about 200 conferences each year and sponsors many educational activities, including computing science accreditation.

Certifications: The society offers three software developer credentials. For more information, visit www.computer.org/certification.

VIRTUAL BOARD OF GOVERNORS MEETING

28 – 29 May

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

revised 2 April 2020

EDITOR: Florian Michahelles, florian.michahelles@siemens.com

DEPARTMENT: IOT NEWS

This article originally appeared in IEEE pervasive computing vol. 17, no. 2, 2018

Carousel Kittens

The Case for a Value-Based IoT

Sarah Spiekermann, *Vienna University of Economics and Business*

The IEEE P7000 Working Group aims to establish a process model for addressing ethical concerns during system design. In this article, the group's vice-chair critically reviews the IoT vision with regard to human values such as freedom, dignity, and privacy, and argues that we need to ramp up requirements engineering to prevent a dystopian IoT future.

In the early 1960s, two neuroscientists carried out one of the first investigations of what we call "embodied cognition" today.[1] Ten pairs of kittens were raised in the dark, except for three hours a day when each pair was placed in a carousel-type contraption that allowed one of them to move freely while the other was carried passively in a basket propelled by the first (see Figure 1). The kittens could not see each other, and the surrounding environment was set up in such a way that both received identical visual stimuli. The active kittens developed normally; the passive kittens not allowed to engage with the world developed serious shortcomings in intelligence and survival ability, such as visually guided paw placement, avoidance of visual cliffs, responsiveness to objects, and so on.

The IoT emerging now—with its automatic doors, virtual concierges, and self-driving cars—makes me feel a bit like that immobile kitten on the carousel, being carried along by technology without any agency of my own. While I appreciate having a safer and more efficient car, I also don't want to lose the visceral thrill of driving.

RETHINKING THE IOT VISION

Human beings are highly complex biological systems with an embodied consciousness that develops intelligence through experience with the world. Technology is making that experience passive rather than active. I'm not sure if it was this distancing of humans from the world Marc Weiser was referring to when he wrote that "the [social] problem [associated with pervasive computing] while often couched in terms of privacy is really one of control."[2] Is it only control I'm losing, or is it autonomy and hence freedom in general? In thinking about the IoT, we must consider the possibility of technological paternalism[3] nudging us into all kinds of behavior. Will smart assistants like Siri or Alexa soon tell us what to eat, who to date, what to buy, and how to spend our time, just as our cars tell us how to drive? Will AIs start to harvest IoT data to analyze our behavior on a massive scale?

I believe it's time to start thinking seriously about these value questions and what's at stake here for humanity. Not everyone shares this critical view though. Many engineers only care about building good technology and prefer to stay out of the muddy philosophical waters of good and evil. Some technology advocates take threats to our established social values lightly. Technology extremists accuse humanists of "speciesism" and embrace a science-fiction-like transhumanistic ideology that regards earth as an "insignificant speck" and reduces humanity to "raw material" for the next-generation computing platform (http://humanityplus.org/philosophy/transhumanist-faq). In the reductionist worldview of a transhumanist, petty values such as human dignity, privacy, or freedom seem to be the subjective preferences of suboptimal information processors called "humans."[4] Instead, presumably "objective" or "unbiased" AIs should make value judgements for us.[5]

This kind of reductionist thinking should be challenged. In fact, we should carefully revisit our idea of what it means to be human. Last year's June issue of *IEEE Spectrum*, "Can We Copy the Brain?," provided the latest insights into how the human brain works compared to machines (https://spectrum.ieee.org/magazine/2017/June). I was impressed by how powerful and super energy efficient our brains are: As we approach the end of Moore's law, today's most powerful supercomputers can only simulate about 1 percent of the brain's neural activity, while consuming 10 billion times more energy in proportion to the brain's size. Our brains have more than 860 billion neurons, which interconnect through one trillion synapses. We used to think that human learning occurred solely by modifying the effectiveness of existing synapses, which forms the basis of machine learning, but now we know that it results from growing new synapses between neurons—completely "rewiring" the brain. And wrap your brain around this: up to 40 percent of the synapses on a neuron are replaced with new ones every day! One author concluded, "While it is true that today's AI techniques reference neuroscience, they use an overly simplified neuron model, one that omits essential features of real neurons, and they are connected in ways that do not reflect the reality of our brain's complex architecture."[6]

In light of this knowledge, we should trust in the power of our own species perhaps a bit more and rethink our vision of the IoT and what its mission should be. We would do well to recall IEEE's mission, "advancing technology *for humanity*," and recognize more explicitly that technology is here to serve us and to help us realize *our* unused potential.

VALUE-BASED DESIGN AND IEEE P7000

Is it possible to systematically engineer systems that respect human values? Yes! In support of this very goal,

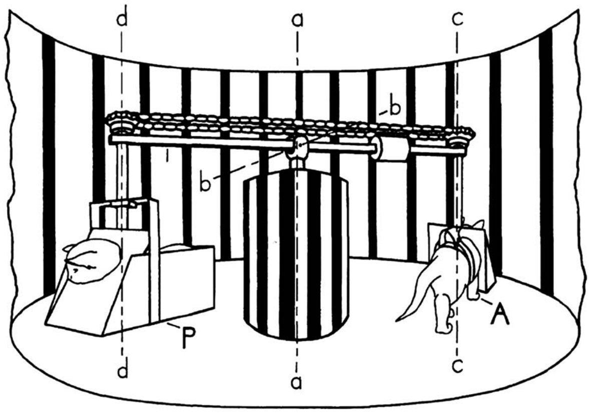

FIGURE 1. Depiction of the carousel kitten experiment, in which one of a pair of kittens is active (A) and the other is passive (B). Source: R. Held and A. Hein, "Movement-Produced Stimulation in the Development of Visually Guided Behavior," *J. Comparative and Physiological Psychology*, vol. 56, no. 5, 1963, pp. 872–876.

IEEE launched the Global Initiative on Ethics of Autonomous and Intelligent Systems (http://standards.ieee.org/develop/indconn/ec/autonomous_systems.html), which involves hundreds of scholars from around the world with multidisciplinary backgrounds connecting the humanities to engineering. One of the core achievements of this initiative is the creation of 10+ official standardization efforts focusing on value issues such as transparency, privacy, algorithm bias, and well-being.

The baseline standard, IEEE P7000, is the world's first "Model Process for Addressing Ethical Concerns during System Design" (https://standards.ieee.org/develop/project/7000.html). Two years into the standardization process, the current draft of P7000 provides a governance framework for first systematically identifying human values and then using risk management to mitigate threats to those values. Supported by a transparency process as well as a stakeholder management process, it leverages the vast corpus of philosophical thought to understand the full spectrum of a new technology's harms and benefits to human beings.

Depending on how P7000 matures, a core trait could be virtue ethics—that is, an emphasis on culturally or socially desirable character traits afforded by technology like courage, friendliness, independence, trustworthiness, and so on. If we're all treated as passive kittens, what will humanity look like 30 years from now? How can we avoid such a future? Once such long-term questions are considered, engineers and innovation teams can identify value and virtue priorities afforded by their technology. A risk approach identifies technical threats to relevant virtues and mitigates them with consistent design specifications. Elsewhere I've referred to this design approach, which aims to maximize positive value potential and minimize value harms for people in IT-rich environments, as *value-based system engineering*.[7]

CONCLUSION

Value-based design and stakeholder management are a good start but alone won't suffice. My latest research on privacy and security engineering practices shows that many engineers struggle in today's organizational environments to include nonfunctional requirements into system design.[8] The amount of time allowed for system development is too short and autonomy is too limited for many engineers to live up to their ethical responsibility. Over 60 percent of a sample of 124 engineers worldwide described themselves as being responsible for privacy and security engineering. However, more than a third reported that they work for an organization with weak privacy and security norms and nonsupportive corporate management. Therefore, businesses need to embrace value-based system engineering as well. If they fail to do and society continues its transhumanistic slide into bright and shiny but virtueless technology, we might all end up as kittens on a carousel.

REFERENCES

1. M. Crawford, *The World Beyond Your Head: How To Flourish In An Age Of Distraction*, Penguin Books, 2015.
2. M. Weiser, R. Gold, and J.S. Brown, "The Origins Of Ubiquitous Computing Research At PARC In The Late 1980s," *IBM Systems J.*, vol. 38, no. 4, 1999, pp. 693–696.
3. S. Spiekermann and F. Pallas , "Technology Paternalism—Wider Implications Of RFID And Sensor Networks," *Poiesis & Praxis*, vol. 4, no. 1, 2005, pp. 6–18.
4. J. Ito, " Resisting Reduction: Designing Our Complex Future With Machines," other, 14 November 2017; https://jods.mitpress.mit.edu/pub/resisting-reduction.
5. W. Youyou, M. Kosinski, and D. Stillwell, "Computer-Based Personality Judgments Are More Accurate Than Those Made By Humans," *PNAS*, vol. 112, no. 4, 2015, pp. 1036–1040.
6. J. Hawkins, "What Intelligent Machines Need To Learn From The Neocortex," *IEEE Spectrum*, 2 June 2017; https://spectrum.ieee.org/computing/software/what-intelligent-machines-need-to-learn-from-the-neocortex.
7. S. Spiekermann, *Ethical IT Innovation: A Value-based System Design Approach*, CRC Press, 2016.
8. S. Spiekermann, J. Korunovska, and M. Langheinrich, "Understanding Engineers' Drivers And Impediments For Ethical System Development: The Case For Privacy And Security Engineering," pending publication WU MIS Working Paper 2, Vienna Univ. of Economics and Business, 2018; http://epub.wu.ac.at/6339.

SARAH SPIEKERMANN is a professor at Vienna University of Economics and Business (WU Vienna), where she chairs the Institute for Management Information Systems, and is vice-chair of the IEEE P7000 Working Group (http://sites.ieee.org/sagroups-7000). Contact her at sspieker@wu.ac.at.

EDITOR: Ipek Ozkaya, Carnegie Mellon Software Engineering Institute, ipek.ozkaya@computer.org

DEPARTMENT: FROM THE EDITOR

This article originally appeared in Software vol. 36, no. 3, 2019

Ethics Is a Software Design Concern

Ipek Ozkaya

The IEEE and Association for Computing Machinery (ACM) joint report "Software Engineering Code of Ethics" summarizes the responsibilities of software engineers as the following: "Software engineers shall commit themselves to making the analysis, specification, design, development, testing and maintenance of software a beneficial and respected profession."[1]

The initial draft of "Software Engineering Code of Ethics" dates back to the early 1990s. The IEEE and ACM worked together and unanimously adopted version 5.2 in 1997. This version has guidelines for the eight core principles including public, product, judgment, and profession.

There have been many changes in the software industry since the mid-1990s. The 2000s marked the rise of the World Wide Web and Internet-based companies. Having their information on the Internet inevitably resulted in increased user awareness of security and privacy concerns. Commercial companies embraced open source software development to accelerate innovation. This resulted in developers paying more attention to concerns such as modularity and system decoupling. The open models of software development, along with business-to-business and consumer-to-consumer Internet businesses, inevitably presented significant design challenges in areas such as data rights, access rights, external system dependencies, security, and ownership to name a few. All of these design concerns are related to ethics.

DOI No. 10.1109/MS.2019.2902592
Date of publication: 16 April 2019

However, we seem to have been paying more attention to ethical conduct as a consequence of the increased availability of big data systems and infiltration of artificial intelligence (AI) into software-enabled systems. Today, most concerns related to ethics regarding software engineering revolve around the ethics of AI and mostly focus on the fair and unbiased use of data. Explainable AI, the verifiability and validation of AI systems, algorithmic bias, and fair and unbiased data models all are valid challenges that concern the public and governments. *AI is software*! All of these challenges have implications for how software is developed, configured, validated, and deployed. Examples in which personal information is leaked, taken without authorization, or exploited to manipulate user choices have many societal implications. The default and often naive expectation of end users has been that the software they use will build ethics and protect their privacy and information. This perception is finally changing as a consequence of many instances of data breaches. The many allegations that Cambridge Analytica accessed the private information of potentially up to 87 million Facebook users, only to be followed with a massive security breach that exposed 50 million users, have been the tip of the iceberg.[2] Verizon's annual Data Breach Investigations Report identified 53,308 security incidents and 2,216 data breaches from 65 countries for 2018.[3]

Increased interest from the public and popular media drive change. For example, ACM revised its "Code of Ethics and Professional Conduct" in 2018,[4] for the first time since 1992. Its origins date back to 1973, when the initial code of ethics for computing professionals was developed. It is quite timely that the software engineering and computing professions understand ethical engineering principles regarding how we practice as well as how we design software systems, with or without AI.

Can these two codes of conduct, the ACM/IEEE's "Software Engineering Code of Ethics" (1997) and ACM's Code of Ethics and Professional Conduct (2018), provide solid principles for those of us who are actively designing and developing software systems? The ethics principles outlined in these two documents are in fact clear, even for our new challenges: increased AI and data-related security, privacy, fairness, and bias issues. For example, the 2018 ACM code includes the following principles:

- Principle 1.4, emphasizing fairness and taking action to not discriminate
- Principle 1.6, regarding privacy expectations
- Principle 2.9, directing the design and implementation of systems that are robustly and usably secure
- Principle 3.7, recognizing and taking special care of systems that become integrated into society's infrastructure.

One viewpoint is that developments in technology, big data, and algorithms, along with the pervasive nature of software, necessitated the changes in the 2018 ACM code. As a pleasant surprise, the 1997 software engineering code also outlines clear guidelines, particularly under its product principles. Here are a few.

- 3.12: Work to develop software and related documents that respect the privacy of those who will be affected by that software.
- 3.13: Be careful to use only accurate data that was derived by ethical and lawful means, and use it only in properly authorized ways.
- 3.14: Maintain the integrity of data, being sensitive to outdated or flawed occurrences.

While ethics is the professional way in which one practices, many ethics challenges can also be treated as key design concerns. For instance, if a civil engineer builds a structure that is below the required structural reinforcement in an earthquake zone, he or she would face legal consequences. Reinforcing adequately for the longevity of the building and safety of its occupants is a given, not a negotiated or discovered requirement. In engineering programs, ethical conduct is taught at the undergraduate level with guidance on design decision making, such as the correct thickness of steel rods to use for the corresponding earthquake zone.

Similar to other engineering disciplines, in software engineering, we should treat ethics and its related concerns as fundamental design constraints. Should the practice of software engineering also consider the legal implications of ethical conduct? Perhaps the release of European Union's General Data Protection Regulation (GDPR) hints that software engineering may be moving in a direction similar to other engineering disciplines as well. GDPR suggests consequence on lack of compliance on the listed data protection rules, including legal and financial consequences of unethical conduct.[10]

THERE WILL ALWAYS BE ADVERSARIAL THREATS, EITHER INTERNAL OR EXTERNAL, THAT WILL BREACH DATA AND ABUSE SYSTEMS AND RESOURCES. HOWEVER, EMBRACING ETHICS AS AN EXPLICIT, NONNEGOTIABLE SOFTWARE DESIGN CONCERN WILL BE A START TOWARD CONSCIOUS PROGRESS.

ETHICS AS AN ARCHITECTURALLY SIGNIFICANT REQUIREMENT

Is it possible to guard a software system from unintended uses by treating ethics as an architecturally significant requirement? There will always be adversarial threats, either internal or external, that will breach data and abuse systems and resources. However, embracing ethics as an explicit, nonnegotiable software design concern will be a start toward conscious progress. Treating ethics as a design concern starts with identifying key quality attributes that all systems must implement.

Deploying successful systems that address business and user goals within cost, resource, and expected quality constraints require tradeoffs. It will be necessary to change our position when it comes to concerns related to ethics; design tradeoffs we make should not compromise these concerns. At a minimum, security, privacy, data management, transferability, and explainability concerns should be addressed for designing ethics in.

Security should be a top priority design concern when ensuring that software meets ethical standards. At a high level, software security defines a system's ability to reduce the likelihood of malicious or accidental actions that result in possibly compromising or losing information. The ISO/IEC FCD 25010 Software Product Quality Standard provides a starting point for defining security and its related aspects.[5] The standard breaks security down into confidentiality, integrity, nonrepudiation, accountability, and authenticity. These characteristics clarify who accesses what area of the system and aim to ensure that data are well identified, traced, and reproduced. There are countless known design approaches for these aspects. A classic resource for an introduction to understanding the design implications of security is *Building Secure Software* by Viega and McGraw.[6]

Privacy in software relates to the ability of end users to have control over and freedom of choice about the collection, use, and disclosure of information about themselves. Software systems should provide functionality that informs end users about the information they are disclosing and gives them control of their own data. Similar to security, existing guidelines and principles help engineers to understand privacy-by-design principles.[7]

Data management is at the core of ethical software engineering. The extent that software and its developers and users follow privacy and security principles in collecting, processing, transferring, tracking, and protecting data determines how well the resulting product fits within ethical boundaries. How data are structured, how they behave, and how they are allocated to other software elements need to be explicitly architected.[8] Who is responsible for ensuring that data are unbiased or unfairly skewed? Everyone who touches it. This is probably where we need the most debate and growth: one's bias is someone else's norm, and today's norms are tomorrow's biases.

Security, privacy, and data management are three concerns for which we have knowledge that we can draw from when considering ethics as a design concern. The recent advances in big data, computing power, and the application of machine-learning algorithms elevate two others: transferability and explainability.

Transferability has economic value, such as in the context of product lines, for example. However, this economic value must be evaluated within the context of the intended use of the transferred context. Applying image recognition advances to detect medical concerns such as melanoma may be welcomed, more so than applying the same software for surveillance. Possible design approaches may include limiting reusability and transferability to unintended contexts.

Explainability is the ability to trace how algorithms and software work and arrive at the conclusions they do. Some design strategies to assist with explainability already exist, such as instrumentation, logging, and monitoring. However, others will need to be developed when software systems also include algorithmic elements such as different applications of artificial neural networks or data classification and cleansing techniques. Currently, design challenges include tracing the steps of data analysis to the outcomes. The ability to trace why algorithm scanning recommends certain profiles over others, or understand why the autonomous software recommends certain actions over others, can have both safety and reliability implications. While transferability and explainability concerns are not new to software engineering, they are key to ethics as a design concern as they may impact how the software is structured and secured.

CALL TO ACTION

Computing and software engineering professionals have already identified what it means to engineer systems for the social good. Yet the human condition gets confused quickly and forgets. The very first principles of both ACM/IEEE's "Software Engineering Code of Ethics" (1997) and ACM's "Code of Ethics and Professional Conduct" (2018) relate to societal good.

- "Code of Ethics and Professional Conduct," Principle 1: "Contribute to society and to human well-being, acknowledging that all people are stakeholders in computing."
- "Software Engineering Code of Ethics," Principle

> 1.03: "Approve software only if they have a well-founded belief that it is safe, meets, specifications, passes appropriate tests, and does not diminish quality of life, diminish privacy, or harm the environment. The ultimate effect of the work should be to the public good."

Businesses operate under laws and regulations. Software is a public commodity. The societal impact of software has taken center stage as a result of advances in big data, machine learning algorithms, and computing power. Aspects related to legal regulation, establishing review boards, and other checks and balances for designing, developing, and consuming software are being discussed by popular media,[9] but it will take time for all of this to fall into its intended place.

While process, regulation, and legal action has its place, as software engineers, we already have a good basis to design ethics into our software. At a minimum, we should consider the following steps:

- Anyone who contributes to the software and computing profession should read and embrace ACM/IEEE's "Software Engineering Code of Ethics" (1997) and "ACM Code of Ethics and Professional Conduct" (2018).
- Ethics should be treated as a fundamental and nonnegotiable design concern. We can expand on characteristics that we already know how to design for and ensure that when we make design choices we do not compromise them.
- Young software engineers should be educated on these software design principles and codes of ethics.

Designing ethics in software is not trivial. There are gray areas, especially when it comes to how to curate and use data, but this is not completely uncharted territory either.

REFERENCES

1. D. Gotterbarn, K. Miller, and S. Rogerson, "Software engineering code of ethics," *Commun. ACM*, vol. 40, no. 11, pp. 110–118, Nov. 1997. doi: 10.1145/265684.265699.
2. M. Isaac and S. Frenkel, "Facebook security breach exposes accounts of 50 million users," *NY Times*, Sept. 18, 2018. [Online]. Available: https://www.nytimes.com/2018/09/28/technology/facebook-hack-data-breach.html.
3. S. Widup, M. Spitler, D. Hylender, and G. Bassett. (2018). *2018 Verizon data breach investigations report*. Verizon, New York. [Online]. Available: https://enterprise.verizon.com/resources/reports/DBIR_2018_Report_execsummary.pdf.
4. Association for Computing Machinery, "*ACM code of ethics and professional conduct*," ACM, New York, 2018. [Online]. Available: https://www.acm.org/code-of-ethics.
5. *Systems and Software Engineering—Systems and Software Quality Requirements and Evaluation (SQuaRE)—Systems and Software Quality Models*, Standard ISO/IEC FCD 25010, 2017.
6. J. Viega and G. McGraw, *Building Secure Software*. New York: Addison-Wesley, 2001.
7. G. Danezis et al., Privacy and data protection by design: From policy to engineering, European Union Agency for Network and Information Security, 2015. [Online]. Available: https://arxiv.org/abs/1501.03726
8. P. Clements et al., *Documenting Software Architectures: Views and Beyond*, 2nd ed. New York: Addison-Wesley, 2010.
9. C. Metz, "Is ethical A.I. even possible?" *NY Times*, Mar. 1, 2019. [Online]. Available: https://www.nytimes.com/2019/03/01/business/ethics-artificial-intelligence.html.
10. European Commission, "2018 reform of EU data protection rules." Accessed on: Mar. 2019. [Online]. Available: https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules_en.

EDITOR: Amit Agrawal, amit.agrawal@me.com

DEPARTMENT: TOOLS AND PRODUCTS

This article originally appeared in IEEE Computer Graphics and Applications vol. 39, no. 1, 2019

Mining Insights From Visual Assets

A Case Study

Zorroa democratizes machine learning by creating an easy to use data mining platform for visual assets.

In a previous paper, we talked about how machine learning (ML) and artificial intelligence were starting to affect the computer graphics industry. In this paper, we feature Zorroa, an ML company that is providing a platform for applying visual analytics to enterprise assets.

In any enterprise, assets are stored in files that are spreadsheets, text documents, images, audio files, or videos. According to Figure 6 in the paper[1] by Reinsel *et al.*, the amount of data stored in "nonentertainment image/video" at the time of writing was almost twice as large as the "productivity data." Even though it is widely recognized that actionable business insights are stored in images, audio, and videos, they have traditionally not been exploited because of lack of a) easy and secure access to the assets, b) algorithms that can be applied to these documents, and c) platforms that bring the assets to the user in an accessible manner.

This is an untapped opportunity that companies are beginning to develop. Existing Digital Asset Management systems (Cumulus, MerlinOne) are starting to integrate ML algorithms into their offerings. Companies, internally or working with integrators such as Accenture, Deloitte, or GreyMeta, are building one-off solutions to specific problems. In addition, specialists like Clarifai (clarifai.com), vidrovr (vidrovr.com), or logograb (logograb.com) provide specific ML technologies. For example, logograb can tell whether a specific logo exists in an image or video.

LANDSCAPE

With the advent of cloud computing, industry has been transitioning their IT infrastructure to the cloud. Expensive on-premises storage solutions are getting replaced by cloud storage solutions provided by the likes of Amazon, Google, or Microsoft. Although, CIOs want to make that transition for economic reasons, security concerns are slowing the move.

In addition, advancements in the ML research are happening on a daily basis. CTOs and CIOs want to leverage these advancements to deliver value for their customers and employees. However, finding ML talent is difficult. According to Nick Patience,[2] lack of skilled talent contributes up to a 36% barrier to AI adoption in the workplace. The second significant barrier is time lost in accessing and preparing the data (16%). So, even when a CIO can find ML talent, they spend their time cleaning and prepping data rather than solving real problems.

Zorroa (zorroa.com) hopes to address these needs by providing a platform [Zorroa Visual Intelligence (ZVI)] that democratizes visual analytics by making enterprise visual data and ML algorithms available together in an easy to use system.

In Figure 1, an enterprise onboards Zorroa with their existing data asset management system. Zorroa provides them with an easy to use suite of ML

DOI No. 10.1109/MCG.2018.2885416
Date of current version 6 March 2019.

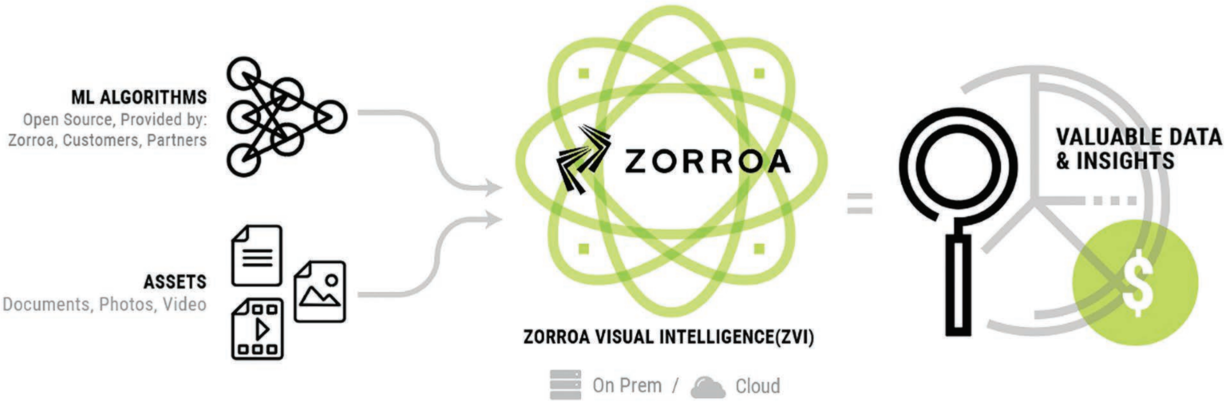

FIGURE 1. How ZVI fits into the landscape.

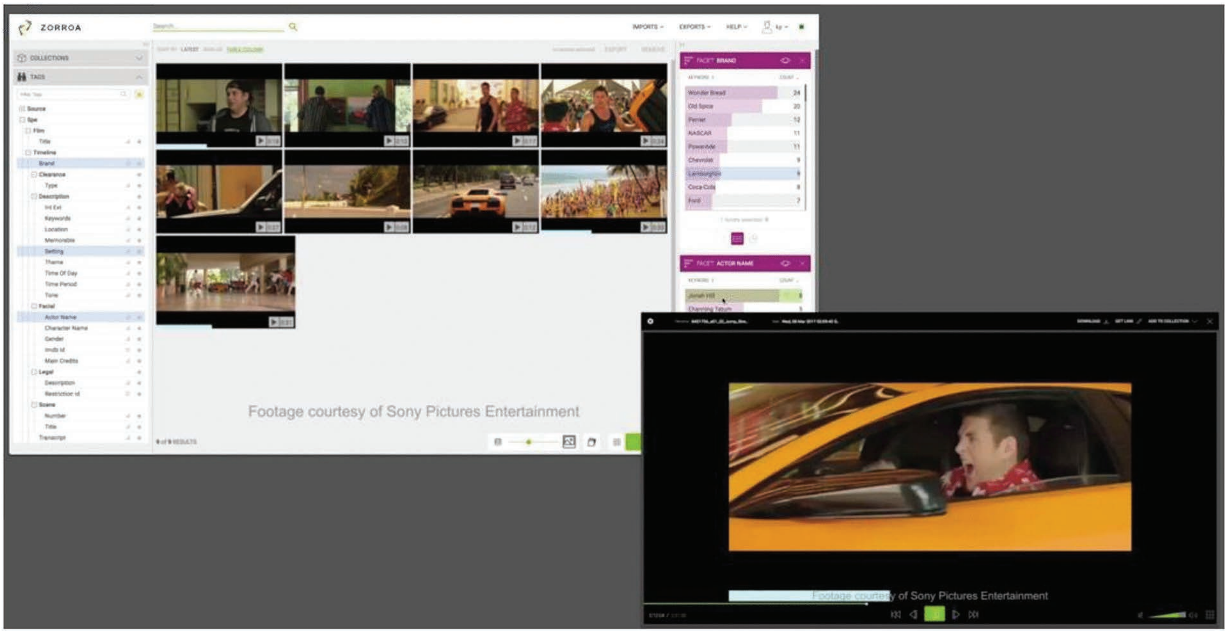

FIGURE 2. Zorroa enabled a studio client to find video clips that included a Lamborghini where an A-list actor said "Faster, faster." (Source: Sony Pictures Entertainment; used with permission.)

algorithms out of the box. The company can then extend the system by leveraging their own resources or hire Zorroa or a third party to create new ML algorithms and integrations. This enables a CIO to quickly get started in using ML within the company and develop the ML talent slowly and judiciously.

This has a second advantage. The CIOs can onboard with an on-premises solution and transparently transition to a cloud hosted solution as they get various approvals from the regulatory bodies.

CASE STUDIES

In Figure 2, we see a case study where a Zorroa studio client wanted to search and monetize their video archive. In this instance, an ad agency was looking for a video clip where an A-list actor driving a Lamborghini says "Faster, faster." The client had a highly detailed metadata database and a facial recognition system. The previous process was extremely manual and required personnel to search both the database and the videos separately. This process

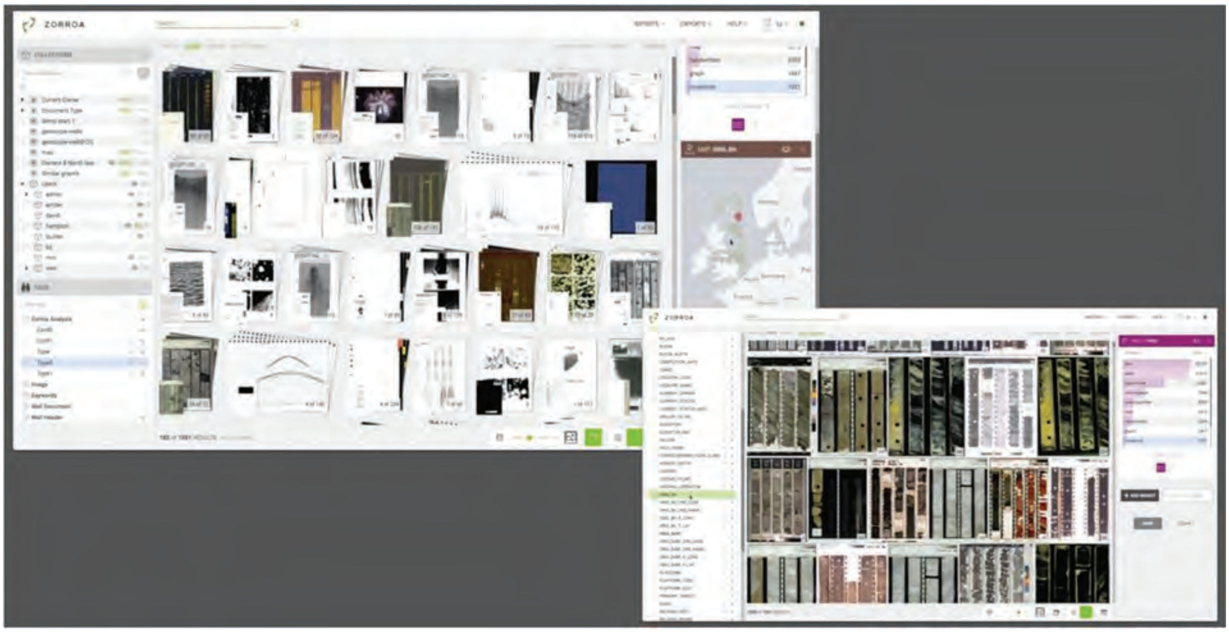

FIGURE 3. Zorroa automatically creates metadata for scanned documents, then clusters, and geolocates the documents allowing the client to rapidly find all the oil wells within a region and all the documents corresponding to these wells. (Source: Common Data Access Limited; used with permission.)

took three people an entire day (27 man-hours) to find the clips they were looking for. Zorroa integrated the database, the metadata, closed caption data, and the video in a way that allows the client to find exactly what they are looking for in a fraction of the time (3 min). In this case the value add by Zorroa was not just providing similarity algorithms that they wrote for clustering video clips (identifying video clips that were "near" each other in a high dimensional space) but also providing easy access to open source ML libraries (OpenCV in this case) that worked together with the client's existing databases and video archives.

In another example (Figure 3), an oil and gas company wanted to use their vast visual asset archive to help drive merger and acquisition (M&A) decisions. The data consisted of paper maps with hand written notes, text documents, and hand-written field reports. ZVI did optical character recognition on these documents, mapped the documents to actual GPS locations and auto-categorized the documents based on the identified text (leveraging tensor flow). Now a geo-physicist could select a region interactively on a map and find the oil wells in the selected region and all associated documents corresponding to these wells. This enables the client in making informed multimillion-dollar M&A business decisions.

WHOLE IS BIGGER THAN THE SUM OF THE PARTS

In both the examples, ZVI provided their customers with the ability to have a platform (Figure 1) where existing data sources, visual assets, and ML algorithms came together into a seamless whole. This is significant because the effect of each new algorithm and data source is not incremental but multiplicative as much more complex queries can be answered. In addition, ZVI's plug in architecture makes it easily extensible and provides for an easy on-ramp for the CIOs into the ML world.

As ML continues to affect the industries around us, Zorroa is a company to watch.

REFERENCES

1. D. Reinsel, J. Gantz, and J. Rydning, "Data age 2025: The evolution of data to life-critical. Don't Focus on Big Data; Focus on the Data That's Big," Apr. 2017.
2. N. Patience, "Skills gap is the biggest barrier to enterprise adoption of AI and machine learning," 451 Research, Sep. 2018.

EDITOR: Yong Cui, cuiyong@tsinghua.edu.cn

DEPARTMENT: STANDARDS

This article originally appeared in Internet Computing vol. 23, no. 5, 2019

Machine Learning for Internet Congestion Control: Techniques and Challenges

Lei Zhang, Yong Cui, Mowei Wang, Zhenjie Yang, and Yong Jiang, *Tsinghua University*

Internet congestion control (CC) remains a cornerstone issue in networking fields. It has attracted much research attention in academia, industry, and Internet standards organization. This article focuses on the machine learning (ML) technologies for Internet congestion control. Specifically, it summarizes the main reasons why network operators should apply ML in congestion control, surveys the latest advances of learning-based CC approaches, and explores challenges of standardizing CC with machine learning. This article provides two aspects challenges of learning-based CC that could motivate researchers to propose novel algorithms and develop standards of Internet CC with advanced ML techniques.

Internet CC is an important networking issue that the Internet Engineering Task Force (IETF) has been paying attention for more than 20 years.[1,2] As of now, research on CC can be divided into three phases. At the first stage, researchers proposed a CC scheme that all flows or users followed, and studied its effectiveness to deal with congestion. It may be called the "homogeneous CC paradigm." Subsequently, CC became the default deployment.[3,4] Many studies try to develop new schemes to improve CC and studied how these new schemes coexist with the default ones. It might be called the "competing CC paradigm." In the latest phase of CC study reviewed in this article, there is no assumption of what schemes are used by others; a flow is trying to learn how to survive well given other traffic. It might be called the "heterogeneous CC paradigm." For the first two phases, those schemes mostly deal with the complexity of network topology, the different number of flows, and their traffic demand/dynamics, which are already very complicated. For the third phase, it is considerably more complicated due to another dimension: how the other flows behave.

Recently, ML has emerged as one of the most prominent new approaches for realizing network control policies. Generally, ML techniques automatically learn policies from historical data and model the mapping from inputs to outputs without predefined rules. Among ML methods, offline learning is suitable for the scenarios where it can be assumed that the behavior of others has "converged" and assumed not to change much. Whereas online learning provides a game situation between the flows or users, these flows or users could play a cooperative environment, trying to achieve some common goals. There are interesting recent works on these lines by ML experts, e.g., the DeepMind people.[5] The CC problems can be cast with either of the above cases. Hence, some research works propose CC schemes with ML techniques and it is necessary to develop IETF standards of learning-based congestion control.

In this article, we present a survey on Internet CC from the ML perspective. First, we describe the reasons to apply ML techniques for congestion control. Then, we survey the state-of-the-art CC schemes and analyze their technical characteristics. We then discuss the challenges for developing the standards of learning-based CC in the real world. We hope that this study can encourage the researchers to design novel algorithms or develop Internet standards for congestion control.

DOI No. 10.1109/MIC.2019.2948520
Date of current version 18 December 2019.

Congestion control	Machine learning method		Objective function (Utility function)	Action	Experimental environment
Remy[9]	Offline learning	A tabular method	Throughput and delay	Cwnds and pacing	NS-2
Indigo[10]		Imitation learning	The ideal cwnds	Cwnds adjustment	Mahimahi
Custard[14]		Trust region policy optimization (TRPO)	Throughput, delay and loss rate	Sending rate	Emulab
Auraro[13]		Proximal policy optimization (PPO)	Throughput, delay and loss rate	Sending rate	Mininet
PCC[11]	Online learning	Rate probing	Throughput and loss rate	Sending rate	GENI Emulab Planetlab
Vivace[12]		Convex optimization	Sending rate, RTT gradient and loss rate	Sending rate	Emulab Mahimahi

TABLE 1. Novel CC with machine learning.

WHY CC WITH MACHINE LEARNING?

ML is suitable and efficient for learning complex behaviors, where it is not easy to find the relationship between the input and the output. Specifically, ML can provide new possible ways to generate control policies by training a learning-based agent. The adoption of ML as a solution in network system is becoming a reality.[9] As of now, the network management research group has successfully proposed two learning-based Internet drafts[6,7] in IETF.

Traditional Internet CCs only consider several metrics as decision signals,[1–4] such as packet loss and round-trip time (RTT). The existing rule-based methods elaborately make use of the above signals but achieve poor throughput when running in links with high stochastic packet loss or network jitter. In fact, decision-making can be affected by many factors, including traffic pattern, link failure, dynamic latency, packet loss, and diverse application requirements. It is difficult to get optimal or near-optimal control policies from complex network behaviors following predefined rules. ML can provide possible ways to generate models via the training approaches. It also has the ability to model the inherent relationships between the inputs and the outputs of the network environments.[8] Among the state-of-the-art techniques in machine learning, deep reinforcement learning (DRL),[20] as one of the latest breakthroughs' techniques, makes it easy to react to multidimensional feedbacks directly from network environment in variable network conditions.[13,14] In the following, we will introduce the representative efforts that conduct CC with ML methods.

NOVEL RESEARCH WORKS OF CONGESTION CONTROL

For Internet CC, the core problem is to make the decision about how and when to send data. Researchers develop flexible strategies using ML approaches to cope with varying network conditions. The most representative research works are shown in Table 1. Remy,[9] Indigo,[10] Aurora,[13] and Custard[14] perform optimization by learning the control rules offline, while PCC[11] and Vivace[12] work in an online learning manner. All of these schemes have different objective functions or utility functions as their optimization objectives. They choose different input signals, output, and ML methods, respectively. They also evaluate under different experimental environment. Next, we analyze the main techniques of these schemes.

Offline Learning

Remy[9] takes the target network assumption and the traffic model as prior knowledge and automatically

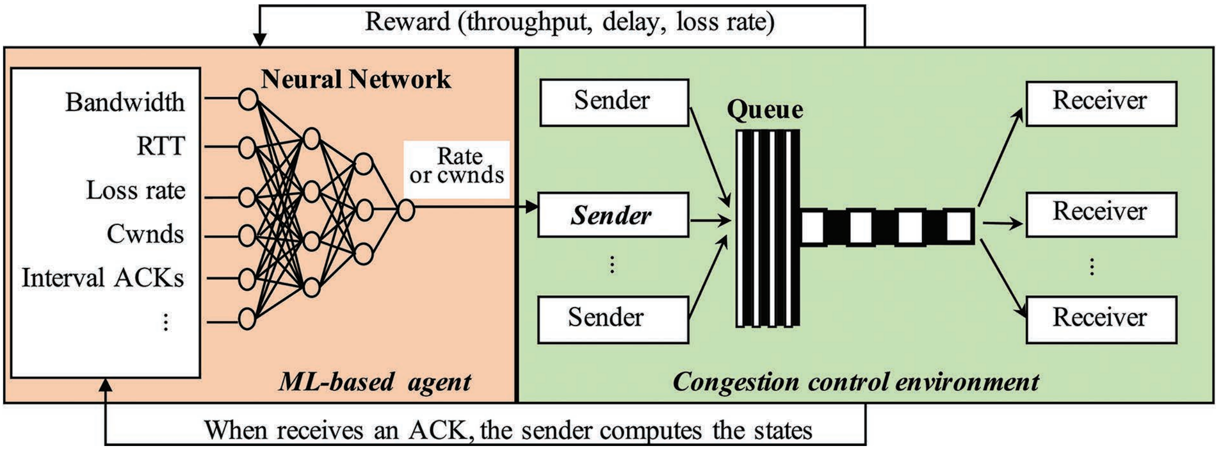

FIGURE 1. Architecture of CC with reinforcement learning.

generates a CC algorithm for the corresponding environment. In the offline phase, Remy uses an objective function to guide the rule generation process. The learned rule, i.e., RemyCC, maps the specially designed network states to the corresponding parameters about the congestion windows (cwnds). Whenever the sender receives an acknowledgment (ACK), RemyCC looks up its mapping rule and changes the cwnds according to the current network state. Although Remy helps us to improve transmission efficiency, its performance could greatly degrade if the network assumptions are violated. Indigo[10] is another method of learning-based CC scheme with the data gathered from Pantheon,[10] a system for evaluating CC schemes. Indigo learns to "imitate" the oracle rule offline. The oracle is constructed with ideal cwnds given by the emulated bottleneck's bandwidth-delay product.

Aurora[13] and Custard[14] employ DRL to generate a policy that maps observed network statistics to choose the sending rate. DRL[20], as a novel ML algorithm, trains an agent which can sample the network state, learn the policy, and improve its behavior by constantly interacting with an environment, as shown in Figure 1. The input of the agent is the network state (e.g., bandwidth, RTT, loss rate, etc.) and the output is the action, i.e., sending rate or cwnds. The goal (termed "reward") of reinforcement learning is to maximize discounted cumulative reward from the environments. Reinforcement learning is suitable for the sequential decision-making problem that can make decisions not only in discrete space (e.g., cwnds), but also in continuous space (e.g., sending rate). Aurora and Custard use different input signals and learn to make the decision by exploration–exploitation behavior. Despite the fact that the offline learning schemes can converge quickly and obtain more information, the general applicability is limited to the network scenarios where they have not been trained for.

Online Learning

PCC[11] and Vivace[12] are based on online learning. They attempt to adopt a trial-and-error mechanism to decide the sending rate. PCC's default objective function involves the throughput and the loss rate, while Vivace adopts a more complex utility function that replaces the absolute value of RTT with the "RTT gradient," i.e., the RTT with respect to time. With the carefully engineered utility function, Vivace aims to guarantee some desirable properties (e.g., fair convergence). Due to the characteristics of online learning, PCC and Vivace provide no-regret guarantees even under complete uncertainty about the environment, i.e., without inferring anything about the relation between policies and the induced utility values.[12] Both PCC and Vivace focus on looking for the change in the sending rate that may lead to the best performance, without directly interpreting the environment or making use of previous experience. Although online learning can react to network conditions quickly, its performance may diminish in some cases as their greedy exploration could be trapped at a local optimum.[22] It should be noted that online learning usually has long convergence time.[19]

CHALLENGES FOR LEARNING-BASED CONGESTION CONTROL

Latest CC schemes are indeed capable of learning useful strategies to adapt to the network environments. However, there are several challenges for Internet standardization. In this section, we attempt to present several issues of CC with ML for standards as follows. First, we discuss the challenges of ML for congestion control. Second, we explore the challenges of CC with machine learning.

Challenges of ML for Congestion Control

Input and Output Space

The space of input and output determine the primary operation of learning-based algorithms. The input space of the existing CC schemes varies greatly. For example, Remy takes the interval of ACKs, the interval of packets sent and RTT as states, whereas PCC takes the sending rate. This provides different information for the learning. A unified interface of state should be provided for standards of learning-based congestion, and its design is very challenging. Meanwhile, the output space also affects the efficiency of learning. Traditional congestion controls (e.g., Cubic[3]) or standards (e.g., IETF RFC 2581[1]) usually make the decision on cwnds. As the problem of bufferbloat becomes more and more serious, recent researchers have proposed to use rate-based transmission.[11] For learning-based congestion control, a large output space is a big challenge that makes it difficult to learn a model with ML methods. Some research works propose to decrease the decision space by reducing the dimension. For example, Indigo[10] adopts the adjustment of previous cwnds (addition, subtraction, multiplication, and division) instead of the value of cwnds.

Experimental Environment

A large amount of data and the experimental environment are important for ML algorithms, especially for offline learning methods. For the CC problem, there is no unified dataset for learning-based algorithms that is open sourced currently. In addition, simulators and emulators in learning-based CC as shown in Table 1, e.g., NS-3[15], Mininet[16], Mahimahi[17], and Emulab[21] can only provide the reproducible and rapid experimentation, but they fail to capture the dynamics of the real world. Although Stanford researchers have proposed the Pantheon[10] platform as the training ground which includes real network paths, it is not easy for researchers to tune their algorithms and use it as a performance benchmark or a performance comparison platform.

Universality

Directly deploying the offline learning-based agent from simulator or emulator can reduce the performance. The general applicability of the trained model is one of the key challenges faced by offline methods. Most offline learning methods assume that the data follow the same distribution which is not the case for real-world traffic flows. Approaches like Remy,[9] Indigo,[10] and Aurora[13] train the model on specific network conditions and perform well, to some degree, across a range of specific test conditions. However, they could not guarantee the performance of the learning-based CC methods when testing outside of the training fields. To develop standards, the offline learning-based CC model should have high general applicability that can adapt to high variance and dynamic traffic environments.

Challenges of CC With Machine Learning

Fairness

Fairness is a crucial consideration for the design of TCP CC schemes.[18] On the Internet, different CC schemes may exist at the same time and interact with each other. However, the existing CC schemes with ML techniques cannot guarantee fairness with legacy TCP. The congestion controls with ML are trained in the environment with their own objective functions. When competing with other protocols, CC with ML cannot dynamically modify their objective functions so that the CC makes decision based on the predefined optimization objective. Further, even if CC approaches are trained in an environment where it competes with other protocols, they might learn to occasionally drop packets to free up network capacity.[13] To develop standards, fairness is the critical factor to be considered in the design of CC with ML methods.

Efficiency and Effectiveness

Today, a large number of applications have high efficiency requirements.[19] The CC algorithms must be robust in the transport layer as mentioned in IETF RFC,[2] and meet the efficiency requirements of real-time transmission. Learning-based CC models (offline generated) have limitations on computational overhead, energy consumption, and response time. However, the current learning-based schemes in Table 1 do not detail the overhead comprehensively. Additionally, offline-based models work under their optimization functions and are required to consider the corresponding fault tolerance control methods in order to expel bad policies. As standards for learning-based approaches, congestion controls with ML in practice require the learning-based model to take the real-time network state and immediately output the near-optimal policy online. The tradeoff between efficiency and effectiveness is important for the performance in practical network scenarios.

Multiple Objectives

As mentioned above, the optimization objective is another core role of learning-based congestion control. The existing congestion controls in Table 1 often use a combination of throughput, latency, and loss rate as the objective function or the utility function. However, once the tradeoffs between the throughput, latency, and packet loss are determined by the designers, the optimization objective of the learning-based CC is fixed. With increasingly complicated and diverse applications, there are distinct and diverse network performance requirements. The learning-based CC algorithms for Internet standards should satisfy the diverse transmission requirements of applications. Therefore, diverse optimization objectives are also considered for standards to handle the different tradeoffs between the performance factors which the applications or users need.

CONCLUSION

Although ML shows great potential in solving CC problems, there is still a long way to go for the industry to use CC with ML directly in practice due to some practical issues of ML for networking. In this article, we first analyze the advantages of using ML in Internet CC. Next, we summarize the latest CC schemes deriving from different learning techniques. However, some issues still remain to be addressed and we discuss the challenges for Internet standardization from the CC and ML perspectives.

ACKNOWLEDGMENT

The authors would like to thank Dah Ming Chiu for his helpful suggestions and comments. This work was supported in part by NSFC Project under Grant 61872211 and in part by the National Key R&D Program of China under Grant 2018YFB1800303.

REFERENCES

1. TCP Congestion Control, IETF RFC 2581. 1999. [Online]. Available: https://tools.ietf.org/html/rfc2581.
2. Open Research Issues in Internet Congestion Control, IETF RFC 6077. 2011. [Online]. Available: https://tools.ietf.org/html/rfc6077.
3. S. Ha, I. Rhee, and L. Xu, “CUBIC: A new TCP-friendly high-speed TCP variant,” *ACM Sigops Oper. Syst. Rev.*, vol. 42, no. 5, pp. 64–74, 2008.
4. S. W. O’Malley, L. S. Brakmo, and L. L. Peterson, “TCP vegas: New techniques for congestion detection and avoidance,” in *Proc. ACM SIGCOMM Conf.*, 1994, pp. 24–35.
5. 2019. [Online]. Available: https://deepmind.com/blog/article/capture-the-flag-science
6. Intelligent Reinforcement-learning-based Network Management, draft-kim-nmrg-rl-05. 2019. [Online]. Available: https://tools.ietf.org/html/draft-kim-nmrg-rl-05.
7. Intelligent Reasoning on External Events for Network Management, draft-pedro-nmrg-intelligent-reasoning-00.” 2019. [Online]. Available: https://tools.ietf.org/html/draft-pedro-nmrg-intelligent-reasoning-00.
8. M. Wang, Y. Cui, X. Wang, S. Xiao, and J. Jiang, “Machine learning for networking: Workflow, advances and opportunities,” *IEEE Netw.*, vol. 32, no. 2, pp. 92–99, Mar./Apr. 2018.
9. K. Winstein and H. Balakrishnan, “Tcp ex machina: Computer-generated congestion control,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 4, pp. 123–134, 2013.
10. F. Y. Yan, et al., “Pantheon: The training ground for internet congestion control research,” in *Proc. USENIX*

Annu. Tech. Conf., 2018, pp. 731–743.

11. M. Dong, Q. Li, D. Zarchy, P. B. Godfrey, and M. Schapira, "PCC: Re-architecting congestion control for consistent high performance," in *Proc. 12th USENIX Conf. Networked Syst. Des. Implementation*, 2015.
12. M. Dong, et al., "PCC vivace: Online-learning congestion control," in *Proc. 15th USENIX Conf. Networked Syst. Des. Implementation*, 2018, pp. 343–356.
13. N. Jay, N. Rotman, B. Godfrey, M. Schapira, and A. Tamar, "A deep reinforcement learning perspective on Internet congestion control," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, vol. 97, pp. 3050–3059.
14. N. Jay, N. H. Rotman, P. Brighten Godfrey, M. Schapira, and A. Tamar, "Internet congestion control via deep reinforcement learning," in *Proc. Conf. Workshop Neural Inf. Process. Syst.*, 2018, pp. 1–10.
15. The network simulator—NS-3. 2001. [Online]. Available: https://www.nsnam.org/
16. N. Handigol, B. Heller, V. Jeyakumar, B. Lantz, and N. Mckeown, "Reproducible network experiments using container-based emulation," in *Proc. 8th Int. Conf. Emerg. Netw. EXp. Technol.*, 2012, pp. 253–264.
17. R. Netravali, et al., "Mahimahi: Accurate record-and-replay for HTTP," in *Proc. USENIX Annu. Tech. Conf.*, 2015, pp. 417–429.
18. Guidelines for Internet Congestion Control at End-points, IETF draft-fairhurst-tsvwg-cc-01, 2019. [Online]. Available: https://tools.ietf.org/html/draft-fairhurst-tsvwg-cc-01
19. B. Raouf, et al., "A comprehensive survey on machine learning for networking: Evolution, applications and research opportunities," *J. Internet Services Appl.*, vol. 9, no. 16, pp. 1–99, 2018.
20. V. Mnih, et al., "Playing Atari with deep reinforcement learning," 2013, *arXiv:1312.5602*.
21. B. White, et al., "An integrated experimental environment for distributed systems and networks," *ACM SIGOPS Oper. Syst. Rev.*, vol. 36(SI), pp. 255–270, 2002.
22. M. Schapira and K. Winstein, "Congestion-control throwdown," in *Proc. 16th ACM Workshop Hot Topics Netw.*, 2017, pp. 122–128.

LEI ZHANG is currently working toward the Ph.D. degree with the Department of Computer Science and Technology, Tsinghua University. Contact her at leizhang16@mails.tsinghua.edu.cn.

YONG CUI is currently a full professor at Department of Computer Science and Technology, Tsinghua University. Contact him at cuiyong@tsinghua.edu.cn.

MOWEI WANG is currently working toward the Ph.D. degree with the Department of Computer Science and Technology, Tsinghua University. Contact him at wang.mowei@outlook.com.

ZHENJIE YANG is currently working toward the Ph.D. degree with the Department of Computer Science and Technology, Tsinghua University. Contact him at yangzj15@mails.tsinghua.edu.cn.

YONG JIANG is a full professor at the Graduate School at Shenzhen, Tsinghua University. Contact him at jiangy@sz.tsinghua.edu.cn.

EDITOR: Erik P. DeBenedictis, Sandia National Laboratories, epdeben@sandia.gov

COLUMN: REBOOTING COMPUTING

This article originally appeared in Computer vol. 51, no. 8, 2018

A Role for IEEE in Quantum Computing

Erik P. DeBenedictis, *Sandia National Laboratories*

Will quantum computation become an important milestone in human progress? Passionate advocates and equally passionate skeptics abound. IEEE already provides useful, neutral forums for state-of-the-art science and engineering knowledge as well as practical benchmarks for quantum computation evaluation. But could the organization do more?

Over the past 35 years, our success in understanding quantum computing has revealed its potential as a new, disruptive computing technology. This technology is based on quantum physics and might solve classes of problems that are intractable today. The physics has been demonstrated to the level of the very earliest "classical" computers, and it's now time to see whether quantum computers can be manufactured at a larger scale and used widely, placing the topic squarely in IEEE's mission space.

IEEE is uniquely positioned to help the public and policymakers understand progress and plan the path forward. If quantum computers are destined to be a big business, students will need to be trained in the new technology so as to become part of the workforce that expands the economy. Commercial success will depend on an open and collaborative dialog on engineering, technical standardization, and policy development, which coincides with IEEE's core businesses: conferences, publications, and standards.

DEVELOP YOUR OWN VIEW ON QUANTUM COMPUTING

Instead of trying to pitch quantum computing or convey skepticism about its feasibility, let me explain the key issue so readers can form their own opinions.

Quantum computing is not expected to make an incremental advance over classical computers, like Moore's law, but might transform the notion of what is computable.

The advance from Roman numerals to place-value number systems thousands of years ago could serve as precedent for the transformation. In modern computer terminology, place-value numbers scale better to complex arithmetic operations. For example, humans can multiply and divide with place-value numbers, while these operations are impractical with Roman numerals. However, the ancients did not know what they were missing because they did not multiply and divide very much. The resulting increased ability to do arithmetic triggered the creation of entirely new areas—science, engineering, business, and products such as computers—that were not anticipated when place-value numbers were devised.

A quantum computer's qubits work like lottery tickets where you pick numbers first and there's a drawing afterwards—like Powerball but with rules dictated by quantum physics. In a quantum computer, the lottery drawing is called measurement and it labels losing lottery tickets with 0 and winning tickets with 1, thereby turning qubits into bits. However, quantum computing takes place before the lottery's drawing when it has not been decided whether a ticket will win or lose. A quantum computer's gates have the effect of swapping some of the picked numbers between pairs of lottery tickets, creating a computational model based on correlated probabilities.

Now think about computations you've done with pencil and paper or programmed on a computer. How often have you thought, "Gee, this computation would be much more efficient with predrawing lottery tickets instead of numbers." If you're anything like me, you've never thought that. Yet it's mathematically indisputable that this type of computation is vastly more efficient than today's computers for some problem classes.

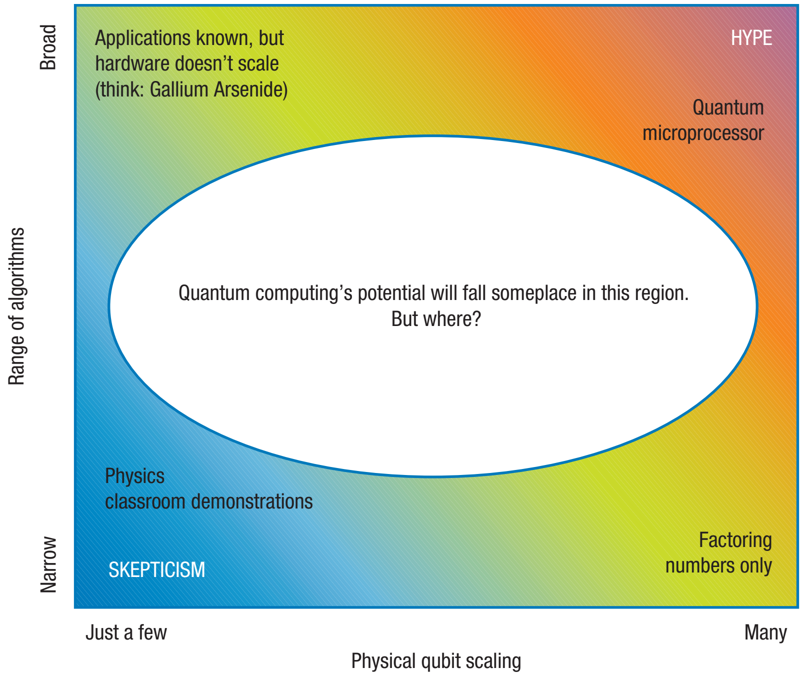

FIGURE 1. One hundred years from now, quantum computing will be at a point in the oval, sandwiched between four impossible scenarios.

Place-value numbers expand the meaning of a digit based on where it appears in relation to other digits, yielding more efficient arithmetic. Likewise, a qubit can compute with correlated probabilities before the qubit is turned into a bit. Quantum computing will trigger the invention of new computer applications, but nobody knows whether there will be enough of them to transform society. The benefit of qubits over bits might become common knowledge eventually, but right now we need a few engineers to figure it out for the first time.

SKEPTICISM, REALITY, AND HYPE

A hundred years from now, quantum computing will likely have found a position in the large white oval shown in Figure 1, sandwiched between four limiting scenarios that will ultimately be dismissed as hype or skepticism.

Figure 1's vertical axis represents the ultimate number of useful quantum algorithms, or algorithms best expressed using the predrawing lottery tickets described in the previous section. We know quantum algorithms are superior for factoring large numbers, yet theory precludes a quantum computer from being the equivalent of a microprocessor with an astronomically high clock rate. Nobody knows how many applications will eventually run best on quantum computers, particularly if society changes in response to quantum computers' ability to solve new problems.

Figure 1's horizontal axis represents our ultimate ability to engineer large-scale quantum computers. Research laboratories have created gate-type

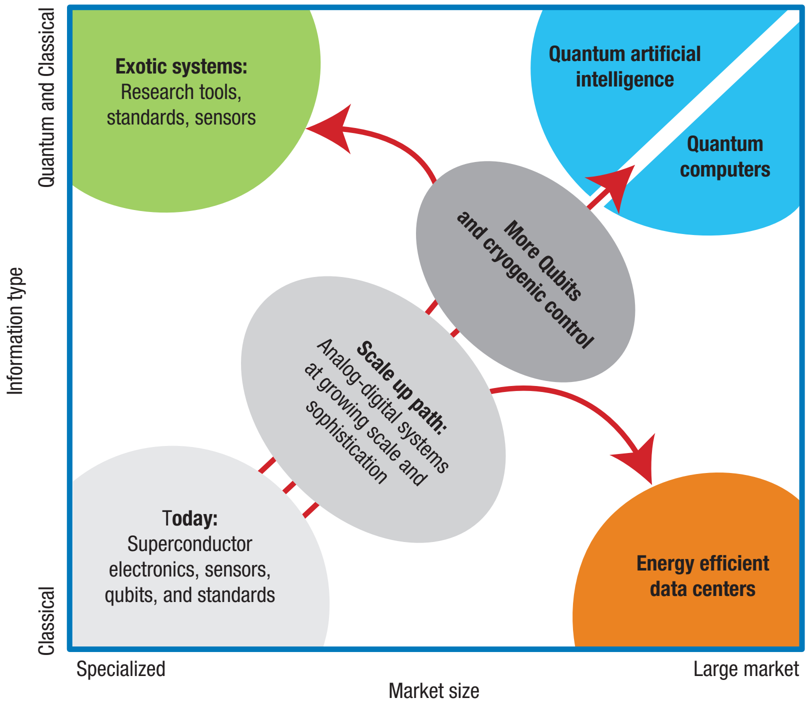

FIGURE 2. Top-level structure for a superconductor electronics and quantum information processing roadmap.

quantum computers with 50–100 qubits, and a larger number of qubits for quantum annealers (although annealers are less capable per qubit). However, this is much less than the billions of active devices in today's microprocessors.

The main debate today concerns where reality lies on the diagonal between skepticism and hype in Figure 1. However, the other diagonal must be considered as well. It's possible that many important quantum algorithms will be found, but the implementation of quantum computers will remain difficult (like, for example, Gallium Arsenide semiconductors, which were eventually set aside). Alternatively, we might master the technology behind quantum computer hardware, but the range of applications will remain limited to narrow problems such as factoring numbers.

THE QUANTUM-CLASSICAL INTEGRATED ENTERPRISE

We don't know whether quantum computing will withstand the test of time, but we're learning the scope of the enterprise if it does. Quantum computing's ecosystem was addressed at a November 2017 workshop—the 20th Biennial US Workshop on Superconductor Electronics, Devices, Circuits, and Systems in Santa Cruz, California—which included working sessions on a growth path. I based Figure 2 on this workshop's discussions, and this figure and the ideas it depicts have become a new branch of IEEE's International Roadmap for Devices and Systems roadmap for quantum computing and its ecosystem.[1]

Until recently, quantum computing breakthroughs were mostly physical science research projects demonstrating particular qubit types, many based on superconductor Josephson junctions. These operate in a cryostat at remarkably low temperatures around 0.01 K and interfaced to the outside world through a handful of coax cables. These results are in stark contrast to current computer engineering practice, which addresses chips, architecture, manufacturability, design tools, and software at the scale of billions of devices.

Future quantum computers will integrate true quantum components with classical control systems, parts of which operate at very low temperatures. This hybrid system will need to adapt many aspects of computer engineering to a previously obscure branch of electronics called *cryoelectronics*. Cryoelectronics principally includes superconductor electronics based on Josephson junctions and semiconductors operating at low temperatures.

The gray structures in Figure 2 are an evolutionary path whereby today's handfuls of qubits and chips with around 100,000 Josephson junctions can scale up; become better integrated; and address practical issues in manufacturing, analog signaling, and design tools. These intermediate systems might be useful for science experiments and niche applications, but are not expected to have large markets.

The colored structures in the corners of Figure 2 represent applications or markets that could split off eventually, including the following:

Truly exotic systems that might advance society even if produced in small quantities, such as spacecraft sensors and gravity wave detectors.

Energy-efficient classical computers for data centers and supercomputers, perhaps exemplified by the current IARPA C3 program.

Quantum computers, which will be a hybrid of quantum and classical control components. This option is divided into quantum computers running human-created algorithms and quantum machine learning,[2] the two divisions probably having different architectures.

HOW CAN IEEE HELP TECHNICALLY?

Today's qubits are unreliable or noisy as a result of imperfect materials and manufacturing, meaning they can only perform a few operations before making a mistake. For example, an ion trap quantum computer whose operations were successful 99 percent of the time—corresponding to a 1 percent error rate—warranted a Nobel prize in 2012.[3] CMOS has an error rate of about 10–21, so there's a lot of room for improvement.

IEEE's quantum roadmap effort should be able to assist these improvements.[1] Starting in the mid-1990s, the semiconductor industry managed the historic rise of CMOS in part through the International Technology Roadmap for Semiconductors, whose principal purpose was to identify the materials science and device physics research necessary to maintain the expected rate of progress. Now called the International Roadmap for Devices and Systems, this roadmap has become part of IEEE's Standards Association (IEEE-SA) and might be able to extend its historical role in orchestrating the development of technology to include quantum computers.

Creating a roadmap requires knowing where the road leads and measuring how fast you're going. The next milestone on the road will be *quantum supremacy*—the point at which a quantum computer can solve a problem not possible for any classical computer.

IEEE-SA also has an effort called P7131 that is developing a metric to measure a quantum computer's capability or quality (http://standards.ieee.org/develop/project/7130.html). At the time of this writing, major research organizations tout the number of qubits in their research-grade quantum computers, such as 49, 50, and 72 qubits, implicitly using qubit count as a metric. Although IEEE-SA will follow a consensus-based process to define a quantum computer metric, I can report that current discussions include combining the number of qubits, qubit stability or operational reliability, and architectural efficiency.[4] There's also an understanding that benchmark programs will be required at some point, or the equivalent of Linpack for the TOP500 Supercomputer list.

HOW CAN IEEE FACILITATE COMMUNICATIONS?

IEEE's main service to the community involves conferences and publications, both of which evolve to embrace new technologies. The Rebooting Computing initiative, which sponsors this column, began in 2013 and annually hosts the International Conference on Rebooting Computing (icrc.ieee.org), which is a venue for reporting research results that include quantum computing. The same event hosts an industry summit for business opportunities.

The industry summit has become a forum for quantum computing announcements, a role its hopes to keep and expand. There's a plan to include a quantum "competition," similar to the Gordon Bell award or the TOP500 list for supercomputers.

I've been in contact with various IEEE societies and councils and have noted their interest in supporting conferences and special journal issues on quantum engineering, but further news on these will have to wait until calls for papers are issued.

An IEEE standard P7130 is also being developed to establish common terminology and notation for quantum computing concepts (http://standards.ieee.org/develop/project/7130.html).

It is becoming increasingly likely that quantum computers will succeed in factoring large numbers and force a change in cryptographic codes, such as the well-known https. However, it is possible that the underlying technology will find other uses, ultimately having a transformative effect on society like the invention of place-value arithmetic. Reality almost certainly lies in between. Given the magnitude of the consequences and the fit to IEEE's technical area and member skills, I'm suggesting that IEEE consider a carefully thought out approach to figuring where, exactly, reality lies between the extremes.

ACKNOWLEDGMENTS

This document has been written as part of an emerging IEEE effort that includes the following individuals: Tom Conte, Paolo Gargini, Scott Holmes, Travis Humble, Garrett Kenyon, Bruce Kraemer, Oleg Mukanov, Andrew Sornborger, Bill Tonti, Elie Track, R. Stanley Williams, and Whurley.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the US Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the US Department of Energy or the US Government.

REFERENCES

1. D.S. Holmes and E.P. DeBenedictis, "Roadmapping Cryogenic Electronics and Quantum Information Processing," In *International Roadmap for Devices and Systems, 2017 Edition, Executive Summary*, IEEE, 2017; https://irds.ieee.org/images/files/pdf/2017/1523037982875-IRDS_CryoE_HolmesDeBenedictis_v4b.pdf.
2. E.P. DeBenedictis, "A Future with Quantum Machine Learning," *Computer*, vol. 51, no. 2, 2018, pp. 68–71.
3. "*The Nobel Prize in Physics 2012: Press Release*," NobelPrize.org, 2 Jul. 2012; https://www.nobelprize.org/nobel_prizes/physics/laureates/2012/press.html.
4. L.S. Bishop et al., "Quantum Volume"; https://pdfs.semanticscholar.org/650c/3fa2a231cd77cf3d882e1659ee14175c01d5.pdf.

ERIK P. DEBENEDICTIS is a technical staff member at Sandia National Laboratories Center for Computing Research. He's a member of IEEE, ACM, and APS. Contact him at epdeben@sandia.gov.

This article originally appeared in Computer vol. 52, no. 10, 2019

The U.S. National Quantum Initiative

Paul Stimers, *K&L Gates*

FROM THE EDITOR

As the readers of this periodical know, quantum computing has made substantial strides during the past decade. A recent special issue of *Computer* (June 2019) discussed the practical accomplishments of the field and its prospects for the future. To put some of those actions in a bigger context, we thought it would be useful for our readers to understand some of the work that is being done in Washington, D.C., to support and develop quantum technology. As a result, we asked one of the industry leaders, K&L Gates partner Paul Stimers, executive director of the Quantum Industry Coalition, to describe the current state of quantum policy. The Quantum Industry Coalition is a trade association, a group that represents the new quantum industry to the U.S. Congress and the U.S. administration. Hence, this article gives an insider's view of the new government programs to support quantum research and development. As with all things that touch government and politics, this kind of support has the potential for stirring up diverse and conflicting opinions. Our intent here is not to debate the effectiveness or ineffectiveness of such government sponsorship of research but to give a description of the new quantum information programs to our readers, who may easily be involved in these programs, as quickly as possible. Any opinions in this article are held by the author and are not necessarily those of the IEEE, the Quantum Industry Coalition, or his law firm. —*David Alan Grier*

Recognizing the potential impact of quantum technology, the U.S. government has enacted legislation to coordinate and accelerate U.S. quantum research and development. This article provides an inside look at the National Quantum Initiative Act and current U.S. quantum policy.

Over the past few years, the quantum fields have begun to move from theory toward practice. We have seen substantial progress in quantum computing, cryptography, communication, clocks, and sensors. An ecosystem of quantum companies is developing, building on research from universities and national labs and working in partnership with some of the world's largest IT companies to advance and commercialize quantum research and development.

As quantum technologies advance, they offer some interesting opportunities to support national security, but they potentially threaten it as well. In the relatively near term, quantum clocks and sensors will be able help the U.S. military navigate in the event that an adversary disrupts or destroys the GPS system. In the medium term, quantum cryptography and quantum communication will be able to help both

DOI No. 10.1109/MC.2019.2930795
Date of publication: 24 September 2019

the United States and its adversaries protect sensitive messages. In the long term, a universal quantum computer has the potential to break most forms of encryption currently in use, putting at risk all sorts of sensitive military and civilian information.

Similarly, quantum technologies may potentially have profound economic impacts. As part of a cloud-based computing toolkit, quantum computers may be able to help companies, researchers, and governments answer questions that traditional computers cannot. Among the most promising applications are drug development, logistics, and traffic optimization. Quantum communication and postquantum cryptography could form the next defensive weapons in the never-ending battle for cybersecurity.

Amazon founder Jeff Bezos argues that it is still "day one" of the Internet. The Quantum Industry Coalition believes that it is still "hour one," and maybe only "minute one," of quantum computing and other quantum technologies. The June 2019 issue of *Computer* devoted to quantum realism addressed, in several respects, both the promise and the challenge posed by quantum technology. At this point, we cannot accurately predict which quantum technologies will be successful, what benefits they will bring, and how they will support a national security agenda. However, the potential impact of these technologies on both economic and national security activities is simply too great to ignore.

The rest of the world is not blind to the potential military and economic benefits of quantum leadership: the European Union, Canada, Australia, Israel, Japan, and other countries have made significant investments in quantum research and development during the past 20 years. No country is investing as heavily as China, however. China's public accomplishments—likely augmented by substantial secret efforts—are impressive. It has developed and demonstrated quantum key distribution via satellite (*Micius*, launched in August 2016) and ground-based quantum communications (the Quantum Beijing–Shanghai Trunk, in use since September 2017) and has announced that it will spend US$10 billion to build a National Laboratory for Quantum Information Sciences in Hefei.[1]

NATIONAL QUANTUM INITIATIVE ACT

Against this background, Congress acted quickly last year to codify a multiagency approach to accelerating and coordinating quantum research and development. The National Quantum Initiative Act (NQIA) was introduced in both the House and the Senate in June 2018 and passed by both bodies overwhelmingly in December 2018, at which point it was quickly signed into law by the president.[2] The speed with which it moved stemmed from a bipartisan, collegial drafting process, substantial stakeholder outreach, and a sense of urgency created by news of foreign achievements in the field.

The NQIA creates a central structure for the National Quantum Initiative (NQI) and authorizes three agencies—the National Institute of Standards and Technology (NIST), the National Science Foundation (NSF), and the Department of Energy (DOE)—to implement the bulk of the program. It authorizes up to US$1.275 billion for the NQI through five years.

Central structure

The act requires the president to implement a program to develop and advance quantum technologies. The outline for this program follows a model that was devised, in 2003, for the National Nanotechnology Initiative, a piece of legislation that was very effective at developing its technology. The quantum program will be based on a 10-year plan to accelerate the development of quantum information science and technology applications in the United States. It will expand quantum research through new investment; increase efforts to provide the training and education needed to build a workforce skilled in quantum; create an interagency process to coordinate the activities of the different federal agencies, liaise with industry, and leverage existing federal investments such as the national laboratories; and create a network of quantum centers throughout the country.

National Quantum Coordination Office

At the center of the NQI is the National Quantum Coordination Office that will be housed in the White House Office of Science and Technology Policy. This coordinating office will act as a hub for all civilian federal quantum activities and a point of connection between federal and nonfederal quantum activities. It will coordinate the NQI's efforts; conduct public outreach; and promote access to and the application of quantum research, technologies, innovations, and resources. Although the 2018–2019 government shutdown slowed its initial formation, the office is staffing up and getting underway.

Subcommittee on Quantum Information Science

The NQIA requires the administration to form the Subcommittee on Quantum Information Science (SCQIS) within the National Science and Technology Council, a cabinet-level group that coordinates scientific research within the federal government. In anticipation of this legislation, the Trump administration moved to create such a subcommittee early last year, putting it under the council's Committee on Science. That step elevates quantum research to the highest levels of national interest.

The SCQIS has members drawn from across the government, including the U.S. Department of Defense (DOD) and the Office of the Director of National Intelligence. Its cochairs are delegates from the DOE, NIST, NSF, and the White House Office of Science and Technology Policy. In autumn 2018, the SCQIS produced its first report, which argued for a "science-first approach" that would start by supporting research, building a workforce, and working with industry. It proposed a grand-challenges approach to research, focusing on "fundamental scientific or technology problems with answers that will be transformative" for the nation and have "broad economic and scientific impact."[3]

The National Quantum Initiative Advisory Committee

The NQIA requires the president to create the NQI Advisory Committee. This group has a somewhat different mission from the SCQIS, as it is intended to include representatives from industry, universities, and federal laboratories who are qualified to provide advice and information on quantum information science and technology research and development, demonstrations, standards, education, technology transfer, commercial application, and national security and economic concerns. It reviews the trends in research and industry and identifies opportunities to improve the NQI. It is required to report to the president and Congress every six months. Although the White House announced, in spring 2019, that the committee would be formed, convening the body required further executive action. The president signed an executive order on 30 August establishing the advisory committee, with membership including the director of the Office of Science and Technology Policy and up to 22 experts appointed by the Secretary of Energy.[7]

FEDERAL AGENCIES

The NQIA requires three agencies to take specific roles in developing policy and promoting research: NIST, the NSF, and the DOE. It creates two new kinds of institutions: multidisciplinary centers for quantum research and education (under NIST) and quantum information science research centers (under the DOE).

NIST

Giving a central role to NIST is not a surprising decision as the agency has been engaged in quantum research and development for more than a decade, and it produced a report, in 2009, that has guided government policies to this point. The NQIA gives NIST three roles that are traditionally within its mandate. First, it requires NIST to deal with measurement issues by supporting and expanding the "research and development of measurement and standards infrastructure necessary to advance commercial development of quantum applications."[2] [See Sec. 201(a)(1).] Second, it requires NIST to use its existing programs to help train quantum scientists and expand the quantum workforce. Third, it requires NIST to establish or expand collaborative ventures with industry and other government agencies. This approach proved effective for nanotechnology and is expected to have similar results for quantum information technology.

The NQIA also requires NIST to convene a quantum consortium of stakeholders to identify the future measurement, standards, cybersecurity, and other appropriate needs for supporting the development of the U.S. quantum information science and technology industry. The act authorizes up to US$80 million per year through fiscal year 2023 for NIST's activities, including the consortium. NIST began this work by partnering with SRI International, Menlo Park, California, to create the Quantum Economic Development Consortium (QEDC). The QEDC met for the first time in October 2018 and has been holding regular stakeholder meetings and workshops. It started with approximately 25 members, including several of the large computing firms, and has been growing since then.[4]

The NSF

The NQIA gives the NSF the natural role of developing research programs and supporting graduate education in the quantum information sciences. The NQIA requires the NSF to carry out a basic research and education program on quantum information science and engineering, which includes awarding competitive grants to universities, nonprofits, and consortia to support basic interdisciplinary quantum research and promote human resources development in all aspects of quantum information science and engineering. In particular, the act requires the NSF use its existing programs to

- use its existing programs to improve quantum education at the undergraduate, graduate, and postgraduate levels and increase participation in the quantum fields
- formulate goals for quantum science, research, and education activities to be supported by the NSF
- coordinate NSF research efforts
- engage with the rest of the government, research communities, and potential users of the information that the NSF produces.

In its 2020 budget request, the NSF asked for US$105 million for quantum information research programs.

Multidisciplinary centers for quantum research and education

Under the NQIA, the NSF is responsible for creating multidisciplinary centers for quantum research and education. The act requires the NSF to open between two and five such facilities. The NQIA authorizes up to US$10 million, out of existing NSF funding, per center per year through fiscal year 2023. The centers will be created through the common competitive process and likely be university based, although the legislation anticipates that there may be consortia of universities that include collaborators from the private sector. The centers are to conduct basic research and education activities to advance quantum science and engineering; support curriculum and workforce development; and leverage industry perspectives, knowledge, and resources. The centers will be authorized for a five-year term with five-year renewals and may be terminated for cause if they underperform.

The DOE

The DOE is also a natural contributor to the NQI. It has a long history of supporting high-performance computing, and its laboratories, such as Sandia National Laboratories, have conducted some of the fundamental research on quantum information. The NQIA requires the DOE to carry out a quantum information science research program. This basic research program will formulate DOE quantum information science research goals, provide research experiences and training for undergraduate and graduate students, coordinate research across existing DOE programs, and engage with the rest of the government, research communities, and potential users of the information the DOE produces.

National quantum information science research centers

To support its basic research, the DOE will create the second form of new research entity, the national quantum information science research centers. In parallel with its requirements for the NSF, the NQIA requires the DOE, through its Office of Science, to establish and operate at least two and as many as five of these centers to conduct basic research to accelerate scientific breakthroughs in quantum information science and technology and support research conducted by the DOE. Some of the existing projects at DOE labs may form the core of the new centers. The NQIA states that the centers are to coordinate with other DOE initiatives, including the nanoscale science research centers, energy frontier research centers, energy innovation hubs, and national laboratories as well as with higher education and industry. As with the NSF centers, DOE centers will be authorized for a five-year term with five-year renewals and may be terminated for cause if they underperform. The NQIA authorizes up to US$25 million per center per year through fiscal year 2023, taken out of DOE funding.

OTHER U.S. QUANTUM WORK

Of course, the United States is not new to the field of quantum technology. Richard Feynman discussed the possibility of quantum computing in his famous 1959 lecture "There's Plenty of Room at the Bottom."[5] Throughout the 1980s, Feynman and others advanced the science and promoted the idea of building quantum computers. In 1994, Peter Shor described an algorithm for factoring large numbers—thereby breaking many of the cryptography systems currently in use—using a quantum computer.

NIST and the DOD held their first quantum information workshops in the mid-1990s. Two NIST partnerships, the Joint Institute for Laboratory Astrophysics and the Joint Quantum Institute, have been doing quantum research since the 1990s and 2000s, respectively. Several of the national laboratories have ongoing quantum research programs. The first mention of quantum information science in the U.S. budget was in 2008 when it was included in the Networking and Information Technology Research and Development Program, under the National Science and Technology Council.

It is widely understood that the DOD and the intelligence community have conducted quantum research for years and that some portion of the research is classified. The 2019 National Defense Authorization Act established a defense quantum information science and technology research and development program to coordinate and accelerate the DOD's quantum research and development efforts.[6] The program was designed to develop and manage a balanced portfolio of fundamental and applied quantum research and transition that research into deployable technology. This year's National Defense Authorization Act is expected to expand the DOD's quantum capabilities further as both the House and Senate have passed versions of the bill that include quantum language supported by the Quantum Industry Coalition. If enacted, the language will enable the DOD to coordinate closely with the civilian elements of the NQI as well as with industry, academic institutions, and national laboratories.

PRINCIPLES FOR FEDERAL ACTION

Our organization, the Quantum Industry Coalition, is designed to represent the American quantum information industry to Congress and the administration and present the industry's point of view to both of those bodies. The coalition strongly supports the NQIA because its members believe that quantum information technology is an important field of research that has the potential to expand the U.S. economy and promote national security and that U.S. leadership in the field is vitally important.

THE NQIA IS A STRONG PIECE OF LEGISLATION THAT IS LIKELY TO SUCCEED.

We believe that the NQIA is a strong piece of legislation that is likely to succeed. First, it follows the examples of other successful programs, such as the National Nanotechnology Initiative. Second, it follows a set of key principles in American science policy that have been tested and proven to be effective. The members of the coalition believe—as do many across industry, academia, and professional societies—that federal action to promote U.S. quantum leadership should have the following characteristics:

- *Set broad goals for quantum research and development*: These should not be imposed top-down by the government but should develop out of an ongoing discussion among industry, academic, civilian government, and military stakeholders. The goals should focus on results that help strengthen the U.S. economy and national security.

- *Exchange information*: Compiling and sharing nonsensitive information about public- and private-sector research and development will enable the government to assess progress toward goals and minimize gaps and overlaps. Participation should be voluntary. Steps should be taken to avoid collecting sensitive information and protect any such information that is collected.
- *Accelerate research and development toward usable results*: Clear objectives and information about current efforts should enable the federal government to direct increased funding toward a balanced mix of fundamental and translational research efforts that will yield usable technology in the medium term as well as advancing quantum science in the long term.
- *Advance U.S. national security*: Only by harnessing the capabilities of a strong and diverse quantum economy will the United States be able to win the national security quantum race. Excessive or unwieldy secrecy and export control requirements will stifle U.S. quantum research and development while spurring innovation overseas.
- *Promote workforce development*: American quantum companies need to have access to a pool of qualified American workers. Educational institutions should be incentivized to respond to projected industry demand.
- *Work with U.S. allies as appropriate*: American leadership does not need to exclude other countries. Many of our closest allies are home to leading companies and research institutions in the field that can help advance U.S. priorities in partnership with industry.
- *Avoid common pitfalls*.
- *Mandating specific technologies*: It is too early to know which technologies will pan out, and the federal government should not play that role at any point.
- *Picking winners and losers in the marketplace*: The government should maintain a level playing field and let competition determine who wins and loses.
- *Crowding out private-sector investment*: The government should leverage private-sector investment, not compete against it.
- *Excessively controlling technology*: Over-classification and stringent export controls will merely push development overseas while slowing it in the United States.

Some of the United States' strongest international competitors have shown a willingness to commit large amounts of money and large numbers of people to this race and back those commitments with a cohesive strategy, an aggressive industrial policy, and a convenient disrespect for intellectual property law. The United States cannot take that approach, and even if it could, it would not do so as well as its competitors. Instead, it must focus on executing the approach that has worked so well before: free inquiry, free enterprise, cooperation, coordination, and investment. If the United States is able to lead the way in quantum technology, the value to the country and the world could be tremendous.

REFERENCES

1. E. B. Kania and J. K. Costello, "*Quantum hegemony? China's ambitions and the challenge to U.S. innovation leadership*." Accessed on: Aug. 3, 2019. [Online]. Available: https://s3.amazonaws.com/files.cnas.org/documents/CNASReport-Quantum-Tech_FINAL.pdf
2. U.S. House. 115th Congress, 2nd Session. *One Hundred Fifteenth Congress of the United States of America: At the Second Session, National Quantum Initiative Act* (Pub. L. No. 115-368). Accessed on: Aug. 3, 2019. [Online]. Available: https://www.congress.gov/115/bills/hr6227/BILLS-115hr6227enr.pdf
3. Executive Office of the President of the United States, "*National strategic overview for quantum information science*." Accessed on: Aug. 3, 2019. [Online]. Available: https://www.whitehouse.gov/wp-content/uploads/2018/09/National-Strategic-Overview-for-Quantum-Information-Science.pdf
4. National Institute of Standards and Technology, "*Quantum Economic Development Consortium (QED-C)*." Accessed on: Aug. 3, 2019. [Online]. Available: https://www.nist.gov/sites/default/files/documents/2018/10/11/4._joseph_broz_plans_for_the_quantum

_economic_development_consortium_qed-c.pdf

5. R. P. Feynman, "*There's plenty of room at the bottom*." Accessed on: Aug. 3, 2019. [Online]. Available: http://calteches.library.caltech.edu/47/2/1960Bottom.pdf
6. U.S. House. 115th Congress, 2nd Session. *One Hundred Fifteenth Congress of the United States of America: At the Second Session, Defense Quantum Information Science and Technology Research and Development Program*, (Sec. 234). Accessed on: Aug. 3, 2019. [Online]. Available: https://www.congress.gov/115/bills/hr5515/BILLS-115hr5515enr.pdf
7. Executive Office of the President. (2019, Aug. 30). "Establishing the National Quantum Initiative Advisory Committee, Exec. Order No. 13,855, 84 FR 46873," *Federal Register*. [Online]. Available: https://www.federalregister.gov/documents/2019/09/05/2019-19367/establishing-the-national-quantum-initiative-advisory-committee.

PAUL STIMERS is the executive director of the Quantum Industry Coalition, Washington, D.C., and a partner at K&L Gates, Washington, D.C. Contact him at paul.stimers@klgates.com.

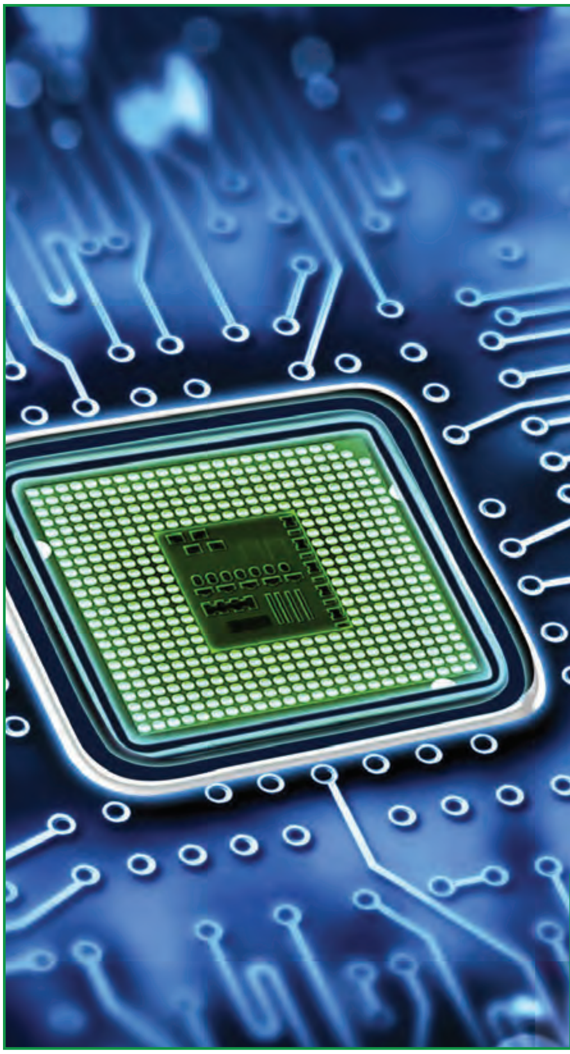


stay connected.

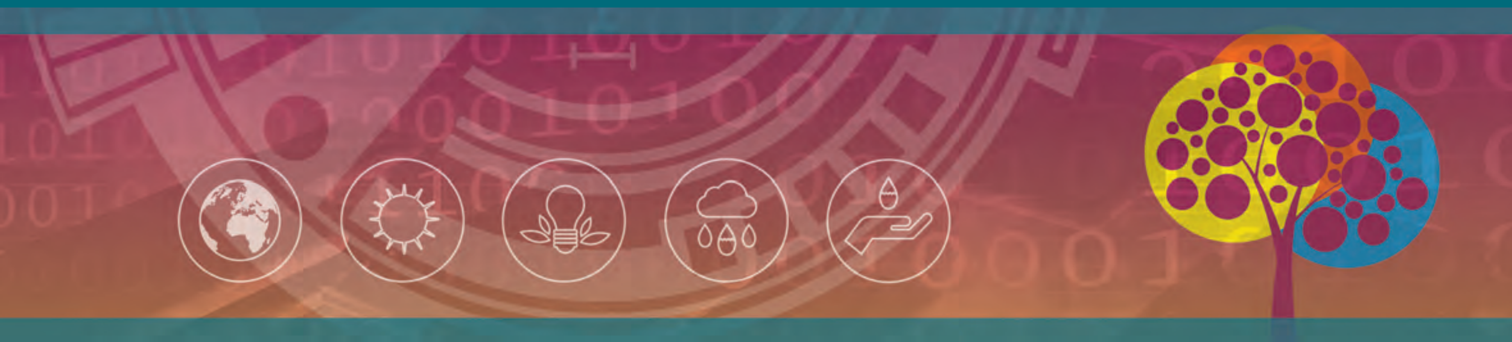

Conference Calendar

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.

JUNE

7 June

- ARITH (IEEE Int'l Symposium on Computer Arithmetic), Portland, USA

14 June

- CVPR (IEEE Conf. on Computer Vision and Pattern Analysis), Seattle, USA

22 June

- CSF (IEEE Computer Security Foundations Symposium), Boston, USA

29 June

- DSN (IEEE/IFIP Int'l Conf. on Dependable Systems and Networks), Valencia, Spain

30 June

- MDM (IEEE Int'l Conf. on Mobile Data Management), Versailles, France

JULY

6 July

- ICME (IEEE Int'l Conf. on Multimedia and Expo), London, UK

8 July

- ICDCS (IEEE Int'l Conf. on Distributed Computing Systems), Singapore

13 July

- COMPSAC (IEEE Computer Software and Applications Conf.), Madrid, Spain

28 July

- CBMS (IEEE Int'l Symposium on Computer-Based Medical Systems), Rochester, USA

AUGUST

1 August

- JCDL (ACM/IEEE Joint Conf. on Digital Libraries), Xi'an, China

14 August

- SmartIoT (IEEE Int'l Conf. on Smart Internet of Things), Beijing, China

16 August

- Hot Chips, Palo Alto, USA

17 August

- ACSOS (IEEE Int'l Conf. on Autonomic Computing and Self-Organizing Systems), Washington, DC, USA

19 August

- RTCSA (IEEE Int'l Conf. on Embedded and Real-Time Computing Systems and Applications), Gangnueng, South Korea

24 August

- FiCloud (Int'l Conf. on Future Internet of Things and Cloud), Rome, Italy

31 August

- RE (IEEE Int'l Requirements Eng. Conf.), Zurich, Switzerland

SEPTEMBER

7 September

- EuroS&P (IEEE European Symposium on Security & Privacy), Genova, Italy

14 September

- CLUSTER (IEEE Int'l Conf. on Cluster Computing), Kobe, Japan

21 September

- ASE (IEEE/ACM Int'l Conf. on Automated Software Eng.), Melbourne, Australia

24 September

- BigMM (IEEE Int'l Conf. on Multimedia Big Data), New Delhi, India

27 September

- ICSME (IEEE Int'l Conf. on Software Maintenance and Evolution), Adelaide, Australia

28 September

- SecDev (IEEE Secure Development), Atlanta, USA

OCTOBER

3 October

- PACT (Int'l Conf. on Parallel Architectures and Compilation Techniques), Atlanta, USA

5 October

- EDOC (IEEE Int'l Enterprise Distributed Object Computing Conf.), Eindhoven, the

Netherlands

- ICSE (IEEE/ACM Int'l Conf. on Software Eng.), Seoul, South Korea

12 October

- ISSRE (IEEE Int'l Symposium on Software Reliability Eng.), Coimbra, Portugal

16 October

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

17 October

- MICRO (IEEE/ACM Int'l Symposium on Microarchitecture), Athens, Greece

18 October

- ICCD (IEEE Int'l Conf. on Computer Design), Hartford, USA
- MODELS (ACM/IEEE Int'l Conf. on Model-Driven Eng. Languages and Systems), Montreal, Canada

19 October

- DFT (IEEE Int'l Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems), Frascati, Italy

21 October

- FIE (IEEE Frontiers in Education Conf.), Uppsala, Sweden

25 October

- VIS (IEEE Visualization Conf.), Salt Lake City, USA

NOVEMBER

2 November

- ICFEC (IEEE Int'l Conf. on Fog and Edge Computing), Melbourne, Australia

6 November

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

8 November

- NAS (IEEE Int'l Conf. on Networking, Architecture, and Storage), Riverside, USA

9 November

- ICTAI (IEEE Int'l Conf. on Tools with Artificial Intelligence), Baltimore, USA
- ISMVL (IEEE Int'l Symposium on Multiple-Valued Logic), Miyazaki, Japan

15 November

- SC, Atlanta, USA

16 November

- FOCS (IEEE Symposium on Foundations of Computer Science), Durham, USA
- LCN (IEEE Conf. on Local Computer Networks), Sydney, Australia

30 November

- ICHI (IEEE Int'l Conf. on Healthcare Informatics), Oldenburg, Germany

DECEMBER

9 December

- CC (IEEE Int'l Conf. on Conversational Computing), Irvine, USA
- AIKE (IEEE Int'l Conf. on Artificial Intelligence and Knowledge Eng.), Irvine, USA

14 December

- BCD (IEEE/ACIS Int'l Conf. on Big Data, Cloud Computing, and Data Science Eng.), Macao
- CloudCom (IEEE Int'l Conf. on Cloud Computing Technology and Science), Bangkok, Thailand

16 December

- BIBM (IEEE Int'l Conf. on Bioinformatics and Biomedicine), Seoul, South Korea