

# COMPUTING edge

- Data Engineering
- Cloud Computing
- Quantum Computing
- History



MARCH 2022

[www.computer.org](http://www.computer.org)





# CALL FOR SPECIAL ISSUE PROPOSALS

*Computer* solicits special issue proposals from leaders and experts from a broad range of computing communities. Proposed themes/issues should address timely, emerging topics that will be of broad interest to *Computer's* readership. Special issues are an important component of *Computer*, as they deliver essential research insights and well-developed perspectives on new and established technologies and computing strategies.

We encourage submissions of high-quality proposals for the 2023 editorial calendar. Of particular interest are proposals centered on:

- offsite educational and business continuity technology challenges,
- privacy related to personal location tracking and surveillance (digital and physical),
- artificial intelligence and machine learning,
- technology's role in disrupted supply chains,
- misinformation and disinformation (fake information—malicious or non-malicious), and
- cyberwarfare/cyberterrorism

Proposal guidelines are available at:

[www.computer.org/csdl/magazine/co/write-for-us/15911](http://www.computer.org/csdl/magazine/co/write-for-us/15911)



## STAFF

### Editor

Cathy Martin

### Publications Portfolio Managers

Carrie Clark, Kimberly Sperka

### Senior Advertising Coordinator

Debbie Sims

### Production & Design Artist

Carmen Flores-Garvey

### Publisher

Robin Baldwin

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

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

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

**Reuse Rights and Reprint Permissions:** Educational or personal use of this material is permitted without fee, provided such use: 1) is not made for profit; 2) includes this notice and a full citation to the original work on the first page of the copy; and 3) does not imply IEEE endorsement of any third-party products or services. Authors and their companies are permitted to post the accepted version of IEEE-copyrighted material on their own Web servers without permission, provided that the IEEE copyright notice and a full citation to the original work appear on the first screen of the posted copy. An accepted manuscript is a version which has been revised by the author to incorporate review suggestions, but not the published version with copy-editing, proofreading, and formatting added by IEEE. For more information, please go to: [http://www.ieee.org/publications\\_standards/publications](http://www.ieee.org/publications_standards/publications)

[/rights/paperversionpolicy.html](#). Permission to reprint/republish this material for commercial, advertising, or promotional purposes or for creating new collective works for resale or redistribution must be obtained from IEEE by writing to the IEEE Intellectual Property Rights Office, 445 Hoes Lane, Piscataway, NJ 08854-4141 or [pubs-permissions@ieee.org](mailto:pubs-permissions@ieee.org). Copyright © 2022 IEEE. All rights reserved.

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

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

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

## IEEE Computer Society Magazine Editors in Chief

### Computer

Jeff Voas, *NIST*

### Computing in Science & Engineering

Lorena A. Barba, *George Washington University*

### IEEE Annals of the History of Computing

Gerardo Con Diaz, *University of California, Davis*

### IEEE Computer Graphics and Applications

Torsten Möller, *Universität Wien*

### IEEE Intelligent Systems

Longbing Cao, *University of Technology Sydney*

### IEEE Internet Computing

George Pallis, *University of Cyprus*

### IEEE Micro

Lizy Kurian John, *University of Texas at Austin*

### IEEE MultiMedia

Shu-Ching Chen, *Florida International University*

### IEEE Pervasive Computing

Marc Langheinrich, *Università della Svizzera italiana*

### IEEE Security & Privacy

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


### IEEE Software

Ipek Ozkaya, *Software Engineering Institute*

### IT Professional

Irena Bojanova, *NIST*

# COMPUTING edge



# 12

Validating  
“Data  
Validation”

# 18

The Next Wave  
in Cloud Systems  
Architecture

# 24

Interhost  
Orchestration  
Platform  
Architecture for  
Ultrascale Cloud  
Applications

## Data Engineering

- 8 Data Labeling for the Artificial Intelligence Industry:  
Economic Impacts in Developing Countries

NIR KSHETRI

- 12 Validating “Data Validation”

NORITA AHMAD AND KEVIN DIAS

## Cloud Computing

- 18 The Next Wave in Cloud Systems Architecture

AMIN VAHDAT AND DEJAN MILOJICIC

- 24 Interhost Orchestration Platform Architecture for  
Ultrascale Cloud Applications

SASKO RISTOV, THOMAS FAHRINGER, RADU PRODAN, MAGDALENA  
KOSTOSKA, MARJAN GUSEV, AND SCHAHRAM DUSTDAR

## Quantum Computing

- 33 Emerging Technologies for Quantum Computing

JONATHAN M. BAKER AND FREDERIC T. CHONG

- 40 Quantum Computing

JOSE LUIS HEVIA, GUIDO PETERSSEN, CHRISTOF EBERT, AND  
MARIO PIATTINI

## History

- 48 Preserving the Past by Industry Participants

JAMES W. CORTADA

- 51 History of the Marching Cubes Algorithm

WILLIAM E. LORENSEN

## Departments

- 4 Magazine Roundup  
7 Editor’s Note: The People behind the Data  
58 Conference Calendar

Subscribe to *ComputingEdge* for free at  
[www.computer.org/computingedge](http://www.computer.org/computingedge).

# Magazine Roundup

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

## Computer

### ***When Scientific Software Meets Software Engineering***

The authors of this article from the December 2021 issue of *Computer* investigate the different levels of abstraction, linked to the diverse artifacts of the scientific software development process, that a software language can propose and the validation and verification facilities associated with the corresponding level of abstraction the language can provide to the user.

## Computing

### ***Reproducibility Practice in High-Performance Computing: Community Survey Results***

In a recent survey of the SC conference community, the authors of this article from the September/October 2021 issue of *Computing in Science & Engineering* collected information about the SC reproducibility initiative activities. Results show that the reproducibility initiative activities have contributed to higher levels of awareness on the part of SC conference

technical program participants, and hint at contributing to greater scientific impact for the published papers of the SC conference series. Stringent point-of-manuscript-submission verification is problematic, as are inherent difficulties of computational reproducibility in HPC. Future efforts should better decouple the community educational goals from goals that specifically strengthen a research work's potential for long-term impact through reuse 5 to 10 years down the road.

## IEEE Annals

### ***Leonardo Torres Quevedo: Pioneer of Computing, Automatics, and Artificial Intelligence***

The search for little-known pioneers of automation, missed fathers or mothers of the computer, unfamiliar founders of computer science, or unrecognized creators of artificial intelligence invites us to look at the past with an open mind. In this article from the July–September 2021 issue of *IEEE Annals of the History of Computing*, the authors provide a comprehensive examination of the

almost unknown contribution of a Spanish pioneer in all those fields during the first two decades of the 20th century, the engineer and mathematician Leonardo Torres Quevedo (1852–1936).

## IEEE Computer Graphics and Applications

### ***DeepGD: A Deep Learning Framework for Graph Drawing Using GNN***

Many graph drawing techniques have been proposed for generating aesthetically pleasing graph layouts, but it remains a challenging task since different layout methods tend to highlight different characteristics of the graphs. Recently, studies on deep learning-based graph drawing algorithms have emerged, but they are often not generalizable to arbitrary graphs without retraining. In this article from the September/October 2021 issue of *IEEE Computer Graphics and Applications*, the authors propose a convolutional graph neural network-based deep learning framework, DeepGD, which can draw arbitrary graphs once trained. It attempts to generate layouts by compromising among multiple prespecified



aesthetics, considering a good graph layout usually complies with multiple aesthetics simultaneously. To balance the tradeoff, they propose two adaptive training strategies, which adjust the weight factor of each aesthetic dynamically during training.

## IEEE Intelligent Systems

### **An Emotional Recommender System for Music**

---

Recommender systems have become essential to users for finding “what they need” within large collections of items. Meanwhile, recent studies have demonstrated that user personality can effectively provide more valuable information to significantly improve recommenders’ performance, especially considering behavioral data captured from social network logs. In this article from the September/October 2021 issue of *IEEE Intelligent Systems*, the authors describe a novel music recommendation technique based on the identification of personality traits, moods, and emotions of a single user, starting from solid psychological observations recognized by the analysis of user behavior within a social environment. Users’ personality and mood have been embedded within a content-based filtering approach to obtain more accurate and dynamic results.

## IEEE Internet Computing

### **Future of HPC: The Internet of Workflows**

---

Driven by convergence with artificial intelligence and data analytics, increased heterogeneity, and a hybrid cloud/on-premises delivery model, dynamic composition of workflows will be a key design criteria of future high-performance computing (HPC) systems. While tightly coupled HPC workloads will continue to execute on dedicated supercomputers, other jobs will run elsewhere, including public clouds and at the edge. Connecting these distributed computing tasks into coherent applications that can perform at scale is what the authors of this article from the September/October 2021 issue of *IEEE Internet Computing* call “Internet of Workflows.”

## IEEE micro

### **The Birth of the Microprocessor**

---

The story starts in February 1968 when the author of this article from the November/December 2021 issue of *IEEE Micro* joined Fairchild Semiconductor R&D Lab in Palo Alto, CA, USA. At that time, nearly all the integrated circuits in production used bipolar technology. Metal-oxide-semiconductor

(MOS) technology was up and coming but still had many problems; it was extremely slow and was not yet reliable. The author believed that the future of digital electronics belonged to MOS technology because one could integrate in the same silicon area ten times more logic gates with half the number of manufacturing steps required by bipolar technology.

## IEEE MultiMedia

### **EGGAN: Learning Latent Space for Fine-Grained Expression Manipulation**

---

Fine-grained facial expression aims at changing the expression of an image without altering facial identity. Most current expression manipulation methods are based on a discrete expression label, which mainly manipulates holistic expression with details neglected. To handle the above-mentioned problems, the authors of this article from the July–September 2021 issue of *IEEE MultiMedia* propose an end-to-end expression-guided generative adversarial network (EGGAN), which synthesizes an image with expected expression given continuous expression label and structured latent code. An adversarial autoencoder is used to translate a source image into a structured latent space. The

encoded latent code and the target expression label are input to a conditional GAN to synthesize an image with the target expression. A perceptual loss and a multiscale structural similarity loss are introduced to preserve facial identity and global shape during expression manipulation.



### **Electronic Textile Gaia: Ubiquitous Computational Substrates Across Geometric Scales**

From in-body implantables to geotextiles and large-area spacecraft blankets, electronic fabric is now poised to operate across geometric scales that span many orders of magnitude, and thus across operational contexts with divergent material resiliency requirements—reaching far beyond the wearable device regime that is typically considered. This article from the July–September 2021 issue of *IEEE Pervasive Computing* reviews the key technical trends and lingering hurdles that are relevant to using functional fibers and e-textiles for operating at disparate scales. The authors focus on leveraging the unique material properties of a textile and the miniaturization of electronic devices in concert with the revolution in mass-manufacturing and digital fabrication technologies used to customize the device at the level of polymer, fiber, fabric, three-dimensional form, and system.



### **A Research Road Map for Building Secure and Resilient Software-Intensive Systems**

Poor software engineering processes can result in insecure and brittle software-intensive systems. A new US agenda addresses this by advancing development and architectural paradigms, and by providing concrete research and development recommendations. The authors of this article from the November/December 2021 issue of *IEEE Security & Privacy* propose that the security community work closely with the software engineering community to realize secure, resilient software-intensive systems.



### **Managing Technical Debt Under Uncertainty**

Managing technical debt in the presence of uncertainty is a fundamental problem and an open challenge. The authors of this article from the November/December 2021 issue of *IEEE Software* provide an overview of the research, practical limitations, and future study of the current state of technical debt management under uncertainty.



### **A Process Model for Service- Oriented Development of Embedded Software Systems**

The concepts of service-oriented computing (SOC) have previously

been used in the embedded systems domain, mostly at the device level in ad hoc manners to achieve advantages of device integration. However, SOC has not been used for the development of embedded software systems (ESS), i.e., software controlling the embedded devices. To fill this gap, a process model is proposed in this article from the September/October 2021 issue of *IT Professional* that allows the systematic development of ESS based on SOC concepts. The proposed process consists of analysis and design phases of embedded software development. The analysis phase is concerned with the collection of system information and preparation for system design. Based on this, the service-based software architecture is developed in the design phase. The effectiveness of the proposed process model is demonstrated through its application in a smart home case study. 🌐

Join the IEEE  
Computer  
Society  
[computer.org/join](https://computer.org/join)







## Editor's Note

# The People behind the Data

**D**ata engineers are responsible for collecting and preparing the data that we rely on in myriad industries and contexts. Despite advances in data analysis software, humans often perform better than algorithms at certain tasks, such as data labeling and validation. This *ComputingEdge* issue reflects on the essential role of data engineers and highlights developments in the field.

*IT Professional's* "Data Labeling for the Artificial Intelligence Industry: Economic Impacts in Developing Countries" describes the labor-intensive work of data labelers and how this new type of job has grown in countries like China and India. *Computer's* "Validating 'Data Validation'" proposes a dashboard as a means for an experienced data

practitioner to monitor automated processes to ensure data quality and accuracy.

Large amounts of data are stored and processed in the cloud. In *Computer's* "The Next Wave in Cloud Systems Architecture," experts discuss emerging data-hosting approaches such as hybrid, multicloud, and edge. In *IEEE Internet Computing's* "Interhost Orchestration Platform Architecture for Ultrascale Cloud Applications," the authors present an intelligent page management method to reduce memory utilization and access time in cloud datacenters.

Both hardware and software will need to adapt to quantum computing. In *IEEE Micro's* "Emerging Technologies for Quantum Computing," the authors aim to improve quantum hardware

scalability by evaluating new technologies such as superconducting resonant cavities and neutral atoms. The authors of *IEEE Software's* "Quantum Computing" argue that software development methodologies must evolve to address a future with quantum technology.

Finally, this *ComputingEdge* issue features two articles about the history of computing. In "Preserving the Past by Industry Participants," from *IEEE Annals of the History of Computing*, the author encourages practitioners in industry to preserve materials related to computing history. "History of the Marching Cubes Algorithm," from *IEEE Computer Graphics and Applications*, recounts the creation of an influential and widely used algorithm in computer graphics. 🎮

# Data Labeling for the Artificial Intelligence Industry: Economic Impacts in Developing Countries

Nir Kshetri, *University of North Carolina at Greensboro*

**A**rtificial intelligence (AI) applications, although at a nascent stage of development, are already having considerable economic and social impacts in developing countries.<sup>1</sup> An equally fascinating and important aspect of the development of the global AI industry is that a high proportion of jobs that require relatively low skills in this industry are being performed by the developing countries. The multibillion-dollar data labeling industry represents an interesting example to illustrate this trend.

According to the professional services firm PwC, AI's contribution to the global economy will reach about US\$16 trillion by 2030 (<https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html>). Accurately labeling data for training machine learning (ML) models is integral to the creation of this value. The global market for data labeling, also known as content labeling or data annotation, is expected to reach US\$5 billion by 2023.<sup>2</sup>

Data labeling activities, however, are extremely time and labor intensive. Developing countries are in a position to take advantage of the opportunities provided by the global AI industry. This phenomenon has thus created a whole new industry of data labeling in developing countries, which is described as "a new type of blue-collar industry" (<https://www.forbes.com/sites/korihale/2019/05/28/google-microsoft-banking-on-africas-ai-labeling-workforce/#34b0fd0d541c>). Data labelers are referred to as the blue-collar workers of the AI age (<https://towardsdatascience.com/the-invisible-workers-of-the-ai-era-c83735481ba>).

## DATA LABELING IN DEVELOPING COUNTRIES POWERING THE GLOBAL AI INDUSTRY

Data need to be cleaned, categorized into appropriate groups, and labeled so that AI algorithms can

recognize patterns.<sup>3</sup> For instance, for ML algorithms to accurately recognize a car, the algorithms may need to be trained with a large number of car pictures ([https://www.vice.com/en\\_us/article/7xyabb/china-ai-dominance-relies-on-young-data-labelers](https://www.vice.com/en_us/article/7xyabb/china-ai-dominance-relies-on-young-data-labelers)). In the most common form of AI—the supervised learning—the algorithms need to be fed with millions of pretagged examples of car pictures until they correctly identify the pictures.<sup>2</sup> These activities need a large amount of human labor. For instance, one hour of video data related to autonomous driving may need as much as 800 man-hours of data labeling work (<https://www.axios.com/ai-data-labeling-billion-dollar-market-409704bc-e63c-4af0-b0d0-44424abcd561.html>). Estimates suggest that data labeling activities account for as much as 80% of the time needed to build AI systems.<sup>4</sup>

Labeling data for some AI apps may need high levels of skills and knowledge. For instance, to develop an AI app to detect cancer on images from a CT scan, experienced radiologists may have to train the algorithms (<https://towardsdatascience.com/the-invisible-workers-of-the-ai-era-c83735481ba>). However, there are many tasks that computers lack the capability to perform as well as ordinary human beings. Even less-skilled workers can easily be trained to perform such tasks. Most data labeling trainings can be completed in a short period of time. For instance, in order to learn their tasks, data labelers at iMerit typically take a one-week online training course via video calls with U.S.-based trainers.<sup>4</sup>

Table 1 provides some examples of data labeling companies operating in developing economies. These firms mainly serve foreign clients. Chinese data labeling firms such as MBH, on the other hand, mainly support the domestic AI industry.

In China, which has gained global prominence in recent years in the AI field, research and development activities to support the growth of the AI industry are conducted in wealthy cities such as Beijing, Shanghai, Hangzhou, and Shenzhen. Most of the data labeling

**TABLE 1.** Some Examples of Data Labeling Companies Operating in Developing Economies.

Data labeling company	Operations and workforce	Profiles of clients
iMerit	Based in India and the U.S. 2500 in data labeling centers in India such as Ranchi, Shillong, Vizag, and Kolkata ( <a href="https://tinyurl.com/y465cqmj">https://tinyurl.com/y465cqmj</a> ).	90% are U.S.-based ( <a href="https://tinyurl.com/yyzn8vjd">https://tinyurl.com/yyzn8vjd</a> ).
Samasource	Offices in San Francisco, New York, the Hague, Kenya, and Uganda. Global staff of 2900: East Africa's largest AI and data annotation employer ( <a href="https://tinyurl.com/wqanmja">https://tinyurl.com/wqanmja</a> ).	Include 25% of the Fortune 50 companies including major automakers and U.S.-based technology companies such as Google, Microsoft, and IBM ( <a href="https://tinyurl.com/qnwlxgo">https://tinyurl.com/qnwlxgo</a> ).
MBH	300,000 in China's backward provinces ( <a href="https://tinyurl.com/yeb7qtbj">https://tinyurl.com/yeb7qtbj</a> ).	Mainly Chinese companies such as the Beijing-based video-sharing social networking platform TikTok.
Playment	Based in India and the U.S. More than 300 000 crowdsourced data labelers ( <a href="https://tinyurl.com/y2dqm3c3">https://tinyurl.com/y2dqm3c3</a> ).	Over 100 customers in more than 12 countries ( <a href="https://tinyurl.com/upp4wsu">https://tinyurl.com/upp4wsu</a> ). Some include Samsung, Didi Chuxing Technology, Alibaba, Drive.ai, and Continental AG. Most clients are in the autonomous vehicle industry ( <a href="https://tinyurl.com/y2dqm3c3">https://tinyurl.com/y2dqm3c3</a> ).

works, on the other hand, are performed in the country's disadvantaged regions such as smaller towns and rural areas in Shandong, Henan, Hebei, and Shanxi provinces in the North.<sup>5</sup> In particular, Henan Province is the epitome of an economy that has experienced an AI-led transformation. The Province's Zhengzhou city has been known for the manufacturing plants of Taiwanese electronics company Foxconn. It was estimated that Foxconn employed about 350 000 people and produced about half of the world's iPhones in Zhengzhou in 2016 (<https://www.businessinsider.com/apple-iphone-factory-foxconn-china-photos-tour-2018-5>).

The growth in data labeling also reflects the influences of a number of trends in the manufacturing sector. First, manufacturing activities have become increasingly automatized. For instance, between 2012 and 2016, Foxconn was reported to deploy tens of thousands of robots, which replaced more than 400 000 jobs. The company's plan includes full automation of 30% of production activities by 2020 (<https://www.scmp.com/economy/china-economy/article/2185993/man-vs-machine-chinas-workforce-starting-feel-strain-threat>).

Second, manufacturing jobs in China are declining due to a decline in the country's economic growth and, hence, demand for products, increasing costs, and competition from other economies (<https://www.nytimes.com/2016/07/23/business/international/china-jobs-donald-trump.html>). For instance, the Boston Consulting Group's study conducted in 2015 found that manufacturing costs in some of China's major

manufacturing hubs were almost at the same levels as in the U.S. Unsurprisingly, many western companies have been moving and relocating manufacturing activities into other developing countries in Asia as well as to the U.S., Canada, and Mexico (<https://www.nytimes.com/2016/07/23/business/international/china-jobs-donald-trump.html>). For instance, factory workers in Bangladesh and Vietnam can be hired for less than a quarter and a half, respectively, of the salary of a Chinese worker.

Finally, there has been a generational shift in preference for work. An increasing number of Chinese millennials do not like dull, boring, and tedious factory jobs (<https://www.latimes.com/world/la-fg-china-millennials-jobs-20190512-story.html>).

The upshot of the above is a decline in industrial manufacturing and exports in China. For instance, Henan's mobile phone exports reduced by 23.7% in January 2019 compared to a year earlier (<https://www.scmp.com/economy/china-economy/article/2188162/foxconn-tale-slashed-salaries-disappearing-benefits-and-mass>). China's economically backward regions are embarking on ambitious plans to develop the data labeling industry. For instance, North China's landlocked province Shanxi, which is among the poorest provinces in China, plans to attract more than 100 data labeling companies and train more than 10 000 workers by 2022. The province's goal is to have a RMB 5 billion (US\$726 million) industry by 2025 (<https://kr-asia.com/data-labeling-jobs-are-coming-to-underdeveloped-regions-in-china-but-can-they-stay>).

Workers who were employed in assembly lines and construction sites before are thus finding new jobs in the Chinese data labeling industry.<sup>6</sup>

Partly as a response to the decline in manufacturing activities, in recent years, data labeling companies are mushrooming in the towns and villages of backward regions such as the Henan province.<sup>7</sup> In the Pingdingshan village of the province, some large data labeling projects are reported to employ tens of thousands of people.<sup>8</sup>

### ECONOMIC AND SOCIAL IMPACTS OF THE DATA LABELING INDUSTRY

Finding high quality AI talents such as ML engineers has been big challenge for companies in developing economies (<https://factordaily.com/india-ai-talent-gap-google-microsoft/>). For instance, India is estimated to have only about 50-75 AI researchers (<https://via.news/asia/artificial-intelligence-development-india/>). According to India's skill assessment company Aspiring Minds' Annual Employability Survey 2019, 80% of Indian engineers were considered to be unfit for a job. The survey also found that only 2.5% of them had skills required by the AI industry (<https://www.businesstoday.in/current/corporate/indian-engineers-tech-jobs-survey-80-per-cent-of-indian-engineers-not-fit-for-jobs-says-survey/story/330869.html>). India has also faced a severe shortage of qualified faculty members to teach AI courses in its universities.<sup>9</sup>

On the other hand, there is an abundant supply of low-skill and low-wage labors in India and other developing countries. For instance, Indian high schools graduated 20 million students in 2017.<sup>10</sup> There are, however, only a few job opportunities available to absorb these graduates. As an example, in a Mumbai city police's job advertisement for 1137 constable positions with US\$357/month salary, which required only a high school education, over 200 000 people applied. Many of the candidates had been trained in highly skilled jobs such as doctors, lawyers, and engineers.<sup>11</sup>

In light of the high rates of unemployment and underemployment in developing countries, the economic impacts of the newly emerging data labeling industry are undoubtedly important. However, controversies have arisen regarding ethical and fair practices of data labeling firms.

Some data labeling firms have been accused of paying low wages to their workers. In this way, organizations engaged in this industry are allegedly facilitating a "new kind of slavery in the digital era" (<https://www.weforum.org/agenda/2019/08/ai-low-skilled-workers/>).

Although some data labeling firms have claimed to produce positive social effects, such claims cannot be easily verified. There are virtually no regulations that govern the working conditions of data labelers. Industry standards related to ethical sourcing are weak. There is no standard for reporting data-labeling firms' social impacts. There is also the lack of validation of such impacts by third parties. There is little hard evidence to counter critics' concerns that these firms claiming that they engage in impact sourcing is nothing more than marketing gimmicks or "impact-washing."<sup>3</sup>

Some data labeling firms such as CloudFactory, DDD, iMerit, and Samasource are members of the Global Impact Sourcing Coalition (GISC), which was founded in 2016. Among other measures, the GISC has established an "impact sourcing standard." The standard defines minimum requirements that the GISC members are to satisfy. It has also provided voluntary best practices for employment. The GISC requires its members' performance on criteria such as nondiscrimination and equal pay to be assessed every two years.<sup>3</sup>

When it comes to promoting ethical and socially responsible behaviors, however, the GISC at best is of questionable effectiveness and value. For instance, the violators face no penalty or sanction if they do not pass the assessment. They do not lose the GISC membership. The GISC members also differ significantly in the terms of the information they publish. Samasource's impact audits report includes indicators such as workforce demographics and the number of people lifted from poverty. The data labeling firm DDD's reports contain information about employees' earnings and increase in lifetime income. CloudFactory had not published social impact report since 2015. As of 2019, iMerit had not published any such reports.<sup>3</sup> The U.S.-based provider of data labeling and annotation services Alegion, which is not a GISC member, has outlined broad targets that it seeks to achieve but lacks specific metrics.<sup>3</sup>

Another challenge that arises here is that unlike fair-trade goods that are directly sold to consumers, data labeling firms provide their services to businesses. These firms thus face little public pressure to be honest, and it is ineffective to pressure them to engage in ethical practices.<sup>3</sup>

### SUMMARY

The rapidly growing global AI industry has created demands for highly skilled jobs such as ML engineers and data scientists, as well as less-skilled works such as those of data labelers. In particular, most AI

systems heavily rely on human-powered data labeling activities. Developing countries provide very large and seemingly bottomless sources and resources to support these activities to boost the global AI industry. Data labelers in these countries are playing a key role in curating the data that powers AI systems.

A notable feature of the data labeling industry is that it does not favor locations with specific cultural contexts. The data labeling market thus is characterized by a low entry barrier for most developing countries. Whereas the outsourcing of call center jobs gravitated to countries with sizable English-speaking populations such as India and the Philippines, English skill is less of a factor in data labeling jobs. Digital literacy is sufficient to participate in most data labeling tasks such as image classification.

Among developing countries, China has emerged as a key global AI player. The country's wealthy regions and big cities lack attractiveness for the development of data-labeling services industry. Such services in the country are thus mostly performed in the poorer and backward regions, which are providing economic incentives for data labeling firms.

What is not clear is whether data labeling firms in developing countries are operating in more or less ethical ways compared to other foreign firms operating in these countries. Some critics have claimed that this industry has features that are akin to slavery. While data labeling firms claim to engage in activities that have positive social impacts, it is not easy to assess the validity of such claims. Data labeling companies have their own definitions as to what are ethical and fair practices. Moreover, the definitions vary widely across them. Thus, we may not be able to take self-reported information provided by data labeling firms as proof positive that these firms are creating more positive social impacts in developing countries compared to other foreign firms. In many cases, the problem of assessing such claims is made more complex by the fact that they do not publish any information or fail to provide relevant information on such impacts. 🤖

## REFERENCES

1. N. Kshetri, "Artificial intelligence in developing countries," *IEEE IT Pro.*, to be published, doi: 10.1109/MITP.2019.2951851.
2. "Data-labelling startups want to help improve corporate AI," *The Economist*, Oct. 17, 2019. [Online]. Available: <https://www.economist.com/business/2019/10/17/data-labelling-startups-want-to-help-improve-corporate-ai>
3. K. Kaye, "These companies claim to provide "Fair-trade" data work. Do they?" *MIT Technol. Rev.*, Aug. 7, 2019. [Online]. Available: <https://www.technologyreview.com/s/614070/cloudfactory-ddd-samasource-imerit-impact-sourcing-companies-for-data-annotation/>
4. C. Metz, "A.I. is learning from humans. Many humans," *The New York Times*, Aug. 16, 2019. [Online]. Available: <https://www.nytimes.com/2019/08/16/technology/ai-humans.html?auth=linked-google>
5. S. Dai, "AI promises jobs revolution but first it needs old-fashioned manual labour—from China," *South China Morning Post*, Oct. 8, 2018. [Online]. Available: <https://www.scmp.com/tech/article/2166655/ai-promises-jobs-revolution-first-it-needs-old-fashioned-manual-labour-china>
6. L. Yuan, "How cheap labor drives China's A.I. ambition," *The New York Times*, Nov. 25, 2018.
7. H. Wu, "China Is achieving AI dominance by relying on young blue-collar workers," *Vice*, Dec. 21, 2018. [Online]. Available: [https://www.vice.com/en\\_us/article/7xyabb/china-ai-dominance-relies-on-young-data-labelers](https://www.vice.com/en_us/article/7xyabb/china-ai-dominance-relies-on-young-data-labelers)
8. C. Cadell, "Faces for cookware: Data collection industry flourishes as China pursues AI ambitions," *Reuters*, Jun. 28, 2019. [Online]. Available: <https://www.reuters.com/article/us-china-ai-data-insight/faces-for-cookware-data-collection-industry-flourishes-as-china-pursues-ai-ambitions-idUSKCN1TS3EA>
9. S. Ravi and P. Nagaraj, "Harnessing the future of AI in India," *Brookings Inst.*, Oct. 18, 2018. [Online]. Available: <https://www.brookings.edu/research/harnessing-the-future-of-ai-in-india/>
10. K. Sheehy, "High school grads in China, India are better prepared for college," *U.S. News.*, Aug. 27, 2012. [Online]. Available: <https://www.usnews.com/education/blogs/high-school-notes/2012/08/27/high-school-grads-in-china-india-are-better-prepared-for-college>
11. N. Bagri, "India is trapping its young people," *Foreign Policy*, May 14, 2019. [Online]. Available: <https://foreignpolicy.com/2019/05/14/india-is-trapping-its-young-people/>

**NIR KSHETRI** is currently a Professor of management with the Bryan School of Business and Economics, University of North Carolina at Greensboro. Contact him at [nbkshetr@uncg.edu](mailto:nbkshetr@uncg.edu).

## DEPARTMENT: DATA

# Validating “Data Validation”

Norita Ahmad, *American University of Sharjah*

Kevin Dias, *MRM*

*We have more access to data than ever before, but it is difficult to make sense of it when the data is incomplete or inaccurate. As such, data validation is necessary to increase data quality for effective decision making.*

Data validation is an essential part of data handling, notably in the fields of data science, artificial intelligence, and the Internet of Things. As the name suggests, it is the checking of quality and accuracy of data prior to processing it. Data validation is not a new topic; however, most data handling processes do not follow a systematized data validation approach and hence lead to small- or large-scale error-prone conclusions.<sup>1-3</sup> The use of a proper data validation process could prevent life-threatening scenarios such as incorrectly classifying a cancerous tissue as benign,<sup>1</sup> prevent the loss of millions of dollars by ensuring the accuracy of a price prediction model,<sup>2</sup> or prevent a fatal accident of a self-driving car that failed to recognize a jaywalking pedestrian.<sup>3</sup> As such, there is an urgent need for efficient and effective data validation procedures.

Today, with the advances in big data and analytics, businesses, governments, and individuals can apply diverse data mining and machine-learning algorithms to bring new business opportunities and improve quality of life. The International Data Corporation predicted that, by 2025, data will grow to 175 trillion gigabytes worldwide, and businesses will become even more reliant on big data for decision making.<sup>4</sup> However, given the nature of big data (that is, huge volume of generated data, fast velocity of arriving data, and large variety of heterogeneous data), the quality of data can easily be compromised.<sup>5</sup> There are anecdotes where collected data was so voluminous

that people eventually discarded it since it derived no benefit.<sup>6</sup>

### WHAT IS DATA VALIDATION?

There are different definitions of data validation but, in general, data validation is a process of delivering clean and accurate data to specific programs, applications, and services.<sup>7</sup> For example, in machine learning, data validation means checking the quality and accuracy of source data before training a new model.<sup>8</sup> Different types of validation can be performed depending on the objectives and constraints of a given data set.


Data validation usually occurs during the transform stage of the extract, transform, and load (ETL) data process, where data are first extracted from a data source (Stage 1); validated, cleaned, merged, formatted, and/or appended with other data set extracts (Stage 2); and finally ready to load (Stage 3) and process as per the given use case.<sup>9</sup>

### WHY DATA VALIDATION?

Depending on a specific industry, there are numerous reasons why data validation is critical to data-driven projects. The two most important reasons that are often taken too lightly are: 1) early detection of errors and 2) the cost of time saved.<sup>5,8,10</sup> These two go hand in hand, but it is important to highlight them independently as various decision makers and data experts tend to exclude data validation due to a lack of knowledge on the many consequences from each of the two.

### Early detection of errors

Validating details of data are necessary to mitigate any project defects. Given that businesses rely on



high-quality data to make critical decisions, they run the risk of basing decisions on data that are not accurately representative of the situation at hand if data validation is not properly performed. According to Gartner, on average, the financial impact of poor data quality on organizations is US\$9.7 million per year.<sup>2</sup>

### Cost of time saved

It is without question that data scientists, analysts, and engineers are in one of the highest paid lines of work. The reported median annual wage for these positions in 2020 was higher than the median annual wage for all other occupations.<sup>11</sup> However, it was reported that one out of three data analysts spend over 40% of their time “vetting and validating their analytics data before it can be used for strategic decision-making.”<sup>12</sup> In the past, data prep tasks have occupied around 80% of a data scientist’s time—challenging the wisdom of asking highly paid data scientists to spend most of their time preparing data instead of using them for the actual analysis and decision making. Imagine the value of time saved if there is a solution to this efficiency gap.

### HOW IS DATA VALIDATION PERFORMED?

The most straightforward rules used in data validation are rules that ensure data integrity, for example, the correct format to enter a phone number or the minimum password length. These basic data validation rules help to uphold standards that will effectively make working with data more efficient. During the data validation process, it is important that the standards and structure of the data model is also well understood. Structured data are data that has been formatted into a well-defined data model while unstructured data are data in a raw form. There is also semistructured data that fall in between structured and unstructured data.<sup>13</sup> Understanding the difference between these types of data will help in

maintaining compatibility with applications and other data sets with which data are integrated and stored.

There are many methods available for data validation. The most common methods are using script and software programs. Writing a script may be an option for those who are fluent in coding languages such as Python and R. Script allows for a better interpretation of the data, such as comparing data values and structure against predefined rules and verifying all the necessary information within the required quality parameters. This method can however be very time consuming depending on the complexity and size of the data set. Today, there are many commercial software programs that can be used to perform data validation. For most

---

*DIFFERENT TYPES OF VALIDATION CAN BE PERFORMED DEPENDING ON THE OBJECTIVES AND CONSTRAINTS OF A GIVEN DATA SET.*

---

people, this is the preferred method since the program has been developed to understand the common rules and file structures used in data validation in the industry. No in-depth understanding of the underlying format is required from the user. However, this method can be a bit costly. As an alternative, a more technical user would opt for an open source software such as SourceForge and OpenRefine because they are more cost-effective.

Another important aspect of data validation is the evaluation of the validity of range measurements. The top use case is in satellites or deep-space exploration systems such as those used by NASA. Given that the system state estimates can be altered significantly by erroneous sensor data, an algorithm which decides whether to assimilate or reject data are required. Algorithms such as least squares, linear combinations, Bayesian, or Kalman filtering can be used to select

sensors that are consistent with the computed estimates and reject those that are far from the predicted values.<sup>14,15</sup>

## WHAT ARE THE CHALLENGES?

Let's face it, data validation is by far the least exciting data-related buzzword that exists. This is a fundamental failure in both industry as well as academia. The only people you see actively talking about it are veterans in the field who have firsthand experience in realizing its importance too late. Furthermore, as of the writing of this column, there is a surplus of data practitioners dominated by millennials, with a demand from employers who have a flawed view of recruiting based on quantity and not quality. It is only being part of a "trend" for 20-year-olds to hone in on their ability to build predictive models rather than

---

*DECISION MAKERS NEED TO HAVE  
FUNDAMENTAL STANDARDS FOR  
QUALITY AND RELIABLE DATA  
BECAUSE THEY ARE CRITICAL FOR  
CORPORATE SURVIVAL.*

---

verify the credibility of their data sources.<sup>5</sup> Additionally, the majority of data analytics teams today is so time constrained in delivering "quick insights" that data validation is unfortunately classified as a waste of time. To further stress on the broad lack of love for data validation in the industry, there are over a hundred different data science/analysis certifications available but how many data validation certifications are there?

The ETL processes are inherently so focused on "getting the job done" with speed and efficiency, with accuracy being a low priority. In the top-down team structures where there is a cascading effect of pressure, the data analyst doing the heavy lifting is left with a difficult choice between spending time sense checking their data or keeping their job. Data governance and data ethics-related roles such as the chief data officer are only recently being birthed into the industry and will hopefully bring more emphasis on the importance of allocating engineering hours specifically to data validation.

Additionally, the management teams often fail to make a clear distinction between a data science data validation and a human-oriented data validation. There is an industry-coined phrase called "human-in-the-loop" validation that is used for situations when it isn't enough to run if/else error tests but also necessary for a person to sense-check the data themselves. Examples of this are phone numbers (that require an actual phone call), email addresses (email bounce-back test), and websites (that need a Google search). When a human successfully makes these validation tests, machine learning can be applied to both learn from the validations and also to assess newly recruited "human data validators."

It is, however, important to note that personally identifiable information (PII) is sensitive, personal data that is protected by data regulatory laws such as the General Data Protection Regulation<sup>16</sup> and can make data validation tedious. This is particularly a challenge with the types of data mentioned earlier (that is, email address and phone number) that can be tied back to a person. Since these regulations make attaining, storing, and using PII data a costly liability for a company, it can be risky when PII data are used as part of the data validation approach or as the data being validated itself.

Perhaps the most challenging of all is scalability. Scalability of data validation is how decision makers find the easiest excuse to turn a blind eye on. It is a costly aspect of the data pipeline that needs more research and attention. To go from 95% accuracy to 99% accuracy, in confidence intervals, can cost a company an exponential amount more versus their standard resourcing cost in data validation due to the human aspect of it. When a business decision maker is more comfortable to go with 95% over 99% accuracy, it is often too late to realize that it could cost thousands of dollars to a startup or billions of dollars to a Fortune 500 company in a worst-case scenario. This needs to change. Decision makers need to have fundamental standards for quality and reliable data because they are critical for corporate survival.<sup>2</sup>

## DATA VALIDATION DASHBOARD

We have discussed some of the essential key concepts and challenges regarding data validation. It is evident that when it comes to data validation, human oversight cannot be completely removed from the practice.



Data validation is a part of a project that exists to help a human with a data-driven decision. For example, before digital marketing existed, traditional marketers with little to no data, would make decisions based on gut instinct and “intuition” from their experiences— invaluable human traits that cannot be learned. Even today, no matter how advanced data validation software is, or automated methodologies are, humans are unlikely to proceed with decisions that are not vetted by another human. The final decision would probably be made by someone at a senior level, with experience and contextual knowledge, who can tell whether the data validation process was successful.

A *data validation dashboard* might serve as a means for an experienced data practitioner to monitor data analysis processes from start to finish. The dashboard could also be used as a tool that, over time, improves recommendations for handling data validation by suggesting the appropriate time to be spent at each step, the number and types of people involved, the format of the data validation technique, and the expected results.

Two of the core functions of such a dashboard would be

- a. a checklist to help decide and sequence the most appropriate tasks to apply for a given problem
- b. a reference list of learned data quality issues filtered specific to the data set/problem at hand.

The dashboard could enable teams or project managers to allocate resources, tasks, and also more effectively supervise the status and the best outcome possible for projects.

Since many people are comfortable with knowing that they understand data validation at the fundamental level, it is most important for you to take away the more subtle nuances about it that this column article highlights. Data validation generally occurs but isn't exclusively at the data source level. Data sources or data sets should be acknowledged as being comparatively “more” or “less” accurate, as perfect accuracy is impossible to conclude. Records and data points, although confirmed accurate, can run the risk of inaccuracy over time if not routinely verified.

As marketing automation and data-driven retargeting becomes more widespread, there will be an increasing need to guarantee the accuracy of data pertaining to real people and business entities. A data analytics operation can produce ground-breaking data-driven solutions or recommendations, but when presented to an already risk-averse business decision maker, the last thing you want is to allow the risk of having employed inaccurate data from the very beginning. 🤖

---

*RECORDS AND DATA POINTS, ALTHOUGH CONFIRMED ACCURATE, CAN RUN THE RISK OF INACCURACY OVER TIME IF NOT ROUTINELY VERIFIED.*

---

## REFERENCES

1. C. Farr, “*This patient’s medical record said she’d given birth twice—In fact, she’d never been pregnant,*” CNBC, Dec. 9, 2018. <https://www.cnn.com/2018/12/09/medical-record-errors-common-hard-to-fix.html> (accessed June 21, 2021).
2. “*Poor - Quality data imposes costs and risks on business, says new Forbes insights report,*” Forbes, May 31, 2017. <https://www.forbes.com/sites/forbespr/2017/05/31/poor-quality-data-imposes-costs-and-risks-on-businesses-says-new-forbes-insights-report/?sh=4eb96655452b> (accessed June 21, 2021).
3. A. Marshall and A. Davies, “*Uber’s self-driving car saw the woman it killed, report says,*” WIRED, May 24, 2018. <https://www.wired.com/story/uber-self-driving-crash-arizona-ntsb-report/> (accessed June 21, 2021).
4. D. Reinsel, J. Gantz, and J. Rydning, “*The digitization of the world from edge to core,*” International Data Corporation, Needham, MA, White Paper, 2018. Accessed: June 19, 2021. [Online]. Available: <https://resources.moredirect.com/white-papers/idc-report-the-digitization-of-the-world-from-edge-to-core>
5. J. Gao, C. Xie, and C. Tao, “Big data validation and quality assurance - Issues, challenges, and needs,” in *Proc. IEEE Symp. Service-Oriented Syst. Eng. (SOSE)*, Mar. 2016, pp. 433–441. doi: 10.1109/SOSE.2016.63.
6. J. Anderson and L. Rainie, “*The future of big data - Main findings: Influence of big data in 2020,*” Pew Research Center, 2012. <https://www.pewresearch.org/internet>

- /2012/07/20/main-findings-influence-of-big-data-in-2020/ (accessed July 8, 2021).
7. "Techopedia explains data validation," Techopedia, 2017. <https://www.techopedia.com/definition/10283/data-validation> (accessed June 15, 2021).
  8. N. Polyzotis, M. Zinkevich, S. Roy, E. Breck, and S. Whang, "Data validation for machine learning," in *Proc. Mach. Learn. Syst.*, 2019, vol. 1, pp. 334–347.
  9. T. Jun, C. Kai, F. Yu, and T. Gang, "The research & application of ETL tool in business intelligence project," in *Proc. Int. Forum Inf. Technol. Appl.*, May 2009, vol. 2, pp. 620–623. doi: 10.1109/IFITA.2009.48.
  10. N. Polyzotis, S. Roy, S. E. Whang, and M. Zinkevich, "Data lifecycle challenges in production machine learning: A survey," *ACM SIGMOD Rec.*, vol. 47, no. 2, pp. 17–28, June 2018. doi: 10.1145/3299887.3299891.
  11. "Salary guide 2020," Robert Half Technology, Menlo Park, CA, 2020. [Online]. Available: [https://www.roberthalf.com/sites/default/files/documents\\_not\\_indexed/2020\\_Salary\\_Guide\\_Technology\\_NA.pdf](https://www.roberthalf.com/sites/default/files/documents_not_indexed/2020_Salary_Guide_Technology_NA.pdf) (accessed June 20, 2021).
  12. C. Taylor, "Structured vs. unstructured data," *Datamation*, May 21, 2021. <https://www.datamation.com/big-data/structured-vs-unstructured-data/> (accessed July 8, 2021).
  13. M. Goetz, G. Leganza, E. Miller, and J. Vale, "Data performance management is essential to prove data's ROI," FORRESTER, 2018. <https://www.forrester.com/report/Build%20Trusted%20Data%20With%20Data%20Quality/-/E-RES83344> (accessed June 20, 2021).
  14. M. Fernandez and H. F. Durrant-Whyte, "An information-theoretic approach to data-validation," in *Proc. IEEE Amer. Contr. Conf.*, June 1993, pp. 2351–2355. doi: 10.23919/ACC.1993.4793308.
  15. A. El-Mowafy, "Diagnostic tools using a multi-constellation single-receiver single-satellite data validation method," *J. Navig.*, vol. 68, no. 1, pp. 196–214, 2015. doi: 10.1017/S0373463314000526.
  16. C. Bernstein, "Personally identifiable information (PII)," *TechTarget*, Feb. 2020. <https://searchsecurity.techtarget.com/definition/personally-identifiable-information-PII> (accessed June 27, 2021).

**NORITA AHMAD** is an associate professor of management information systems at American University of Sharjah, Sharjah, United Arab Emirates. Contact her at [nahmad@aus.edu](mailto:nahmad@aus.edu).

**KEVIN DIAS** is a performance and analytics manager at MRM, Toronto, M5V 0N6, Canada. Contact him at [kevinrosedias@gmail.com](mailto:kevinrosedias@gmail.com).



**IEEE COMPUTER SOCIETY**  
**Call for Papers**

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

**GET PUBLISHED**  
[www.computer.org/cfp](http://www.computer.org/cfp)

 **IEEE**



**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](http://www.computer.org)

**OMBUDSMAN:** Direct unresolved complaints to [ombudsman@computer.org](mailto: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](mailto: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 17 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](http://www.computer.org/certification).

## **BOARD OF GOVERNORS MEETING**

TBD

## **EXECUTIVE COMMITTEE**

**President:** William D. Gropp

**President-Elect:** Nita Patel

**Past President:** Forrest Shull

**First VP:** Riccardo Mariani; **Second VP:** David S. Ebert

**Secretary:** Jyotika Athavale; **Treasurer:** Michela Taufer

**VP, Membership & Geographic Activities:** Andre Oboler

**VP, Professional & Educational Activities:** Hironori Washizaki

**VP, Publications:** David S. Ebert

**VP, Standards Activities:** Annette Reilly

**VP, Technical & Conference Activities:** Grace Lewis

**2021–2022 IEEE Division VIII Director:** Christina M. Schober

**2022–2023 IEEE Division V Director:** Cecilia Metra

**2022 IEEE Division VIII Director-Elect:** Leila De Florian

## **BOARD OF GOVERNORS**

**Term Expiring 2022:** Nils Aschenbruck,

Ernesto Cuadros-Vargas, David S. Ebert, Grace Lewis,

Hironori Washizaki, Stefano Zanero

**Term Expiring 2023:** Jyotika Athavale, Terry Benzel,

Takako Hashimoto, Irene Pazos Viana, Annette Reilly,

Deborah Silver

**Term Expiring 2024:** Saurabh Bagchi, Charles (Chuck) Hansen,

Carlos E. Jimenez-Gomez, Daniel S. Katz, Shixia Liu,

Cyril Onwubiko

## **EXECUTIVE STAFF**

**Executive Director:** Melissa A. Russell

**Director, Governance & Associate Executive Director:**

Anne Marie Kelly

**Director, Conference Operations:** Silvia Ceballos

**Director, Information Technology & Services:** Sumit Kacker

**Director, Marketing & Sales:** Michelle Tubb

**Director, Membership & Education:** Eric Berkowitz

**Director, Periodicals & Special Projects:** Robin Baldwin

## **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](mailto:help@computer.org)

**Los Alamitos:** 10662 Los Vaqueros Cir., Los Alamitos, CA 90720;

**Phone:** +1 714 821 8380; **Email:** [help@computer.org](mailto:help@computer.org)

## **MEMBERSHIP & PUBLICATION ORDERS**

Phone: +1 800 678 4333; Fax: +1 714 821 4641;

Email: [help@computer.org](mailto:help@computer.org)

## **IEEE BOARD OF DIRECTORS**

K. J. Ray Liu, *President & CEO*

Saifur Rahman, *President-Elect*

John W. Walz, *Director & Secretary*

Mary Ellen Randall, *Director & Treasurer*

Susan "Kathy" Land, *Past President*

Stephen M. Phillips, *Director & Vice President, Educational Activities*

Lawrence O. Hall, *Director & Vice President, Publication Services and Products*

David A. Koehler, *Director & Vice President, Member and Geographic Activities*

James E. Matthews, *Director & President, Standards Association*

Bruno Meyer, *Director & Vice President, Technical Activities*

Deborah M. Cooper, *Director & President, IEEE-US*

## DEPARTMENT: PREDICTIONS

# The Next Wave in Cloud Systems Architecture

Amin Vahdat, *Google*

Dejan Milošević, *Hewlett Packard Labs*

*Twenty-five years ago, no one expected services to be seamlessly and consistently delivered across campuses and edge. It is now possible to envision machine learning-powered services capable of scaling to the reliability, performance, and security requirements of billions of users worldwide.*

**B**ecause of their visionary nature but also their highly volatile success-failure ratio, predictions have always been popular among a wide spectrum of audiences. The COVID-19 pandemic has turned this popularity into a necessity. Time constraints and a lack of thorough experimentation necessitated a dependence on predictions: which vaccines to use, when to close or open countries, when and how to reopen offices, and so on. Suddenly, predictions became a way of life.

After 10 years of making technology predictions for the IEEE Computer Society (see the press releases at <https://www.computer.org/press-room>), two successful special issues on technology predictions in *Computer* (December 2020 and July 2021), and panels at numerous events (SWITCH 2020; the IEEE Computer Society Computers, Software, and Applications Conference 2020; and so forth), we have decided to initiate the column "Predictions."

We decided to invite a reputable guest for each "Predictions" column. A guest should have a demonstrated vision for the future of computing and a track record of delivering on that vision. Our ideal guest should have a combination of academic rigor, deep technology knowledge, and an understanding of business implications.

Amin Vahdat, an engineering fellow and vice president at Google, is a perfect match for the inaugural

"Predictions" column. He leads systems infrastructure at Google, and prior to that, he was a professor of computer science at the University of California San Diego and Duke University. Vahdat brings a wealth of experience in compute, operating systems, accelerators, storage, and networking. In his professional role, it is essential to predict the workload evolutions, customer demands, and trends in technology.

In this article, Vahdat takes us on a prediction tour of the future of systems infrastructure, addressing topics such as the emergence of new accelerators, impact of the growth of data, disaggregated data center designs, the evolving roles of operating systems and programming languages, application delivery models, the roles of networking, and much more. I hope that you enjoy reading this column as much as I enjoyed working with him to deliver it.

**Dejan Milošević:** Interesting technology and business implications are driven by the imminent end of Moore's law, such as a plethora of innovative accelerators, the rise of photonics, and the introduction of new memory technologies. Which one of these is a fundamental disruptor and which is a temporary transition?

**Amin Vahdat:** The computing industry has delivered incredible new business and societal capabilities through sustained, exponential improvements in the scale and performance efficiency of the underlying hardware and software. However, the underlying hardware trends that have powered this march over the past few decades are slowing or stopping with each



generation of hardware, including general-purpose CPUs, dynamic RAM (DRAM), storage devices, and network elements, delivering incrementally less benefit relative to the previous generation, alarmingly with lower levels of reliability in many cases. This will mean that software optimizations, including novel algorithms, protocols, and replication strategies, will play an important role in maintaining the cadence of regular system-efficiency improvements.

I believe that all the technologies you point to, and more that are yet to come, are fundamental disruptors we will need to accommodate. For the first time in decades, a range of new technologies is actually a must have—rather than a nice to have—to sustain exponential growth in compute capability. Current capabilities, like real-time language recognition and translation, would simply not have been possible from either a cost or power standpoint without hardware acceleration through accelerators like tensor processing units and GPUs. Similarly, exploding demand for video consumption on a range of devices and networks could not be effectively supported without hardware-supported video coding. These examples are just the beginning of an increasing array of segment-specific hardware architectures that can deliver step-function improvements to domains such as analytics, high-performance computing, and real-time serving.

On the memory and storage front, we are hitting a wall in available improvements in cost per gigabyte, especially as we try to simultaneously meet the needs of different applications that are individually latency-, bandwidth-, and capacity-bound on the same hardware. We now have underlying hardware technologies that can strike tradeoffs between cost, latency, bandwidth, and capacity somewhat independently, meaning that we will deploy servers with a range of memory and storage configurations, enabling higher-level scheduling infrastructure to map applications to available hardware configurations in the most cost- and performance-effective manner.

Photonics have always played a critical role in the scale of available computation in the data center. Today, silicon photonics, or integration of photonics onto integrated circuits, is emerging as a requirement for continued expansion of power-efficient data rates. We are already seeing this play out in traditional transceiver design, with the next question being the role silicon photonics will play in mainboard and ASIC design, with reliability for hundreds of optical engines integrated on a single board or ASIC being one of the most pressing challenges.

**Milojicic:** Data center design has been motivated by cost and performance for quite some time. Given the rate of evolution of individual technologies, is there an opportunity for rethinking interfaces to optimize performance? For example, accelerating access to data and/or interconnections by deploying computations closer to it. The so-called in- or near-memory computations and similar for interconnects with smart network interface controllers (NICs), data processing units, and so on?

**Vahdat:** Yes, absolutely. The rate of growth of data and the increasing computation required to summarize and transform data into suitable formats is growing well beyond what our CPU budgets will allow for power- and cost-efficient processing. Today, we strive to treat data as a large, location-independent “blob” with intermediate processing steps composed together in a manner not far from classic UNIX-pipe composition. While convenient, this means that data must often move back and forth across kilometers of compute within a large campus, compressed, encrypted, decompressed, decrypted, transformed, processed, and replicated dozens of times. Processor architects optimize pJ/bit, and it is likely that we will need similar measures in considering the cost of end-to-end processing of every ingested byte of data.

The resulting metrics will lead to a number of necessary innovations in system design. First, some amount of computation power will be needed closer to data stores with closely coupled data ideally placed closer to one another in a logical topology. In the extreme, new hardware will colocate server-class processing power or even hardware accelerators in the same server or rack enclosure as the data. However, more generally, such sophisticated processing will require deep understanding of the logical data processing pipeline and an orchestration and runtime layer able to place and migrate data computation according to overall data lifetimes. Understanding the provenance of derived data will open opportunities to create efficient, intermediate formats, perhaps allowing different processing elements to handle a range of native formats and summaries, ideally with a “data compiler” capable of placing the right preprocessing at the right point for a specified data flow. Smart NICs and their generalization will form the nervous system in the above runtime, likely managing both the movement of the data from location to location but also its intermediate processing.

While there is substantial complexity required to realize this, the next wave of efficiency improvements will come from end-to-end performance considerations rather than lower-level measures such as million instructions per second and storage capacity/input-output per second. How much processing, energy, and storage is required to deliver an insight or result in the end-to-end composition? There are integer factors of performance, cost, and energy efficiency available in our infrastructure when viewed through this lens.

**Milojicic:** Do we need new operating systems, programming languages, programming environments, and middleware to make all this seamlessly work, or would those from the previous era suffice?

**Vahdat:** Software will drive the success of this next wave in infrastructure evolution. What we will see is that the shape of compute containers can change dynamically at runtime, perhaps adding and removing memory, storage, or accelerators based on the needs of the services currently scheduled on the server.

Real-time performance monitoring will support isolation among tenants while accounting for antagonists at multiple levels of the system hierarchy; for example, L3 or DRAM capacity versus bandwidth requirements, and behaviors of individual applications. The level of malleability in the composition of servers will need to fundamentally change from the operating systems’ current view of a fixed “hardware” environment from boot to shutdown.

At the distributed systems level, services will require much more predictability and determinism in accessing remote storage and compute resources. This will, in turn, require a runtime capable of delivering isolation among millions of concurrent communication channels, all while ensuring efficiency and the mapping of requests to the underlying replica or resource best capable of fulfilling them.

**Milojicic:** How can we most effectively accomplish hardware–software co-design to account for optimizations adequately at different levels of the stack?

**Vahdat:** The level of visibility we have into the dynamic nature of real data center computation is incredibly limited. We understand what industry benchmarks like SPEC and TPC-H look like in isolation and know how to optimize them for hardware. However, actual data center workloads are increasingly heterogeneous, multi-tenant, and distributed, substantially reducing the predictive power of existing benchmarks.

Moving forward, we will need to start with much deeper hardware measurement infrastructure to characterize workloads in the wild. This can, in turn, support new benchmarks that account for wider variability in computational structure, potential hardware offload capability, and the fundamental distributed and multi-tenant nature of modern computation. Single-server, single-tenant benchmarks cannot be the basis for projecting the value of future infrastructure. An open question is the required scale, heterogeneity, and variability for sufficient predictive power.

**Milojicic:** To optimize the data center cost, disaggregated design has been taking off, with accelerators being separated from compute—deployed in racks-size units—disaggregated memory is emerging and storage has been deployed (for example,

storage area network and network attached storage) for quite some time. Is this trend going to continue in the future?

**Vahdat:** The interesting thing about disaggregation is that, while it holds fundamental technical appeal, including for me personally, its progression has been rather limited since we deployed the first disaggregated storage solutions. Systems like GFS (the Google File System) showed how to treat thousands of hard drives spread across the data center as the underlying storage for a distributed and disaggregated storage system. However, the bandwidth and latency of hard drives are modest compared to the speeds of data center networks and, as importantly, the software layers responsible for managing I/O requests.

True disaggregation of SSDs has been slower to progress in part because local, dedicated devices have high bandwidth and low latency relative to data center network speeds but also because there can be little software on the path between a client request and device access while maintaining the illusion of “local” access to a disaggregated device.

While accelerators are deployed in rack-scale units, they are also typically deployed with dedicated servers and often dedicated secondary networks. There is very little virtualization support for accelerators today.

I do believe that the next level of end-to-end system efficiency and flexibility will require support for disaggregation. However, it will also require some breakthroughs in hardware/software organization because the assumptions built up over decades about local device access over a dedicated PCIe link will be hard to break.

**Milojicic:** Can you tell us how networking is evolving in the data center and how it is affected by the trends and needs of data and compute?

**Vahdat:** The network today constitutes the smallest portion of our data center spend, but it offers some of the biggest challenges and opportunities. For example, the network is often the largest source of large-scale failures and extended downtime. Because it fundamentally connects computation and storage together at scale, a failure in the network most easily

cascades resulting in large-scale outages. Similarly, managing the lifecycle of the network from turnup to upgrades to turndown is often the most complex and toilsome for the operations team. Tying the two together, many network outages are correlated with network operations.

Given the increasing societal reliance on compute infrastructure, and the thousands of seemingly independent cloud services, we multiplex onto shared underlying network infrastructure, the reliability of the network must fundamentally improve. Since, in the end, individual network components and systems are unavoidable at cloud’s scale of deployment and its rapid rate of evolution, solutions likely lie in designs that ensure hard levels of network isolation and multiplexing of services onto multiple independently operated network infrastructure.

Network performance and performance predictability will also be critical, as we discussed earlier. The disaggregated and data-centric data center will require not just higher performance and lower latency but predictability and isolation under a range of highly variable and bursty communication patterns. This will require us to continue the evolution of software-defined networking to enable visibility all the way to the end applications and their composite, multinode communication patterns, with microsecond-granularity actuation loops allocating bandwidth and rate limits to meet the real-time application SLOs, focusing on remote procedure calls and coflows rather than lower-level metrics focused on packets.

**Milojicic:** Data and compute delivery models have evolved from on premise to public cloud, to hybrid cloud to edge. What is the next step in delivery models?

**Vahdat:** One important note is that, while there is a lot of enthusiasm around emerging compute delivery models from public to hybrid to edge cloud, we are still in the very early stages of the migration to these emerging hosting approaches. The vast majority of computing and storage still runs in enterprise data centers with many challenges that must be overcome before we can get to baseline modernization of digital infrastructure.

This is one reason why hybrid connectivity and the ability to seamlessly integrate on-premises

infrastructure with cloud hosted and managed infrastructure is so critical to enabling more rapid migration. Here, there are many practical and research challenges to providing a “one-network” administrative view across multiple sites and cloud providers with individual services ideally, transparently, and incrementally migrating from one side to another.

As such, the next step in the delivery model must take on the software tooling, monitoring, and management necessary to make hybrid, multicloud, and edge-deployment scenarios seamless from the perspective of individual businesses. Similarly, application developers will need the runtime support and consistent APIs to enable them to write once and run anywhere.

**Milojicic:** What is the role of virtualization, and how is it evolving? From bare metal to virtual machines to containers and most recently to functions as a service (serverless), virtualization has been increasingly getting more hardware support. What is the next step in evolution?

**Vahdat:** Yes, this is a great observation. We are seeing the continuing evolution of virtualization to the point where individual bare-metal servers can be securely configured for individual customers, potentially allocated and reallocated at fine time scales. The hardware architecture, from the BMC, network connectivity, root of trust, and more to allow this in a multitenant data center has been understudied in academia and underdeveloped industry wide.

On the other extreme, we will need to move beyond “static slice of hardware” virtualization capability to more flexible and dynamic container shapes in support of higher-level application structure that benefit from dynamic scale out under failures and bursty access conditions. These containers will need to deliver requisite levels of isolation and security among colocated tenants, all while managing tail latency. As discussed earlier, this will mean that the compute, accelerator, memory, storage, and network resources available to a container will come and go with both the application and the container OS, evolving to understand and manage this new dynamism in line with the trends toward disaggregation of the underlying hardware components across the data center.

**Milojicic:** Will 5G and follow-on technologies affect data center designs? For example, will 5G make edge much more real time, putting more pressure on data centers? How much will edge impact data center design and evolution of computing technologies?

**Vahdat:** 5G promises to bring affordable, lower latency, and predictable bandwidth connectivity to a whole new range of computing devices. Data rates will continue to explode, accelerating from the already-impressive place we are at today. One of the key questions will then be the tradeoff between: 1) very low-latency reaction times to data generation/ the sensing to actuation loop, 2) data provenance and sovereignty requirements for the data, 3) the costs of wide area network transport, and 4) the inherent efficiencies associated with large-scale, centralized computing campuses built from the ground up for power and cooling efficiency relative to more expensive, smaller-scale, and less-optimized colocation facilities located in the heart of metropolitan areas. For example, can we afford to pay the tens of milliseconds in speed of light latency required to transport newly generated data to the nearest large-scale campus? Does the data need to be stored and processed within the jurisdiction of a particular company or country?

The reality is that the tradeoff will be application and scenario specific, which makes the emergence of 5G particularly interesting from the perspective of dynamically configuring machine learning inference, general-purpose computing, storage, and communication pipelines across local, metropolitan, and wide area network infrastructures.

**Milojicic:** Historically, there was a dichotomy between large-scale distributed systems and large parallel systems. It appears that the former have won, except for the needs of high-performance computing (HPC) and high-end analytics systems. Going forward, for economic reasons, there appears to be a convergence of general-purpose compute, artificial intelligence, HPC, and data analytics systems. Will a unified system design be possible to meet the requirements of traditional HPC and analytics applications?

**Vahdat:** Yes, this is another really nice observation. Historically, parallel applications have required



predictable, isolated, homogeneous, single-tenant execution environments to achieve reasonable performance given the regular nature of the underlying computation and communication loops. Distributed data center applications have, on the other hand, been loosely coupled, fault tolerant, multitenant, and capable of running on multiple generations of hardware and low-level software simultaneously.

While both evolutionary paths have been pragmatic, we are seeing a convergence between the two, also for pragmatic reasons. High-performance computing and machine learning are increasingly migrating to cloud environments where the availability of large-scale batch resources and efficient virtualization to achieve uniformity is making it attractive to run in more heterogeneous and unpredictable environments, with software and hardware providing the illusion of a uniform, dedicated execution environment. At the same time, the availability of efficient, microsecond-scale communication stacks with hardware support for isolated multitenancy is exposing a number of opportunities for traditional distributed applications to achieve much higher levels of performance and efficiency through reduction of expensive caching and data denormalization.

**Milojicic:** Any closing thoughts that you would like to leave our readers with?

**Vahdat:** It's a really exciting time to be working in computing infrastructure. It was about 20–25 years that a combination of academic research and industrial breakthroughs laid the foundation of modern Internet service delivery. Back then, no one even dared dream that services would be seamlessly and consistently delivered across campuses and edge consisting of tens of thousands of commodity servers running open source operating systems all interconnected by a dedicated wide area and data center networks comparable in scale to the public Internet, with software-managed computing and storage providing the illusion of a single system image across these same exascale computing platforms.

And yet, this seemingly impossible vision is now considered standard practice. It is now correspondingly simple, or at least possible, to dream of new planetary-scale, interactive services that can

seamlessly scale to the reliability, performance, and security requirements of billions of users across the globe for enterprises, large and small.

At the same time, the underlying forces and design patterns that have powered this last revolution in computing are slowing or stopping. Without the winds of exponential growth in compute and storage capacity at fixed cost at our back with the simultaneous headwind of increasing exponential growth of data rates and processing needs, we as a community have to develop fundamentally new design points focused on end-to-end distributed systems efficiency metrics, powered by new segment-specific hardware and increasingly sophisticated dynamic runtimes.

**M**y personal excitement is fueled by the belief that this shift in thinking will enable a whole new set of capabilities that perhaps no one dares dream today. One higher-level prediction I feel comfortable with is that the new capabilities will afford a shift from planetary-scale interactive services to the real-time generation of proactive insights driven by secure, privacy-preserving data analytics frameworks capable of shifting through an ever-increasing explosion of available data, affording fundamental benefits from health care to the sciences to manufacturing and more. 🌐

**AMIN VAHDAT** is an engineering fellow and vice president for systems infrastructure at Google, Mountain View, California, 94043, USA. Contact him at [vahdat@google.com](mailto:vahdat@google.com).

**DEJAN MILOJICIC** is a distinguished technologist at Hewlett Packard Labs, Palo Alto, 94306, California, USA. Contact him at [dejan.milojicic@hpe.com](mailto:dejan.milojicic@hpe.com).



# Interhost Orchestration Platform Architecture for Ultrascale Cloud Applications

Sasko Ristov  and Thomas Fahringer , *University of Innsbruck, 6020 Innsbruck, Austria*

Radu Prodan , *University of Klagenfurt, 9020 Klagenfurt, Austria*

Magdalena Kostoska  and Marjan Gusev , *Ss. Cyril and Methodius University, Skopje 1000, Macedonia*

Schahram Dustdar , *TU Wien, 1040 Vienna, Austria*

*Cloud data centers exploit many memory page management techniques that reduce the total memory utilization and access time. Mainly these techniques are applied to a hypervisor in a single host (intra-hypervisor) without the possibility to exploit the knowledge obtained by a group of hosts (clusters). We introduce a novel interhypervisor orchestration platform to provide intelligent memory page management for horizontal scaling. It will use the performance behavior of faster virtual machines to activate prefetching mechanisms that reduce the number of page faults. The overall platform consists of five modules—profiler, collector, classifier, predictor, and prefetcher. We developed and deployed a prototype of the platform, which comprises the first three modules. The evaluation shows that data collection is feasible in real-time, which means that if our approach is used on top of the existing memory page management techniques, it can significantly lower the miss rate that initiates page faults.*


**D**ynamic memory page management techniques, such as memory deduplication, page faults management, memory overcommitment, memory ballooning, or hot-swapping, rely on the cooperation of the virtual machines (VMs) hosted on a single physical host.<sup>1</sup> Although all these techniques provide autonomous and automatic memory page management that reduces the total memory utilization and memory access time, their application domain is still limited within a single host. On the other hand, cloud environments use horizontal scaling such that hundreds or thousands of VMs (*VM-siblings*) of the same image work on the same problem. These VM-siblings use the same guest operating system, the same code segment, and many memory pages of the same data segment are identical or similar, thereby conducting similar memory access patterns.

Since VM-siblings may be scheduled on multiple hosts, the state-of-the-art intrahost memory management methods cannot be fully exploited. The goal of this article is to introduce an interhypervisor orchestration platform, which uses the knowledge obtained by VM-siblings that are hosted on different hosts and potentially use it in real-time to reduce the memory page fault ratio, for negligible resource utilization overhead and latency.

## RELATED WORK

Wei-Zhe Zhang<sup>2</sup> presented an automatic memory control of multiple VMs that dynamically adjusts allocation according to the used memory by the VMs, while Qi Zhang<sup>3</sup> used a shared memory pool for fast just-in-time memory page recovery. Although these techniques reduce the number of memory page faults, they apply on a single host without being aware of other VM-siblings at other hosts.

Hines<sup>4</sup> and Tasoulas<sup>5</sup> use the estimation of application memory requirements for memory balancing and distribution. Their techniques automate the



**FIGURE 1.** Definition of VM-cluster and VM-siblings (VM-leader and slack-VMs).

distribution of the memory across VMs, but again limit the hypervisor's level on a single host. Additionally, they quantify the amount of memory without a qualitative estimation of the specific page accesses in the near future for each VM-sibling.

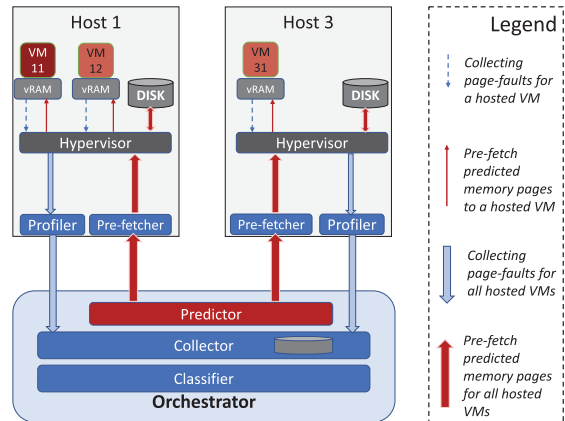
Several researchers presented an orchestration for multiple hypervisors. Gopalan *et al.*<sup>6</sup> introduced the span virtualization, which allows multiple hypervisors to control the memory of guest's OS concurrently. Still, this orchestration is on a single host. Fecade *et al.*<sup>7</sup> orchestrated multiple hypervisors in mobile cloud computing. They used a Bayes-based classifier to predict failures in hypervisor and to prevent VM failures by migrating them to another host. However, the authors verified their approach with simulation, without real implementation and without considering the network latency.

Prefetching is a commonly used technique in memory management. Ren *et al.*<sup>8</sup> introduced an asynchronous prefetching mechanism to speculatively prefetch the dirty pages from a primary VM on a secondary VM on another host without disrupting its execution. While this algorithm shortens the sequential dependence when VM checkpoints are generated and transmitted to a VM on another host, still, it uses one-to-one mapping between the primary and secondary VMs without considering memory access patterns from other VM-siblings.

## INTERHYPERVISOR ORCHESTRATION PLATFORM

### Terminology

Before diving into details, we explain the used terminology. Let  $VM_{11}$ ,  $VM_{12}$ , and  $VM_{31}$  denote three VMs hosted on two hosts (Host1 and Host3), as presented in Figure 1, such that the first index identifies the host, while the second one determines the VM on that host. For example,  $VM_{12}$  represents the second VM hosted on Host1. Let all VMs are a part of horizontal scaling, which means they run the same application, either for different input data or serve requests generated by different users.



**FIGURE 2.** Design of the interhypervisor orchestration platform.

**Definition 1. (VM-cluster).** A VM-cluster is a set of VMs that are scaled horizontally and usually deployed across multiple hosts.

For example, Figure 1 illustrates a part of a data center, where VMs are grouped in: VM-cluster1 with three VM-siblings, while VM-cluster2 and VM-cluster3 with four. Each VM-sibling runs the same application, which is scaled horizontally across three hosts.

**Definition 2. (Specific VMs within a VM-cluster).**

A VM-leader is the fastest VM-sibling that leads the execution within one VM-cluster. Slack-VMs are all other VM-siblings that perform slower than the VM-leader within a specific VM-cluster.

Each VM-cluster identifies a single VM-leader at each point of time, while all the other VM-siblings are slack-VMs. However, these roles may change in time for a specific VM.

Without losing generality, we assume a heterogeneous environment where the VMs and hosts perform at different speeds. Figure 1 presents that  $VM_{11}$ ,  $VM_{32}$ , and  $VM_{14}$  are VM leaders of the corresponding VM-clusters 1, 2, and 3. For example, besides the VM leader  $VM_{11}$ , the VM-cluster1 contains also other VM-siblings (slack-VMs including  $VM_{12}$  and  $VM_{31}$ ).

### Platform Architecture

The interhypervisor platform orchestrates hypervisors of a group of horizontally scaled VM-clusters, as presented in Figure 2. To orchestrate VM-siblings in the horizontal scaling, the first step of the orchestrator is to classify all VMs in VM-clusters according to the information from the cloud controller. Further on, the orchestrator communicates with the agents on each

host to gather the necessary access data that caused page faults for each hosted VM-sibling. Finally, the orchestrator detects patterns in the page faults to determine the VM-leader of each VM-cluster, whose behavior will be used later for prediction and prefetching the memory pages of the *slack-VMs*.

The proposed interhypervisor orchestration platform consists of two main processes: *i) a bottom-up data collection and management, and ii) a top-down control process*. This paradigm is particularly suitable for concurrent management of these two processes. For example, hypervisors can manage VMs on the same host, while our orchestrator supports global views of VM-clusters, combining and analyzing data of all VM-siblings of a single VM-cluster scattered on different hosts. On the other side, the top-down control process allows flooding the information about prefetching memory pages that already generated page faults at the VM-leader, thereby reducing the page faults ratio for the slack-VMs and significantly improving their performance.

The orchestration platform consists of five modules: *classifier, collector, and predictor* within the central part of the orchestrator, together with *profiler and prefetcher* distributed on each host. Since the orchestrator collects data about page faults from profilers and sends the predicted memory pages for all hosted VMs back to prefetchers hosted on multiple hosts, the modules of the orchestrator should be deployed on servers with a higher bandwidth and a small network diameter to the orchestrated hosts. Our approach with the centralized orchestrator and distributed profilers does not require a direct communication between VMs' operating systems.

The orchestrator will have access to memory usage info for all orchestrated VMs, which may open security and privacy risks. In order to mitigate such risks, the profiler sends anonymized data that do not contain information about the owners of the VMs, their addresses, or credentials.

The orchestrator is only a logical representation and any of its three modules may be hosted on a separate server. Moreover, the modules may be containerized and managed with Kubernetes for higher scalability. The orchestrator implementation is not affected by the existing cloud infrastructure as it works on a lower (hypervisor) level.

### **Classifier**

The classifier groups all VMs in VM-clusters, such that each VM member of horizontal scaling becomes a VM-sibling within a specific VM-cluster. Clustering is a dynamic process in the cloud ecosystem regularly

performed by the classifier, where VMs are instantiated, replicated, migrated to another host and terminated. The classifier can be extended to cluster VMs that are not a part of horizontal scaling or VM-clusters without VM-leader, if they have a similar memory access pattern.<sup>16</sup>

### **Profiler**


Each host deploys a profiler that communicates with the local hypervisor to collect data about page faults and swapping for each hosted VM, and sends it to the collector within the centralized orchestrator. The profiler can be built based on iBalloon,<sup>9</sup> which provides efficient intrahypervisor VM memory balancing within a consolidated host. iBalloon runs a lightweight monitoring process (daemon) in each VM of the host that gathers information about its memory utilization. This technique improves the overall performance for memory-intensive applications with less than 5% CPU overhead tradeoff, compensated due to the CPU under utilization compared to the main memory.

### **Collector**

The collector is a simple module that gathers the data from all profilers in a single and persistent centralized *knowledge base*, which contains memory access data of all VM-siblings within each VM-cluster. The main challenge of the collector is to determine the size and length of historical data considered by the predictor. Accordingly, the collector splits the received data into hot and cold parts. The predictor uses the hot part to determine the memory pages to be prefetched at VM-siblings, while the cold part helps strategic planning of future data center design and maintenance through offline analysis. Extending the iBalloon, one can run a daemon that collects all data from each host's memory profiler. Although the concept of cold and hot memory pages<sup>10</sup> provides performance overhead, we can still use this concept to reduce the memory access interception.

### **Predictor**

The predictor is the brain of the orchestrator that exploits the collected data of the knowledge-base. It possibly uses machine learning-based techniques (beyond the scope for this article) to predict the memory page accesses for each VM-sibling within a VM-cluster. Since the cloud environment is a heterogeneous one, both the VM type and the underlying host computation resources (CPU, RAM) must be considered by the predictor to further improve the prediction accuracy. The centralized predictor can exploit data for memory access patterns of all VM-siblings within a



**FIGURE 3.** Platform prototype implementation

VM-cluster, regardless whether they are hosted on a single host or scattered across several hosts. With this knowledge, the predictor can estimate more accurately which memory pages will be accessed by slack-VMs and inform the prefetchers and hypervisors accordingly to prefetch those memory pages into guest main memory.

### Prefetcher

The predictor submits the information to the prefetcher on each host to initiate prefetching of memory pages from local disks to all locally hosted VMs. The prefetcher issues a command to the Hypervisor to prefetch (swap-in) a memory page into each VM to reduce the memory page fault.

## PLATFORM PROTOTYPE IMPLEMENTATION AND EVALUATION

We implemented three modules (the classifier, profiler, and prefetcher) to investigate the effectiveness of the orchestration platform for data collection. The goal of the initial platform prototype was to evaluate how many memory page faults can be generated, profiled, and transferred to the collector, both for traditional and virtual environment.

### Platform Prototype Implementation

Figure 3 presents the implementation of our platform prototype. We deploy the profiler on two hosts, each installed with XenServer (v7.6.0) and Ubuntu 16.04.6 amd64. The hosts (quad-core CPU, 4 GB RAM, SSD) are connected using NAT and 1 Gb/s network. Each host administers a VM-pool Host 0x/dom0x and each VM gets a VMID, which is used to specify it as a source of transferred data in the profiler. The collector and the classifier are deployed on the orchestrator, which has 6 GiB RAM and is also installed with the same Ubuntu.

We used SysBench benchmark tool to generate intensive memory allocation and page faults in Ubuntu with a single thread, which generated a total of 100GiB memory blocked in blocks of 1MB each:

```
sysbench --num-threads=1 --test=memory
--memory-block-size=1M --memory-total-size=100G run
```

Generated page faults were profiled by systemtap, which was extended to submit the VMID. The following listing presents an entry from the page fault, and shows when (including microseconds) and on which host a process generated a page fault, which could be either write (w) or read (r). Additionally, we submit the type of the page fault, i.e., minor or major. The size of each entry was always the same (51B).

```
ID:Timestamp:PID:fault_address:fault_access:kind:
fault_time
1:1591116972164574:30134:140013325101104d:w:
minor:1
```

For sending the profiled data, we used Apache Kafka (v2.2.0) and Zookeeper configured with default ports and JDK (1.8.0) on all brokers in the system. Systemtap sends data to the Kafka producer at Host 0x/dom0x, which is collected by the Kafka consumer at the collector. Finally, the collected stream data was stored in a file and the collector (written in C) writes it in MySQL. The classifier specifies VMIDs and groups them in a VM-cluster. Our current classifier prototype uses only one VM-cluster as we measure the data transfer rate.

### Evaluation

We conducted two groups of experiments. The first group of experiments was intended to investigate the cap varying the number of sources (VMs and hosts), i.e., how many entries the platform prototype is able to generate and send them to the collector without using the profiler systemtap. The second group of experiments determined the overhead of introducing the profiler, which resulted in lower number of records that were collected in real-time. We run each experiment in two different environments (bare metal and virtual). We denote experiments as B1, B1P, B2, B2P, V1, V2, and V1P, where B and V denote the environment (bare metal or virtual), 1 or 2 sources (VMs or hosts), and P denotes experiments with the profiler.

Table 1 presents the achieved throughput of each experiment. We observe that sending data without the profiler achieved 17 million entries/s for one node as a source, which is the maximum from all experiments since we used only one node. Introducing another node (B2) reduced the throughput to 12 million entries/s. On the other side, the profiler (B1P) reduces the throughput to 250000 entries/s with one node and is stable even with two nodes because the streaming cap of 17 million entries/s is not reached.

Virtual environment reduced the throughput by half for the experiment with one VM compared to the

**TABLE 1.** Results of the evaluation. Presented number for the throughput shows the million of entries received by the collector per second.

Experiment	Average throughput ( $\cdot 10^6/s$ )
B1	17
B1P	0.25
B2	12
B2P	0.25
V1	8.5
V1P	0.078
V2	0.25

equivalent B1 and 3.2 times with the profiler. An interesting observation is the higher deviance in the virtual environment.

## Discussion

Since the memory access time at the host is around 50 ns,<sup>11</sup> we estimate around 20 million memory accesses/s if all accesses are page hits, without any page misses and swaps. However, we are mostly interested in TLB misses and page faults, which happen in a range between 0.1 – 1%.<sup>11</sup> This leads to a maximum number of 200000 records/s that need to be stored and processed. We assume the worst case where all TLB misses are also page table misses (page faults), which is opposite to the write count disparity feature. On the other side, the total number of memory pages is usually smaller than 200000, as the legacy page size of 4 kB is nowadays abandoned to reduce the TLB misses. More precisely, all modern CPU architectures and operating systems support memory page size of 2 MB, while some even in the range of GiB. In a virtual environment, extended page table (EPT) faults are handled within 2.4  $\mu$ s.

We selected the current state-of-the-art streaming platforms to evaluate the feasibility of the new proposed platform, such as Apache Kafka, which can handle at each profiler up to 800000 of 100B-long messages per second, regardless of the data size (even up to 1.4 TB). Our platform achieved the maximal throughput of 17 million entries/s, 51B each.

Although the profiler can collect and update the page faults at each host, Zhang *et al.*<sup>12</sup> specify a hybrid hardware and software tracing mechanism to collect and profile last-level TLB misses, up to cache line granularity of 64B. Moreover, another challenge is to collect data from all profilers to the central orchestrator. For example, for a network overhead of only 1% in  $1s^{-1}$ , we can submit 1300 records, 100B each. Although profilers can group several messages into a few larger

ones to reduce the packet header overhead, the total bandwidth remains in a similar range.

Let us analyze the price to be paid to achieve increased performance. At each host, the platform runs both Kafka to utilize a portion of computing resources. While the CPU is usually underutilized in data centers, the memory is often a bottleneck. Additionally, there is a small network overhead depending on the number of pages transmitted. This bottleneck is visible when more nodes or VMs are used (see Table 1).

## FURTHER IMPLEMENTATION CHALLENGES

### Network Overheads

The designers and programmers must consider the network latency and bandwidth that also impact the quality of the gathered data. Another challenge is the network heterogeneity, especially its latency, as hosts can be connected through a single physical switch, while others through several with higher latency. Recent high-speed and high-throughput memory and network, such as byte-addressable non-volatile memory express (NVMe) over fabrics (NVMe) reported negligible application performance degradation.<sup>13</sup> For example, the latest networking generates very low latency of only 1  $\mu$ s and a very high bandwidth up to  $200s^{-1}$ .<sup>14</sup> These trends in the networking allow possibilities for broader dynamic memory page management through the network and make our orchestration platform feasible.

### Prediction Implementation Challenges

The amount of data analyzed by the prediction process impacts its performance. For example, a large history of records has less impact over the current memory accesses due to many context switches that may happen in the meantime. On the other side, considering a small amount of historical records may not be enough as a slack-VM may perform much worse than the VM-leader and, thus, a memory access pattern cannot be detected or valid for this particular case.

The predictor must propose the pages to swap to the disk and avoid being prefetched to the other VM-siblings. Various application types can also show different behavior.

The behavior of each VM fluctuates due to cloud performance instability<sup>15</sup> and therefore, there is no simple and appropriate function for modelling variations in memory page accesses of a VM. Nevertheless, we exploit the fact that caches and memory paging are not directly mapped but associative, which means

that even a relaxed prediction performance still diminishes the memory page fault rate.

The prediction accuracy is affected by other VMs of other VM-clusters running other jobs on the same host. For example, VM<sub>13</sub>, VM<sub>14</sub>, and VM<sub>15</sub>, which share the same Host 1 memory with VM<sub>11</sub> and VM<sub>12</sub>, will affect their memory access and page faults (see Fig 1). This may make the prediction of VM<sub>31</sub> page access less accurate based on VM<sub>11</sub>'s access pattern. This problem is analyzed by Nemati *et al.*<sup>16</sup> They introduced inter- and intracluster similarity metrics to discover distinct groups of workloads with negligible CPU and memory overhead.

The interhypervisor orchestrator needs to predict a vector of memory page accesses for each VM-sibling, as follows.

- › Characterize the VM-cluster using a set of parameters that reflect the memory pages accesses for each VM-sibling.<sup>16</sup>
- › Estimate the memory access of each VM-leader and use it for VM-siblings. Additionally, considering the heuristics with the write count disparity, only a few and frequently updated memory pages could reduce page faults even more.<sup>17</sup>
- › Experimentally evaluate the various machine learning methods (including random forest or similar) considering the tradeoff among latency and resource utilization overhead versus performance.

## Prefetching Challenges

Modern operating systems and hypervisors support memory pages of various size, such as small (4 KiB), medium (2 MiB), and even large of up to 1GiB. While such plethora of heterogeneous memory page sizes may improve the performance,<sup>18</sup> it may convolute both prediction and prefetching. For example, memory access pattern of a VM-leader that uses memory pages of medium size differs from the access pattern of slack-VMs that use small sized ones.

On the other side, write memory accesses on VM-leaders may be logged by the hardware using the Page-Modification Logging<sup>19</sup> on commodity Intel processors. This hardware-assisted enhancement allows the hypervisor to monitor the modified (dirty) memory pages directly while running VMs, thereby distinguishing the "write-hot" and "write-cold" memory pages in real-time.

## CONCLUSION AND FUTURE WORK

We have introduced a platform that enhances the memory page management techniques to reduce the

page faults and increase the performance of virtualized data centers based on an interhypervisor (interhost) approach. State-of-the-art techniques implemented in today's virtualized environments include host-based prefetching and memory swapping concepts. Currently, these approaches are implemented for a single host and cannot be exploited for VMs spread over different hosts.

The interhost orchestration platform has a potential to open up new research directions in cloud data center memory management. Our approach enhances the memory page management implemented for an intrahypervisor solution by an interhypervisor platform, as a more efficient dynamic technique intended for cloud data centers. It can be efficiently implemented for VMs running an application that implements large horizontal scaling among different hosts.

The basic principle of the new proposed approach for the interhypervisor orchestration platform is to detect memory access patterns that generate page faults within VMs hosted on different hosts. Exploiting the patterns in gathered data, the orchestrator may predict which memory pages will be accessed in the near future and, therefore, may avoid generating page faults at the other VMs (slack-VMs).

The platform prototype was developed including the three modules classifier, profiler, and collector. The initial evaluation showed the effectiveness of the platform to collect generated entries about page-faults with a maximal throughput of 17 million entries/s. Although the initial prototype of the profiler reduced the throughput to 250000 entries/s of a host, still the platform prototype was able to reach the estimated 200000 page faults/s,<sup>11</sup> even with a low power hosts and 1 s<sup>-1</sup> network.

We are currently working in the PRE-FETCH project on implementation of the other two modules the predictor and prefetcher and will investigate the effectiveness of the overall interhost orchestration platform. The initial experiments with the random forest predictor showed a promising accuracy. 🤖

## ACKNOWLEDGMENTS

This work was supported in part by bilateral AUT-MKD project PRE-FETCH: Pro-active Memory Management with Page Faults Prediction in Clouds, financed by Austrian Federal ministry of Science, Research and Economy (BMWFW), and the Ministry of Education and Science of Republic of North Macedonia, and in part by the ASPIDE project funded European Union's Horizon 2020 research and innovation program under Grant Agreement 801091.

## REFERENCES

1. H. Liu *et al.*, "Hotplug or ballooning: A comparative study on dynamic memory management techniques for virtual machines," *IEEE Trans. Parallel Distrib. Syst.*, vol. 26, no. 5, pp. 1350–1363, May 2015.
2. W. Z. Zhang, H. C. Xie, and C. H. Hsu, "Automatic memory control of multiple virtual machines on a consolidated server," *IEEE Trans. Cloud Comput.*, vol. 5, no. 1, pp. 2–14, Jan.–Mar. 2017.
3. Q. Zhang *et al.*, "MemFlex: A shared memory swapper for high performance VM execution," *IEEE Trans. Comput.*, vol. 66, no. 9, pp. 1645–1652, Sep. 2017.
4. M. Hines *et al.*, "Applications know best: Performance-driven memory overcommit with Ginkgo," in *Proc. IEEE CLOUDCOM*, 2011, pp. 130–137.
5. E. Tasoulas, H. Haugerund, and K. Begnum, "Baylocator: A proactive system to predict server utilization and dynamically allocate memory resources using Bayesian networks and ballooning," in *Proc. 26th Large Installation Syst. Admin. Conf.*, 2012, pp. 111–122.
6. K. Gopalan, R. Kugve, H. Bagdi, Y. Hu, D. Williams, and N. Bila, "Multi-hypervisor virtual machines: Enabling an ecosystem of hypervisor-level services," in *Proc. Conf. Usenix Annu. Tech. Conf.*, 2017, pp. 235–249.
7. B. Fekade, T. Maksymyuk, and M. Jo, "Clustering hypervisors to minimize failures in mobile cloud computing," *Wireless Commun. Mobile Comput.*, vol. 16, no. 18, pp. 3455–3465, 2016.
8. S. Ren, Y. Zhang, L. Pan, and Z. Xiao, "Phantasy: Low-latency virtualization-based fault tolerance via asynchronous prefetching," *IEEE Trans. Comput.*, vol. 68, no. 2, pp. 225–238, Feb. 2019.
9. Q. Zhang *et al.*, "iBalloon: Efficient VM memory balancing as a service," in *Proc. IEEE Int. Conf. Web Serv.*, Jun. 2016, pp. 33–40.
10. W. Zhao, Z. Wang, and Y. Luo, "Dynamic memory balancing for virtual machines," *SIGOPS Oper. Syst. Rev.*, vol. 43, no. 3, pp. 37–47, Jul. 2009.
11. D. Patterson and J. Hennessy, *Computer Organization and Design RISC-V Edition*, 1st ed. San Mateo, CA, USA: Morgan Kaufmann, 2017.
12. J. Zhang, Y. Liu, H. Li, X. Zhu, and M. Chen, "PTAT: An efficient and precise tool for tracing and profiling detailed TLB misses," *ACM Trans. Embedded Comput. Syst.*, vol. 17, no. 3, pp. 62:1–62:17, May 2018.
13. Z. Guz *et al.*, "NVMe-over-fabrics performance characterization and the path to low-overhead flash disaggregation," in *Proc. ACM Int. Syst. Storage Conf.*, 2017, pp. 16:1–16:9.
14. S. Gugnani, X. Lu, and D. K. D. Panda, "Swift-X: Accelerating OpenStack swift with RDMA for building an efficient HPC cloud," in *Proc. 17th IEEE/ACM Int. Symp. Cluster, Cloud Grid Comput.*, 2017, pp. 238–247.
15. R. Mathá, S. Ristov, T. Fahringer, and R. Prodan, "Simplified workflow simulation on clouds based on computation and communication noisiness," *IEEE Trans. Parallel Distrib. Syst.*, vol. 31, no. 7, pp. 1559–1574, Jul. 2020.
16. H. Nemati, S. V. Azhari, and M. R. Dagenais, "Host hypervisor trace mining for virtual machine workload characterization," in *Proc. IEEE Int. Conf. Cloud Eng.*, 2019, pp. 102–112.
17. X. Chen *et al.*, "Refinery swap: An efficient swap mechanism for hybrid DRAM-NVM systems," *Future Gener. Comput. Syst.*, vol. 77, no. C, pp. 52–64, 2017.
18. T. Mason, T. D. Doudali, M. Seltzer, and A. Gavrilovska, "Unexpected performance of Intel Optane dc persistent memory," *IEEE Comp. Architecture Lett.*, vol. 19, no. 1, pp. 55–58, Jan.–Jun. 2020.
19. Intel Corporation, "Page modification logging for virtual machine monitor white paper," 2015, Accessed on: Sep. 20, 2020. [Online]. Available: <https://www.intel.com/content/www/us/en/processors/page-modification-logging-vmm-white-paper.html>

**SASKO RISTOV** is currently a Postdoctoral University Assistant with the University of Innsbruck, Innsbruck, Austria. His research interests include performance modeling, optimization, scheduling, and resource management in distributed and parallel systems. He received the Ph.D. degree from Sts. Cyril and Methodius University, Skopje, North Macedonia, in 2012. He is the corresponding author of this article. Contact him at [sashko@dps.uibk.ac.at](mailto:sashko@dps.uibk.ac.at).

**THOMAS FAHRINGER** is currently a Full Professor of computer science heading the Distributed and Parallel Systems Group, University of Innsbruck, Innsbruck, Austria. His research interests include software architectures, programming paradigms, compiler technology, performance analysis, and prediction for parallel and distributed systems. He received the Ph.D. degree in 1993 from the Vienna University of Technology, Vienna, Austria. Contact him at [tf@dps.uibk.ac.at](mailto:tf@dps.uibk.ac.at).



**RADU PRODAN** is currently a Full Professor of distributed systems with the Institute of Software Technology, University of Klagenfurt, Klagenfurt, Austria. His research interests include performance, optimization, and resource management tools for parallel and distributed applications. He received his Ph.D. degree in 2004 from the Vienna University of Technology, Vienna, Austria. He is a member of IEEE. Contact him at [radu@itec.aau.at](mailto:radu@itec.aau.at).

**MAGDALENA KOSTOSKA** is currently an Associate Professor with the Sts. Cyril and Methodius University, Skopje, North Macedonia. She received the Ph.D. degree from the Ss. Cyril and Methodius University in 2014. Her research interests include cloud computing and Internet of Things. Contact her at [magdalena.kostoska@finki.ukim.mk](mailto:magdalena.kostoska@finki.ukim.mk).

**MARJAN GUSEV** is currently a Full Professor with the University Sts. Cyril and Methodius, Skopje, North Macedonia. He received the Ph.D. degree from University of Ljubljana, Ljubljana, Slovenia, in 1992. His research interests include Internet of Things, cloud computing, and eHealth solutions. Contact him at [marjan.gushev@finki.ukim.mk](mailto:marjan.gushev@finki.ukim.mk).

**SCHAHRAM DUSTDAR** is currently a Full Professor of Computer Science heading the Distributed Systems Group, TU Wien, Vienna, Austria. His work focuses on Internet technologies. He is an IEEE Fellow, a member of the Academia Europaea, and an ACM Distinguished Scientist. Contact him at [dustdar@dsg.tuwien.ac.at](mailto:dustdar@dsg.tuwien.ac.at).

## ADVERTISER INFORMATION

### Advertising Coordinator

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

### Advertising Sales Contacts

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

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

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

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

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

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

Heather Bounadies  
Email: [hbuonadies@computer.org](mailto:hbuonadies@computer.org)  
Phone: +1 623-233-6575

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

Marie Thompson  
Email: [marie.thompson@computer.org](mailto:marie.thompson@computer.org)  
Phone: +1 714-813-5094

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

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

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


**Submit your paper today!**

Visit [www.computer.org/oj](http://www.computer.org/oj) to learn more.



# Emerging Technologies for Quantum Computing

Jonathan M. Baker , University of Chicago, Chicago, IL, 60637, USA

Frederic T. Chong , University of Chicago, Chicago, IL, 60637, USA and also Super.tech, Chicago, IL, 60615, USA

*Despite promising proof of concept demonstrations, currently available quantum hardware suffers from fundamental scalability limitations. The field is relatively new and it is imperative to consider emerging technologies in the space and evaluate their unique tradeoff spaces to determine how to best close the gap between current devices and target applications. Here, we explore three recent developments on this front. First, we consider extensions to currently available hardware, which allow the use of higher level states, beyond the usual binary, which when used temporarily can confer circuit-level advantage. Second, we consider the use of superconducting resonant cavities to reduce hardware requirements to implement quantum error correction protocols. Finally, we consider the use of neutral atoms, which offer unique strengths and weaknesses. It is valuable to evaluate new technology early and often to determine the best path toward scalability.*

Quantum computing is an emerging technology, so it may seem some sort of hyperbole to consider new emerging technologies for this paradigm, but we are already at the horizon of some potentially game-changing capabilities. Although current machines have shown impressive success with devices based upon trapped ions and superconducting transmons, it is unclear what the eventual winning technologies will be. In this article, we will consider neutral atoms and superconducting resonant cavities and extensions to currently available technology, and their implications for quantum architectures.

Quantum hardware is still in its relative infancy, boasting tens of qubits as opposed to the thousands to millions needed to execute important algorithms for unordered search and factoring.<sup>1,2</sup> Most available systems have struggled to scale beyond their prototypes while simultaneously suppressing gate errors and increasing qubit coherence times. This limits the types of programs which can execute effectively, let alone perform error correction. It is unclear whether

systems composed of superconducting qubits or trapped ions, the major industry players will take the lead.

Device success is predicated on the increase in the number of qubits, reduction of gate errors below error correction thresholds, and increase in qubit coherence times. In the near term, this translates into larger, deeper programs with improved output distributions. But, in the long term, this translates into qubits which are protected from noise inherent in operating quantum systems in the form of encoded logical qubits.

In recent years, there have been a number of improvements throughout the compilation pipeline both close to the hardware and high-level circuit optimization. These optimizations aim to reduce gate counts, circuit depth, and communication costs as proxies for increasing the size and quality of programs executable on currently available hardware.

An alternative approach is to evaluate the viability and tradeoffs of new quantum technology and associated architectures. Despite tremendous efforts at both the hardware and software level to minimize the effects of noise, it is unclear if any of the available hardware will be able to scale as needed and no clear winner on underlying hardware has emerged. Even further, it is unknown whether this hardware is best

suiting to execute the desired programs and while it is often the case that we adapt compilation to the hardware, i.e., transforming applications into the right shape for execution, an alternative approach is to explore how to design new architectures which are better suited for the applications we want to run.

Evaluating new hardware technology at the architectural level is decidedly different, though intrinsically coupled to the development of quantum hardware. Device physicists' goal is often to demonstrate the *existence* of high quality qubits and operations in prototypes, while the architect's goal is to evaluate the systems level ramifications of this technology, exploring the inherent trade-off spaces to determine viability in the near and long term. Perhaps most critically, this architectural design exploration leads to important insights about what is most important for hardware developers to optimize. For example, if some limitations can be effectively mitigated via software, hardware developers can focus on other more fundamental issues. This process of codesign, by evaluating new technology early and often, is central to accelerating scalability.

In this work, we detail three key developments on this front. First, we discuss the temporary use of qudits which can be used to accelerate key circuit components, an example of evaluating a fundamental architectural question of computing radix.<sup>3,4</sup> Second, we explore the use of local "memory-equipped" superconducting qubits to reduce hardware requirements to implement error correction, an example of application-driven hardware design.<sup>5</sup> Finally, we address the use of neutral atoms as a competitive alternative to industry focuses, an example of using software to mitigate fundamental hardware limitations to both guide development and accelerate scalability.<sup>6</sup>

### EXTENDING THE FRONTIER VIA INTERMEDIATE QUDITS

Quantum computation is typically expressed as a two-level binary abstraction using qubits. However, superconducting and trapped ion qubits based quantum architectures are not intrinsically binary and have access to an infinite spectrum of discrete energy levels. Typical implementations actively suppress higher level states. Some gate implementations have made use of higher level states to implement multiqubit interactions, while others are designed to actively mitigate leakage, a source of error which leaves the qubit state left in higher level states.

Higher radix computation has been explored previously in classical computation, but is usually used in specialized applications and general computation


obtains only limited (constant) advantage. Similar work in quantum computation has demonstrated a constant advantage; by expressing an  $N$  qubit circuit as one using  $M = N/\log_2(d)$  qudits, where a qudit is a  $d$ -level system.<sup>7</sup>

This advantage is still critical. Both qubits and gates between them are error prone. As devices scale with more qubits, the relative connectivity of these qubits decreases requiring an increasing number of error-prone gates to interact arbitrary pairs of qubits. Therefore, reducing the hardware resource requirement enables larger programs to be executed while requiring fewer gates.

Full translation from one radix to another (for example binary to ternary) offers constant advantage, however, more clever usage can offer even more. Many quantum algorithms make use of ancilla, additional free bits which are used to store temporary information. In some cases, they can provide asymptotic improvements in the depth of circuit decompositions, highlighting an important space-time tradeoff in quantum programs. By using additional space, in the form of additional qubits or devices, we can reduce the total execution time of the program. When programs are time-limited, for example, by prohibitive coherence times, it is critical to decompose circuits to minimize depth. However, this requires that we maintain a delicate balance as using ancilla limits the total size of programs, which can be executed if a large portion of qubits must be reserved.

Real quantum hardware will have a limited number of qubits so it is critical to make the most of them to enable computation of larger, more useful programs sooner. An alternative approach to full program translation is to make use of qudit states *temporarily* during binary computation, which extends the number of available computational qubits by reducing the ancilla requirements of key circuit elements while still maintaining low-depth decompositions. By accessing higher level states, the computation is subject to a wider variety of errors. If used properly, the amount gained outweighs this cost. Importantly, most quantum hardware does not currently support the execution of multiqubit gates, those with more than two inputs, and instead they must be decomposed into circuits using only one and two input operations.

Here we detail two specific uses. The first is a generalized Toffoli decomposition, a circuit which computes the reversible AND of many input bits.<sup>4</sup> The most efficient decompositions use many ancilla to temporarily store the ANDs between pairs of inputs, recursively. Rather than using a full extra qubit, we temporarily allow our qubits to access the  $|2\rangle$  state




**FIGURE 1.** Toffoli AND using temporary qutrits (bottom). The value in the circle indicates on which value to execute and the symbol in the box indicates which operation is executed, for example to add 1 modulo 3. Typical Toffoli decompositions use 6 two-qubit gates. Figure from Gokhale *et al.*<sup>4</sup>

and store the temporary AND results *locally*. The simplest version of this is to look at the Toffoli gate. First we can execute a 1-controlled +1 gate. This will cause the second input to be in the state  $|2\rangle$  if and only if *both* of the inputs were  $|1\rangle$  to begin. Then, executing a 2-controlled X gate onto the target will flip the target if and only if both inputs were  $|1\rangle$  which is exactly what a Toffoli gate should do. Following this operation, we want to ensure our inputs return to being only a superposition between the lowest two levels (i.e., return to being a qubit), so we perform some uncomputation. This circuit construction is found in Figure 1. The idea is similar in the general case—store the ANDs of inputs locally as  $|2\rangle$  rather than on another ancilla.

It turns out with temporary use of qutrits, we can obtain the same asymptotic depth and gate count as the most efficient decomposition using ancilla without using any ancilla. Specifically, we obtain a logarithmic depth decomposition using no additional space in the form of extra devices. The best known decomposition using only qubits requires linear depth with a large coefficient making it impractical to use. While effective, this strategy has fairly narrow uses requiring hand optimization to be effective.

Second is a decomposition of a reversible adder, which computes the reversible sum of the two inputs onto the second input. With a technique called “compression” we can *generate* the required ancilla in place to obtain logarithmic depth, the best known depth for decompositions with ancilla.<sup>3</sup> The simplest example is to consider the storage of two qubits into a single ququart (four level system). The two qubits and the single ququart can each be written as a superposition of four basis states. Therefore, all of the information of the two qubits can be stored in a single ququart. By allowing one of the qubit’s to access two additional levels, we store all of the information locally leaving the other qubit in the  $|0\rangle$  state, effectively an ancilla bit. This technique means we can generate ancilla local to the computation, using unused qubits



**FIGURE 2.** The compression of 2 qubits into a single ququart and generating an ancilla,  $|0\rangle$ . input two qubits, A and B, and produces a single ququart and an ancilla  $|0\rangle$ . To retrieve the stored information we can do the inverse of this operation using *any* ancilla for the second qubit. Using this type of compression circuit can produce clean ancilla on demand by storing unused data temporarily in higher states. Figure from Baker *et al.*<sup>3</sup>

nearby. A circuit for this compression is found in Figure 2.

This technique has been shown to be powerful for arithmetic circuits. For example, addition circuits can be decomposed in logarithmic depth without any external ancilla. The circuit is broken in a constant number of blocks, where each block is decomposed individually using the other unused blocks to generate the ancilla needed for an efficient decomposition. This block-based construction along with compression is very powerful. If efficient circuit decompositions are known for qubit circuits requiring ancilla and need only a constant amount of information to move between blocks, we can use this compression technique to obtain the same asymptotic depth as the known decomposition.

Both of these techniques free up more of our limited hardware for computation, rather than dedicating device space as ancilla which inherently limits the maximum size of computation which can be performed. Current hardware is often calibrated for use only with qubits but has higher level states available already. While classically multivalued and mixed-radix computing is niche, it is important to evaluate the architectural ramifications of these strategies for quantum computation. By temporarily accessing already available higher states, we can accelerate common circuit components and reduce space overhead extending the frontier of what can be computed.

### VIRTUALIZING LOGICAL QUBITS WITH LOCAL QUANTUM MEMORY


Current quantum architectures do not tend to make a distinction between memory and processing of quantum information. These architectures are viable, however, as more and more qubits are needed the scalability challenges become apparent. For example, many industry players make use of superconducting

transmon qubits, which suffer from fabrication inconsistency and crosstalk during parallel operations. To scale to the millions of qubits needed for error correction, a memory-based architecture can be used to decouple qubit-count from transmon count.

It is imperative to explore the use of technologies outside of what is currently commercialized or made available via cloud services. For example, recently realized qubit memory technology which stores qubits in superconducting resonant cavities may help to realize exactly this memory-based architecture.<sup>8,9</sup> In these devices, qubit information can be stored in the cavity and when an operation needs to be performed it can be transported to the attached transmon. Local memory is not free. Stored qubits cannot be operated on directly and all operations must be mediated via the transmon by first loading the information. This further prohibits parallel operations in qubits in the same cavity requiring serialization. The realization of this technology alone is not sufficient to understand its viability and instead dedicated architectural studies are needed.

Here we explore a proof-of-concept demonstration of its viability by virtualizing surface code tiles (one example of quantum error correction) in a 2.5-D memory-based architecture.<sup>3</sup> Logical qubits can be stored at unique virtual addresses in memory cavities when not in use and are loaded to a physical address in the transmons for computation or to correct errors, similar to DRAM refresh. Figure 3 depicts the proposed architecture with surface tiles mapped to unique virtual addresses. To minimize the use of transmons, the goal is to take logical qubits stored in a plane and find an embedding of that plane in 3-D, where the third dimension is limited to a finite size. The simplest embedding procedure is to slice the plane to form patches, one for each logical qubit and stack them into the layers so that each shares the same physical address and therefore transmons. This procedure is not the most space efficient, since the ancilla needed to detect errors occupy their own unique transmons which are unnecessary. A compact embedding reduces the physical transmon cost by an additional 2× at the cost of some additional time to detect errors.

The initial benefit is clear—this design requires many fewer physical transmons by storing many logical qubits in the same physical location. There are other advantages such as a faster transversal CNOT, which is traditionally executed using a sequence of many primitive surface code operations. Furthermore, this design requires fewer total qubits to distill  $|T\rangle$  states, which are commonly used for universal quantum computation. These improvements translate to



**FIGURE 3.** A fault-tolerant architecture with random-access memory local to each transmon. On top is the typical 2-D grid of transmon qubits. Attached below each data transmon is a resonant cavity storing error-prone data qubits (shown as black circles). This pattern is tiled in 2-D to obtain a 2.5-D array of logical qubits. The key innovation here is to store the qubits that make up each logical qubit (shown as checkerboards) spread across many cavities to enable efficient computation. Figure from Baker *et al.*<sup>5</sup>

gains in important algorithms by reducing execution time and resource requirements, specifically a 1.22× speedup for Shor’s algorithm or allowing it to run on smaller hardware.

Physical qubit savings alone is not sufficient to guarantee its competitiveness. We must ensure the error thresholds (approximately how good physical operations need to be in order for errors to be efficiently corrected) are as good, if not better, than standard implementations and that all of the necessary operations, such as single qubit gates and an entangling two qubit gate, can be executed. Embedding the surface code in the 2.5-D architecture permits all of the same lattice surgery operations to be executed and error detection, although with some delay between each cycle as only one logical qubit can undergo a round of error correction at a time. Despite this additional delay, the thresholds of the embedded code are comparable to a standard surface code implementation.

Error correction protocols are essential for the execution of large-scale quantum programs. The surface code is one such code, designed with currently available architectures in mind. It is a low threshold code which requires only local operations on a 2-D grid.<sup>10,11</sup> However, this application is better matched with this 2.5-D architecture, which allows qubits to be virtualized and stored in local memory. This highlights a

distinct role of the computer architect to discover fortuitous matches between application and emerging hardware technology. In this case, we match error correcting codes to an architecture which reduces resource requirements.


## ARCHITECTURAL TRADEOFFS IN EMERGING TECHNOLOGY—NEUTRAL ATOMS

Current hardware implementations face unique scalability challenges, such as high gate error rates or high crosstalk error with densely connected qubits, or other fundamental challenges in controllability. While these devices have been useful as proof of concept demonstrations of small-scale near-term algorithms, it is unclear whether any of them in present form will be able to execute the large-scale computation needed for quantum speedup. These limitations are fundamental and current compilation approaches to reduce gate counts and depth are insufficient for long-term scalability.

Evaluating the viability of new hardware is essential.<sup>6</sup> Architectural studies which fully explore their unique tradeoff spaces is key for finding the best way to accelerate beyond prototypes. One such alternative to superconducting qubits or trapped ions is neutral atoms.<sup>12</sup> These arrays of individually trapped atoms can be arranged in one, two, or even three dimensions.<sup>13–16</sup>

This technology offers distinct advantages, which make it an appealing choice for scalable quantum computation. Mainly, atoms can interact at long distances, which leads to reduced communication overhead yielding lower gate count and depth programs. Furthermore, these devices may be able to execute high fidelity multiqubit ( $\geq 3$  operands) operations natively. For example, a Toffoli gate can be implemented directly between three qubits without needing an expensive decomposition.<sup>17</sup> These benefits do not come for free. Long range interaction requires proportionately large regions of qubits to be restricted leading to reduced parallelism requiring gates to be serialized with mitigating some of the reduced depth benefits. These unique properties are depicted in Figure 4.

Evaluating new quantum computing technology is especially challenging. Neutral atoms offer some clear advantages over other gate-based models, however, the current demonstrations display gate errors and coherence times which are worse than competitors. With physical properties lagging years behind, the appeal is dampened, but there is no fundamental



**FIGURE 4.** Examples of interactions on a neutral atom device. (a) Interactions of various distances are permitted up to a maximum. Gates can occur in parallel if their zones do not intersect. The interaction marked with green checks can occur in parallel with the middle interaction. (b) The maximum interaction distance specifies which physical qubits can interact. Compiler strategies suited for this variable distance are needed for neutral atom architectures. (c) Neutral atom systems are prone to sporadic atom loss. Efficient adaptation to this loss reduces computation overhead. Figure from Baker *et al.*<sup>6</sup>

limitation to these properties. When studied with all else being equal, the unique properties of neutral atom hardware claim a dominant position.

Exploring the use of new hardware also serves a critical role in the development of the platform itself. Neutral atom systems suffer a potentially crippling drawback—atoms can be lost both during and between program execution. When this happens, measurement of the qubits at the end of the run is incomplete and requires the output of the run to be discarded. The standard approach to coping with this loss is simply to run the program again after reloading all of the atoms in the array. Typically programs are run thousands of times and so each run which must be discarded incurs a twofold cost—we have to perform an array reload and we have to execute an additional run. This is especially bad when execution time is limited by atom reload rate, which is significantly longer than the actual program run.

Fortunately, software solutions can effectively mitigate this increased run time overhead. The idea is to simply keep track of when atoms are lost during the execution of a program, which can be detected quickly via fluorescence. If an atom loss occurs, we throw away that run and adjust how the program is compiled to the hardware. Ideally we do not want to recompile the entire program as this will be expensive, often more expensive than the reload itself, which should be treated as the worst case. Better solutions simply adjust the compiled program by adjusting the placements of the qubits and addition in a small number of extra communication operations.

*THERE IS A LARGE GAP BETWEEN CURRENTLY AVAILABLE QUANTUM COMPUTING HARDWARE AND THE APPLICATIONS WE WANT TO RUN. HOWEVER, QUANTUM COMPUTING SYSTEMS HAVE MADE TREMENDOUS STRIDES IN RECENT YEARS IN A SIGNIFICANTLY COLLABORATIVE AND INTERDISCIPLINARY EFFORT BETWEEN PHYSICISTS, MATHEMATICIANS, COMPUTER SCIENTISTS, AND MANY MORE.*

The best strategies directly take advantage of the neutral atom benefits. For example, these architectures support long range interactions, though most communication reduction does not require the use of the maximum allowed interaction distance. Therefore, by compiling the program to less than the maximum gives us flexibility in the final compiled program as atoms are lost over time requiring only a virtual remapping of program qubits and no extra operations. The best coping mechanisms sustain large numbers of loss by minimizing the total number of reloads, having low compilation time overhead, and adding a small number of gates (minimizing the number of error prone operations added).

Atom loss is a fundamental limitation of neutral atom architectures. Probabilistic loss of atoms is inherent in the trapping process itself and prior hardware studies have focused on *hardware* solutions to reduce this probability of loss. However, software solutions can effectively mitigate problems due to loss. Demonstrating effective mitigation strategies is critical to the overall development of the platform—by solving fundamental problems at the systems level with software, hardware developers can focus on solving and optimizing other problems. Codesign of quantum systems is key to accelerate the advancement of quantum computing technology.

### CONCLUDING REMARKS

There is a large gap between currently available quantum computing hardware and the applications we want to run. However, quantum computing systems have made tremendous strides in recent years in a significantly collaborative and interdisciplinary effort between physicists, mathematicians, computer scientists, and

many more. Here we have discussed some of the central roles of the computer architect in the development of scalable hardware. First, we must determine which are the right abstractions for the job. Typically, quantum computation has been expressed exclusively in terms of two-level systems, however, this abstraction is artificial and *temporarily* making use of higher level qudit states yields efficient low depth circuit decompositions freeing up additional space on the hardware allowing us to execute larger programs sooner. New technologies which more easily allow access to these states can be powerful. Second, we must use applications to guide the design architectures, which best implement them. New, memory-equipped transmon technology has a fortuitous match with lattice surgery based surface codes, trading some serialization for an efficient error correction implementation. This architecture reduces hardware requirements enabling demonstrations of error correction sooner. Third, we must evaluate new technology as it is developed at the systems level to determine its viability for long-term scalability and develop software solutions to new technologies' fundamental limitations to guide hardware developers. For example, neutral atoms offer many unique advantages while also suffering from unique disadvantages and it is important to understand if architectures based on this technology are viable in either the near or long term. Furthermore, software can mitigate the downsides of atom loss, lending hardware developers the space to work on other aspects of the hardware. New technology inspired each of these ideas but the techniques developed here are general and could be applied to other similar technology as they emerge. 🤖

### ACKNOWLEDGMENTS

This work was supported in part by EPIQC, an NSF Expedition in Computing, under Grant CCF-1730449; in part by STAQ under Grant NSF Phy-1818914; in part by DOE Grants DE-SC0020289 and DESC0020331; and in part by NSF OMA-2016136 and the Q-NEXT DOE NQI Center. Disclosure: F. Chong is Chief Scientist at Super.tech and an advisor to Quantum Circuits, Inc.

### REFERENCES

1. L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. Annu. ACM Symp. Theory Comput.*, 1996, pp. 212–219.
2. P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM J. Comput.*, vol. 26, no. 5, pp. 1484–1509, Oct. 1997.



3. J. M. Baker, C. Duckering, and F. T. Chong, "Efficient quantum circuit decompositions via intermediate qudits," in *Proc. IEEE 50th Int. Symp. Multiple-Valued Logic*, 2020, pp. 303–308.
4. P. Gokhale, J. M. Baker, C. Duckering, N. C. Brown, K. R. Brown, and F. T. Chong, "Asymptotic improvements to quantum circuits via qutrits," in *Proc. 46th Int. Symp. Comput. Archit.*, 2019, pp. 554–566.
5. J. M. Baker, C. Duckering, D. Schuster, and F. Chong, "Virtual logical qubits: A compact architecture for fault-tolerant quantum computing," *IEEE Micro*, vol. 41, no. 3, pp. 95–101, May/June 2021.
6. J. M. Baker, A. Litteken, C. Duckering, H. Hoffmann, H. Bernien, and F. T. Chong, "Exploiting long-distance interactions and tolerating atom loss in neutral atom quantum architectures," in *Proc. ACM/IEEE Annu. Int. Symp. Comput. Archit.*, 2021.
7. A. Pavlidis and E. Floratos, "Arithmetic circuits for multilevel qudits based on quantum Fourier transform," 2017, *arXiv:1707.08834*.
8. C. T. Hann *et al.*, "Hardware-efficient quantum random access memory with hybrid quantum acoustic systems," 2019, *arXiv:1906.11340*.
9. R. Naik *et al.*, "Random access quantum information processors using multimode circuit quantum electrodynamics," *Nature Commun.*, vol. 8, no. 1, 2017, Art. no. 1904.
10. C. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, "Surface code quantum computing by lattice surgery," *New J. Phys.*, vol. 14, no. 12, 2012, Art. no. 123011.
11. D. Litinski, "A game of surface codes: Large-scale quantum computing with lattice surgery," *Quantum*, vol. 3, p. 128, 2019.
12. M. Saffman, "Quantum computing with atomic qubits and Rydberg interactions: Progress and challenges," *J. Phys. B, At., Mol. Opt. Phys.*, vol. 49, no. 20, 2016, Art. no. 202001.
13. M. Endres *et al.*, "Atom-by-atom assembly of defect-free one-dimensional cold atom arrays," *Science*, vol. 354, no. 6315, pp. 1024–1027, 2016.
14. D. Barredo, S. De Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, "An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays," *Science*, vol. 354, no. 6315, pp. 1021–1023, 2016.
15. H. Kim, W. Lee, H.-G. Lee, H. Jo, Y. Song, and J. Ahn, "In situ single-atom array synthesis using dynamic holographic optical tweezers," *Nature Commun.*, vol. 7, no. 1, pp. 1–8, 2016.
16. D. Barredo, V. Lienhard, S. De Léséleuc, T. Lahaye, and A. Browaeys, "Synthetic three-dimensional atomic structures assembled atom by atom," *Nature*, vol. 561, no. 7721, pp. 79–82, 2018.
17. H. Levine *et al.*, "Parallel implementation of high-fidelity multiqubit gates with neutral atoms," *Phys. Rev. Lett.*, vol. 123, no. 17, 2019, Art. no. 170503.

**JONATHAN M. BAKER** is currently working toward a Ph.D. degree in computer science at the University of Chicago, Chicago, IL, USA. He is currently interested in many interdisciplinary aspects in the design and evaluation of quantum computer architectures, spanning the software-hardware stack from application to device control. Baker has degrees in computer science, chemistry, and mathematics from the University of Notre Dame, Notre Dame, IN, USA. Contact him at [jmbaker@uchicago.edu](mailto:jmbaker@uchicago.edu).

**FREDERIC T. CHONG** is the Seymour Goodman Professor with the Department of Computer Science, University of Chicago, Chicago, IL, USA and the Chief Scientist at Super.tech, Chicago, IL, USA. He is also Lead Principal Investigator for the EPiQC Project (Enabling Practical-scale Quantum Computing), an NSF Expedition in computing. He was a faculty member and Chancellor's Fellow at the University of California Davis in 1997–2005. He was also a Professor of Computer Science, Director of Computer Engineering, and Director of the Greenscale Center for Energy-Efficient Computing at the University of California Santa Barbara from 2005 to 2015. Chong received a Ph.D. degree from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1996. He is a recipient of the NSF CAREER Award, the Intel Outstanding Researcher Award, and 10 best paper awards. Contact him at [chong@cs.uchicago.edu](mailto:chong@cs.uchicago.edu).

# Quantum Computing

Jose Luis Hevia, Guido Peterssen, Christof Ebert, and Mario Piattini

This article originally  
appeared in  
**Software**  
vol. 38, no. 5, 2021

## FROM THE EDITOR

Quantum computing has become a reality. Quantum computers are available to everybody via cloud service or simulation. Toolkits are available that invite practitioners to start their own quantum software projects and thus get used to this novel technology. In this article we evaluate technologies to help developers to start their own quantum software business. Practical guidance is provided from our own quantum technology projects. I look forward to hearing from you about this column and the technologies that matter most for your work.—Christof Ebert

Have you ever tried to retrieve that forgotten key code for your suitcase? After one year without traveling, many of us found themselves having forgotten the combination and manually trying all permutations. The same situation, but more complex, would be to systematically try identifying that forgotten access code for an online app that you had not used for a while. Cyberattackers are doing exactly this, of course, at high speed and with increasing computing performance. The recommended security key length is thus getting longer by the year. Yet, the stepwise process to achieve this is tedious or consumes lots of computing power. Now imagine that all of these possible states could be tried in a single step. This would be good for your own number lock, but frightening for our security infrastructure.

The promise of quantum computing is to vastly accelerate such complex algorithms.<sup>1</sup> Today, even supercomputers fail on high algorithmic complexity because many algorithms still work in sequences which build on results of a previous step. Of course, they use massively parallel hardware and algorithms, but networking imposes limits as does the memory needed to hold the myriad combinations of real-world problems.

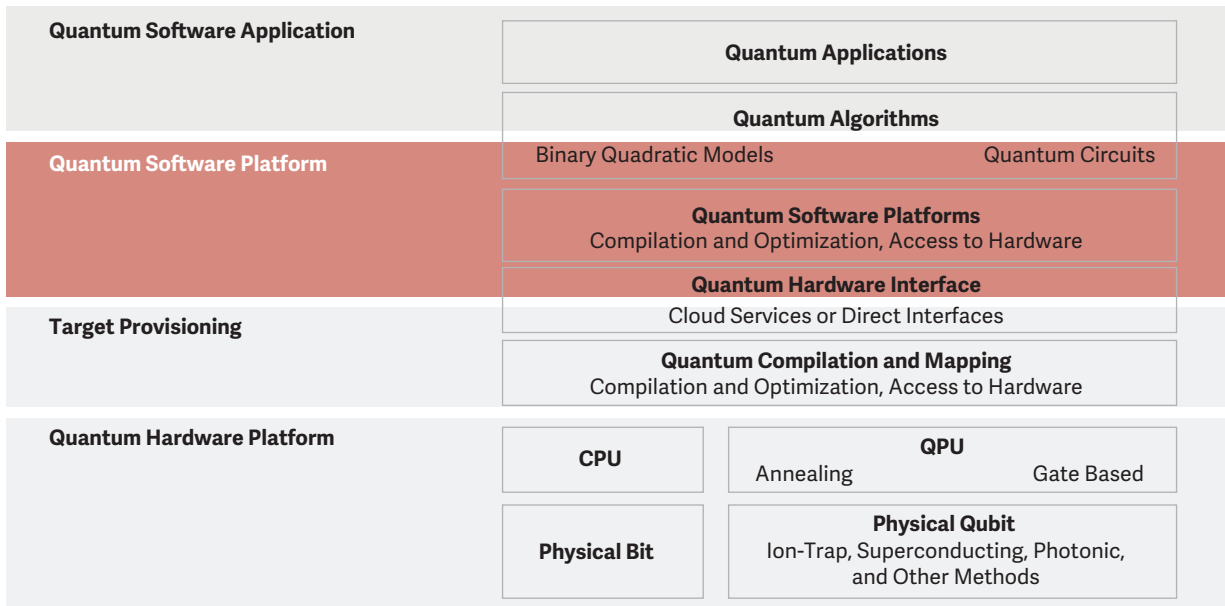
Quantum computing has left the research domain and is ready to use in industry practice. It will rapidly advance industry applications in fields such as data science, pattern recognition, and cybersecurity.<sup>1-4</sup> Yet the actual software development for quantum computing is hampered by a lack of appropriate methods and insufficiently scalable technology.<sup>2</sup>

*ACTUAL SOFTWARE DEVELOPMENT  
FOR QUANTUM COMPUTING  
IS HAMPERED BY A LACK OF  
APPROPRIATE METHODS AND  
INSUFFICIENTLY SCALABLE  
TECHNOLOGY.*

## QUANTUM COMPUTING

Quantum computers use atomic-scale effects, such as electron spin, as underlying information.<sup>1</sup> Quantum computing uses what are called *quantum bits* or *qubits*: bits that are held in superposition and use quantum principles to complete calculations. A binary digit is always in one of two definite states, that is, either zero or one. Qubits are in a superposition of these classic binary states of zero and one.

Superposition is the ability of qubits to be in more than one physical state at a time, which allows us to



**FIGURE 1.** The quantum software and hardware stack.

parallelize combinations. Multiple qubits can also become entangled. If you measure the state of one qubit entangled with another qubit, the outcome of measuring the other qubit is correlated in some way with the first, even if the two qubits are far apart. Superposition and entanglement are used together for quantum computation. Yet these effects also create practical challenges with real quantum computers: they require sophisticated lab environments; also, the information might decay when the state of the system is captured.

How do we deduce the result from these superimposed states? Many quantum algorithms first create superpositions of an exponentially large number of logical states. Interference is used in that the incorrect answers of the specific problem destructively interfere and no longer appear in the final output, leaving behind only the correct answer.

From an algorithmic point of view, quantum computing can solve problems of higher complexity than classical computing—faster and also with cost and energy savings. The first quantum computers were built in the late 1990s. For practical usage, we distinguish quantum simulators, in which the quantum algorithm is simulated on classical hardware (CPUs), and real quantum computers, with quantum processing units (QPUs), in which qubits are built using a wide

variety of technologies: ion-trap, superconducting, and photonic methods, among others.<sup>1,4,5</sup>

Within quantum computers we find mainly two categories (Figure 1): quantum annealing computers, such as the D-Wave computers (suitable for running optimization problems since finding the largest or smallest value of an indicator can translate into minimizing the energy of a system), or gate-based computers, such as those from Google, IBM, Rigetti, IonQ, and Honeywell.<sup>5</sup> All of this strongly affects the way in which applications are developed. Indeed, there are also two approaches: those based on building binary quadratic models for solving a problem or those based on the construction of quantum circuits based on gates.

Today quantum hardware vendors such as IBM, Rigetti, and Google deliver some 100 qubits on a laboratory scale.<sup>4,5</sup> This is impressive and demonstrates how fast the technology is evolving, but it is not yet sufficient to run actual software applications. Therefore, the quantum applications that we envisage today are separating the actual hardware stack from the software tier (Figure 1).

We expect that quantum computers will scale up at a pace like that of Moore's law. For the short term, a quantum network accessible by cloud services could show results from a software perspective. By connecting individual quantum devices, a quantum

*BECAUSE OF THE EXTREME PARALLELISM OF QUANTUM ALGORITHMS, SOME MASSIVE PARALLEL CHALLENGES, SUCH AS DATA SCIENCE AND PATTERN RECOGNITION, CAN BE ACCELERATED BY QUANTUM COMPUTING.*

supercomputer could be created. A bigger step forward is a quantum network based on entangled qubits for fast information exchange. Cybersecurity is an obvious application domain of such a network to facilitate quantum key distribution with a cryptography protocol relying on interlinked quantum particles.

## QUANTUM COMPUTING APPLICATIONS

Applications of quantum computing are manifold. Because of the extreme parallelism of quantum algorithms, some massive parallel challenges, such as data science and pattern recognition, can be accelerated by quantum computing.

Examples include, for example, identifying the optimal route of a delivery car or fleet of trucks to save on time and fuel costs. Or an investment company may need to balance its portfolio risk with numerous possible combinations of shares with different individual performances and related cluster risks. Pharmaceutical researchers need to simulate molecules to better understand drug interactions, even if they do it using only known constraints and reported deficiencies. The latter is our case study in “QHealth.”

On the dark side, massive parallel algorithms will also facilitate hacking any current cryptographic key with much less effort than is currently assumed. Shor’s algorithm can factor large prime numbers down into two smaller ones.<sup>6</sup> This is a very useful property for breaking encryption since the Rivest–Shamir–Adleman (RSA) system of encryption depends on factoring large prime numbers. Already today, major cybersecurity algorithms are anticipating such quantum hacking and vastly enhancing the key length. Postquantum cryptography has started to be researched with encryption techniques that would operate and not be broken even with much larger quantum computers. Most of the encryption systems in modern cryptocurrencies

are built on elliptic curve cryptography rather than RSA because elliptic curves are harder to crack than RSA—at least by classical computers. Current blockchain-based e-currencies thus use signatures that require the Elliptic Curve Digital Signature Algorithm (ECDSA). However, quantum computers seem to challenge ECDSA. With enough qubits, Grover’s algorithm can break elliptic curve cryptography even more easily than you might break RSA.<sup>6</sup> As Grover’s algorithm also accelerates mining, one further application is the evolution in bitcoin mining from GPUs, field-programmable gate arrays, and application-specified integrated circuits toward quantum computers.

Novel quantum computation protocols are currently developed toward enhanced security. In such protocols, the client will encrypt its data so that the host or cloud computer cannot learn anything about them yet can still perform the calculation. After the computation, the client will then decrypt the data again to get the real results of the calculation. Yet another application is a performance boost in network algorithms by using entangled qubits, which allows them to simultaneously calculate independent of their distance apart.<sup>1,7</sup> The latter field of study is not yet mature, with distances only in the meter range and the entangling of only a few qubits, but the effects would be overwhelming if future networking no longer needed physical networks.

Given our analogy with Moore’s law, large enough quantum computers will appear within a few years. Shor’s algorithm works with a quantum computer offering 10–100k qubits. Using Grover’s algorithm for database searching and hacking ECDSA will require some 100k qubits. All this assumes steady growth and the mastering of quantum-specific challenges, such as the noise and error rates caused by the inherent quantum effect that observing a superimposed quantum state will influence its result, as described by the Schrödinger’s cat thought experiment. Given the longevity of embedded computing and the exponential growth rate, now is the time to prepare our software and IT systems for the impacts of quantum computing such as postquantum cryptography.

## QUANTUM SOFTWARE DEVELOPMENT

To utilize quantum computing, quantum hardware

## QHEALTH

Quantum technology can be applied to multiple questions where data science meets algorithmic complexity. The aim of QHealth is to improve the quality of life of older adults. It correlates genetic and other variables related to a person's health history. The health history is analyzed as a function of the patient's drug consumption history, the reactions the older adult has experienced, and his/her physiological and genetic limitations.

The challenge in such an analysis involves the complexity of genetic precondition on one hand and also the number of drugs being used as part of the normal treatments of elderly persons. Even when looking only to the impacts of medication, there are multiple interactions and contraindications. For each active ingredient, in addition to variables such as genetic biomarkers, haplotypes, phenotypes, and so on, we must consider specific personal variables about the patient, such as sex, age, weight, blood pressure, recent drug history, and specific health impacts, among others. The underlying data analytics soon become intractable with classical computing.

QHealth builds a hybrid quantum system combining health-care applications and data analytics with quantum computing. Quantum technologies carry out optimizations and simulations whose realization in classical hardware is not possible in acceptable timescales. This hybrid system, in combination with classical health applications, will give its outputs to medical professionals involved in prescribing drugs to elderly adults. In a further extension, we also envisage application in the

case of younger persons with difficult drug treatment and health conditions, trying to reduce the negative impacts of drugs due to their correlation and mutual side-effects when used in combination. Using the case histories and the socioeconomic and genetic variables of the persons being analyzed, we can then also make recommendations for suitable drug treatments and provide risk assessment before they are prescribed.

Using quantum technology for health care will vastly increase the possibilities to assess and optimize medical treatment applications, especially for persons who need multiple medications for coexisting illnesses. The proposed approach that we currently industrialize not only improves life and medical treatment but also has a financial impact because it will optimize the investments that health systems make in financing drugs and address the adverse effects that drugs often generate.

QHealth is founded by the Center for the Development of Industrial Technology (CDTI) of the Ministry of Science and Innovation of Spain and the European Regional Development Fund, in the 2020 CDTI Missions Program, with a total budget of several million euros. It involves a multidisciplinary team of researchers and technologists from the aQuantum, Gloin, and Madrija companies and the University Institute of Biosanitary Research of Extremadura in collaboration with the Pharmacogenetics and Personalized Medicine Unit, the University of Extremadura, and the University of Castilla-La Mancha.

vendors offer full stacks for the development of quantum software. As those are typically hardware specific, there are also third-party suppliers that provide platforms that claim to be hardware agnostic.

The quantum software platforms are portrayed in red in Figure 1. They offer the following functions:

- › They provide users access to quantum computers to perform quantum computations via cloud services.
- › They provide abstractions between the underlying hardware and the actual software applications. This includes libraries to facilitate using the quantum computer either in simulation or as actual hardware.
- › They offer development kits and computational platforms to ramp up end-user proficiency.
- › They support software engineers in developing and testing their quantum algorithms.
- › They enhance the reliability and performance of physical quantum computers. An inherent weakness of any quantum computing system is the errors in the transition from digital to quantum states. Random errors can occur due to the currently used hardware. Error-correcting software increases the stability and reliability of quantum computers.

Table 1 provides an overview of the currently available quantum software technologies. Toolkits

**TABLE 1.** Quantum software development platforms.

Product Functionality	D-Wave Leap–Ocean	Fujitsu Quantum-Inspired Services	Google Cirq	IBM Quantum Experience and Qiskit	Microsoft Azure QDK and Q#	
URL	<a href="https://www.dwavesys.com/take-leap">https://www.dwavesys.com/take-leap</a>	<a href="https://www.fujitsu.com/es/services/business-services/digital-annealer">https://www.fujitsu.com/es/services/business-services/digital-annealer</a>	<a href="https://quantumai.google/cirq">https://quantumai.google/cirq</a>	<a href="https://quantum-computing.ibm.com">https://quantum-computing.ibm.com</a>	<a href="https://azure.microsoft.com">https://azure.microsoft.com</a>	
Hardware agnosticity	NO Only D-WAVE	NO Only Fujitsu	NO Only Google	NO Only IBM	YES	
Programming language	Python Platform-specific language	Python	Python, platform-specific language	Python QASM Platform-specific language	Q# Python	
Integrated development environment	LEAP for executing quantum algorithms	None. It depends on Jupyter and Python	None. It depends on Jupyter and Python	QEXPERIENCE for executing quantum algorithms	Visual Studio Code	
Optimization	YES	NO	YES	YES	NO	
Modularity	YES if Python is used	YES if Python is used	YES if Python is used	YES	YES	
Out of the box functions	NO	NO	NO	NO	YES	
Service integration	API for executing solvers as a service	API for executing solvers as a service	API for executing circuits as a service	API for executing circuits as a service	API for executing circuits as a service	
Third-party software	Jupyter Strangeworks QuantumPath	Jupyter Strangeworks QuantumPath	Jupyter Strangeworks Zapata Orquestra QuantumPath	Jupyter Strangeworks Zapata Orquestra QuantumPath	Jupyter Strangeworks Zapata Orquestra QuantumPath	
High-level libraries	YES	YES	YES	YES	YES	

URL: uniform resource locator; QDK: Quantum Development Kit; OOTB: out of the box; API: application programming interface.

*BUILDING AND EVEN USING A QUANTUM COMPUTER INVOLVES A HIGH INVESTMENT BECAUSE OF THE UNDERLYING QUANTUM HARDWARE STACK.*

from hardware suppliers are typically specific to their underlying hardware. Manufacturers provide both local simulators as well as cloud resources to access real machines.

Building and even using a quantum computer involves a high investment because of the underlying quantum hardware stack. Since actual quantum computing hardware is much too expensive and complex, most available software platforms are based on cloud services. However, there are very few manufacturers

capable of providing quantum services close to what is currently needed in terms of software business. Also, each manufacturer brings its own solutions, architectures, and specific hardware–software dependencies. To date there are no de facto standards for building an appropriate quantum software stack. In Figure 1 we have attempted to at least provide some abstraction levels between the different functional tiers.

Although there are many algorithms for quantum computers, it requires a good understanding of the underlying theory and technology to determine which algorithm can be used in a certain situation. Even if a suitable algorithm is transferred from traditional data science, its conversion into an executable program requires competence in the environment of the respective quantum computer, which data scientists and software engineers typically do not have.

Microsoft, IBM, and Google have their own

	Rigetti Forest	Xanadu– Strawberry Fields and Penny Lane	Orquestra	Quantum Inspire	QuantumPath	Quantum Programming Studio	Strangeworks QC
	<a href="https://www.rigetti.com/quantum-computing/">https://www.rigetti.com/quantum-computing/</a>	<a href="https://strawberryfields.ai/">https://strawberryfields.ai/</a>	<a href="https://www.zapatacomputing.com/orquestra/">https://www.zapatacomputing.com/orquestra/</a>	<a href="https://www.quantum-inspire.com/">https://www.quantum-inspire.com/</a>	<a href="https://www.quantumpath.es/">https://www.quantumpath.es/</a>	<a href="https://quantum-circuit.com/">https://quantum-circuit.com/</a>	<a href="https://strangeworks.com/">https://strangeworks.com/</a>
	NO Only RIGETTI	Partially IBM	YES, but few integrated providers	YES, mainly connected with IBM Quantum Experience.	YES	NO. The circuit is only directly exportable to Rigetti's hardware	YES
	Python QUIL Platform-specific language	Python	Python	Python cQASM	Python Q#	Multiple languages	Python
	FOREST for executing quantum algorithms	None It depends on Jupyter and Python	NO. It depends on Jupyter and Python	Quantum Experience	YES	YES	YES
	YES	YES	NO	YES	YES	NO	NO
	YES	YES if Python is used	YES, if Python is used	YES, if Python is used	YES	NO	YES
	NO	NO	NO	NO	YES	NO	YES
	API for executing circuits as a service	API for executing circuits as a service	NO	API for executing circuits as a service	API for executing circuits as a service	PARTIAL	PARTIAL
	Jupyter Strangeworks Zapata Orquestra QuantumPath	Jupyter Strangeworks	Jupyter	Jupyter QuantumPath	Jupyter Visual Studio Java	NO	NO
	YES	YES	YES	YES	YES	YES	YES

respective environments, namely, Q#, Qiskit, and Cirq, which use the Python programming language. Microsoft's Quantum Development Kit (QDK) delivers user-friendly code libraries, a debugger, and a resource estimator to assess how many qubits an algorithm will require. Each manufacturer provides its own access rules to the environments and its versions of approved languages. IBM offers access to a five-qubit machine free of charge. More powerful machines are available in its Quantum Network. Microsoft offers access to other companies' quantum computers through its Azure Quantum platform.

Two distinct development technologies are visible: quantum gates and quantum annealing. Most vendors offer an integrated development environment, but they are intended more as an environment for experiment and executing independent quantum algorithms/circuits than a business development environment. Most

of the toolkits also include some quantum software optimization features, but usually modules are unconnected elements in a traditional or online file system, such as GitHub, or http-accessible files.

Several third-party tools can be connected to these toolkits, and high-level libraries are included. These libraries are sets of extensions to the programming language that encapsulate the manufacturer's specific components in a high-level unit: data normalizations, circuit classes, gates, calculation functions and utilities, error control, and many more. They are included because of the R&D and large investments of each manufacturer in this technology. In addition to these valuable resources, suppliers add through the Internet extensive repositories of information, code, algorithms, training materials, and a long list of other types of resources that make access to their quantum technologies much easier.

Third-party software platforms, although they bring quantum resources closer to the business world, do not yet provide the necessary core elements in the lifecycle and architecture of hybrid systems. While their capabilities and tools are good, it is necessary to invest time and effort in investigating how to fit them into a complete rigorous software lifecycle, to improve the productivity and ensure quality quantum software development. These development environments are hardware agnostic as they are intended to serve as development tools for various end-user environments. They are evolving as these kinds of toolkits try to create, in most cases, an integrated development environment. Only a few of them provide optimization facilities and out-of-the-box (OOTB) functionality.

In our evaluation (Table 1) we reflect these different attributes and functionalities. OOTB functionality reflects whether the toolkit is standalone or whether it needs to install third-party software to be able to produce software professionally. Regarding service integration, all of the platforms provide an application programming interface (API)—in the case of quantum gate-based computers for executing quantum circuits as a service and in the case of quantum annealing ones for executing solvers as a service.

## CHALLENGES IN USING QUANTUM SOFTWARE

Designing software for quantum computers requires additional skills compared to creating software for traditional computers. To benefit from the fast pace of quantum hardware evolution, it is urgent that we mature the technology and methods for quantum software. It is not enough to stress the importance of quantum software;<sup>1</sup> we must go a step further and raise the awareness of quantum software engineering (QSE).<sup>5,7</sup> Distinguishing different layers of complex systems by simulation and networked smaller elements will allow us to target innovation in parallel for the underlying hardware and software.<sup>3</sup>

Quantum software with industry-scale performance, robustness, and reliability will mean a next level in software technology. We strongly believe that quantum computing could also bring a new “golden age” to software engineering.<sup>8</sup> But it is necessary to address all the challenges and opportunities faced in QSE<sup>7</sup> and adapt or create the necessary models,

standards, or methods to help us in the creation of new quantum systems and the migration of current ones.<sup>9</sup> One step in such advances is having the right development toolkits and knowing their characteristics.

Quantum software platforms and toolkits are difficult for practical industry usage. They do not bring much context support to the quantum algorithm generation, assuming that the quantum software engineer will know how to incorporate each product to its corresponding platform. In the meantime, collections of quantum software algorithms are available, such as the quite exhaustive quantum algorithm zoo.<sup>6</sup> So, to be able to work with the different quantum hardware, it is necessary to be knowledgeable about the requirements and libraries of each one of them.

## WHERE DO WE GO FROM HERE?

Software and system technology innovation will further evolve at a fast pace in fields such as nanotechnology, biotechnology, genomics, and quantum computing.<sup>3</sup> Today we can already use quantum computers and profit from their huge computation capacity to solve problems considered very difficult or unaffordable for “classic” computing. Quantum computing speeds up the process of solving algorithms that require massive parallel computations and so allows us to better simulate nature. All of this brings very new, disruptive, and potentially useful innovations.

Our focus here is on quantum software platforms to get started in industry-scale software engineering. Quantum hardware suppliers have provided software technologies for their respective computers and quantum effect simulators. Results look promising as there are several platforms available which allow a smooth learning curve.

The state of quantum technology is improving at an accelerating rate. To produce useful and trusted quantum software, applications still must solve relevant issues, such as the resolution of qubits and the control of their errors. The results of quantum machines create new types of errors, and we must learn to interpret the results. However, each vendor provides the results in different ways, which again leads to the need to rely on a particular vendor or to build a homogenized channel to consolidate the results.

The importance of professional software engineering for quantum computing has been neglected



so far.<sup>7-9</sup> New software engineering methods must be conceived based on experiences from software engineering for data science and machine learning.<sup>3,9</sup> They must be enhanced to manage specific quantum challenges, such as uncertainties, noise, and interpretation. Along those lines, development tools are not yet suitable from a business point of view. The resources resulting from the use of vendor software development kits are individual elements that are not yet incorporated into enterprise development resources.

Software technology and development methodologies need to advance to make these assets part of a complete quantum software project lifecycle. The increasing awareness of quantum computing applications demands the production of quality quantum software. Without proper software technology platforms and suitable software engineering methods, quantum software remains a mere research topic. Especially in trusted environments, such as medicine, and others where defects will have severe consequences, quantum software must prove the same high-quality standards that we demand from any other software.

**P**hysics Nobel laureate and quantum pioneer Niels Bohr once remarked: “Those who are not shocked when they first come across quantum theory cannot possibly have understood it.” His observation still applies today, especially in using quantum effects to actually produce software. There is still a way to go to deliver quality quantum applications. Yet now is the time to start. To scale from research to industry, quantum software must adopt sound software engineering methods for the development of quantum software—and enhance them as we once did when scaling agile development. Good enough may be sufficient for today, but it is certainly not for tomorrow. 🤖

## REFERENCES

1. J. D. Hidary, *Quantum Computing: An Applied Approach*. New York: Springer, 2019.
2. *Quantum Software Manifesto*. <https://www.qusoft.org/quantum-software-manifesto> (accessed June 20, 2021).
3. C. Ebert and B. Tavernier, “Technology trends: Strategies for the new normal,” *IEEE Softw.*, vol. 38, no. 2, pp. 7–14, Mar. 2021. doi: 10.1109/MS.2020.3043407.
4. European Quantum Flagship, “*Strategic research agenda*.” <https://digital-strategy.ec.europa.eu/en/library/quantum-flagship-major-boost-european-quantum-research> (accessed June 20, 2021).
5. “*Quantum devices and simulators*.” IBM. <https://www.research.ibm.com/ibm-q/technology/devices/#ibmq-20-tokyo> (accessed June 20, 2021).
6. “*Algebraic and number theoretic algorithms*.” Quantum Algorithm Zoo. <https://quantumalgorithmzoo.org/> (accessed June 5, 2021).
7. J. Zhao, “Quantum software engineering: Landscapes and horizons,” July 14, 2020, arXiv:2007.07047v1 [cs.SE].
8. M. Piattini, G. Peterssen, and R. Pérez-Castillo, “Quantum computing: A new software engineering golden age,” *ACM SIGSOFT Softw. Eng. Newslett.*, vol. 45, no. 3, pp. 12–14, June 2020. doi: 10.1145/3402127.3402131.
9. M. Piattini, M. Serrano, R. Pérez-Castillo, G. Peterssen, and J. L. Hevia, “Towards a quantum software engineering,” *IT Prof.*, vol. 23, no. 1, pp. 62–66, Jan.-Feb. 2021. doi: 10.1109/MITP.2020.3019522.



**JOSE LUIS HEVIA** is the quantum chief technology officer and the software architect and software solutions IT manager at Alhambra IT, Madrid, 28037, Spain. Further information about him can be found at <https://www.aquantum.es>. Contact him at [jluis.hevia@alhambrait.com](mailto:jluis.hevia@alhambrait.com).



**GUIDO PETERSSEN** is the quantum chief operating officer and the director of R&D and software solutions at Alhambra IT, Madrid, 28037, Spain. Further information about him can be found at <https://www.aquantum.es>. Contact him at [guido.peterssen@alhambrait.com](mailto:guido.peterssen@alhambrait.com).



**CHRISTOF EBERT** is the managing director of Vector Consulting Services, Stuttgart, 70499, Germany. He serves on the editorial board of *IEEE Software* and is a Senior Member of IEEE. Further information about him can be found at <https://twitter.com/christofebert>. Contact him at [christof.ebert@vector.com](mailto:christof.ebert@vector.com).



**MARIO PIATTINI** is the quantum chief research officer and leader of the Alarcos Research Group, University of Castilla-La Mancha, Ciudad Real, 13001, Spain. Further information about him can be found at <https://www.uclm.es>. Contact him at [mario.piattini@uclm.es](mailto:mario.piattini@uclm.es).

## DEPARTMENT: ANECDOTES

# Preserving the Past by Industry Participants

James W. Cortada , Charles Babbage Institute, University of Minnesota, USA

FOR over 40 years practitioners in the world of IT and historians of computing and information have come together in these *Annals* to share their understanding of computing's history. Participants have their own reasons for involvement with the *Annals* and have their own approaches to studying and reporting history. Historians typically have different research agendas than practitioners, often dealing with much broader concerns than the practitioner. While some practitioners may have explicit research agendas, many others may just enjoy reading about the history they lived through or are collecting that history in maybe somewhat accidental ways. For instance, many preserve records of their research and development projects, companies and jobs they have engaged with, gray literature (company published manuals and brochures) that came into their sphere, even parts of machines, printouts of software, coffee mugs, certificates, and other documents. They do this collecting either because they are naturally "pack rats" or because they have a feeling or belief that it is important for them to preserve such documents and artifacts. In over 40 years of working with such participants in computing, I have been amazed at how many store materials relevant to the history of computing, also how many do not yet know how important these are or how to ensure their survival after they have passed from the scene. It may be useful for such participants to hear my story.

I have been collecting paper and other records, largely about IBM, for nearly a half century. I would like to suggest that many other IT practitioners can do the same. I stumbled on a two-prong strategy that proved fun, relevant, and easy to execute. In the beginning it was about selfishly preserving documents relevant to my personal role at IBM, where I worked from 1974 to 2013. These included my initial job offer letter,

congratulatory notes on joining IBM, later others complimenting my work, award certificates and appraisal notes, among others. Some e-mails I thought historically important went into the collection too, most notably the ones that circulated in IBM the week of the 9/11 attacks. Priceless! Many, if not most, colleagues did the same.

We either tossed these documents into a cardboard box for years or meticulously put them in a file cabinet in folders. I began with the box, later on folders. Along the way I also accumulated material items, such as my 100 Percent Club pins, which indicated that I had achieved my sales targets—a very big deal in IBM, especially for those of us eager to move up the corporate career ladder. Mementos of visits to laboratories and factories also were tossed into a box, along with IBM logoed coffee mugs, pen, and pencil sets given out as sales awards, and a few binders with IBM training materials. Yes, postcards, key chains, and other credenza dust collectors made it into the boxes too, three by the time I retired.

As so many readers of the *Annals* have done, I accumulated books dealing with the history of my employer, which I read as they were being published. It was easy to do, because when I was coming up, the number of publications about the history of IBM and its industry were few, so always fun and novel to read. By the time I finished digesting the multiple technical histories of IBM written by Emerson Pugh and his colleagues in the 1980s and early 1990s, I had become sensitized to the kinds of materials not in the company's archives but that could be potentially useful to future historians. That is when I started to collect methodically, much the way a hobbyist/collector might art, old radios, or stamps. As I looked around, I saw collectables all over the place.

When IBM decided to no longer advertise its PCs with Charlie Chaplin cartoons, that day I noticed a PC poster with the classic tag line "A Tool for Modern Times" hanging on the wall in my building. I asked the local administrative manager if I could have it; he said yes, and today it hangs on a wall in my home. Good luck trying to find one anywhere; but I saved one. Since the 1950s the most famous poster in IBM was of a green duck entitled, "How to Stuff a Wild Duck," which referred

to IBM encouraging technical wild ducks to create new information technologies. I wanted one, and eventually two decades later, obtained one from an IBMer who no longer wanted his. Try and find that one; the last copy I saw was for sale on eBay for nearly \$1000.

In 1984 or 1985, the IBM sales office in Nashville, Tennessee, decided it was too cluttered and so held a contest to see who could throw out the greatest amount of old paper. I had the presence of mind to realize here was an opportunity to preserve old IBM publications (e.g., about S/360s and S/370s), photographs of the office from the 1950s onward, even a few certificates congratulating people on various achievements. That effort resulted in two boxes of materials and a “heads up” call to corporate archives that there was a nearly complete set of manuals about the iconic S/360 family of computers. The archivists moved swiftly to save those materials. I stored my 2 boxes of Nashville materials in the basement of our home.

Over the years, colleagues learned that I was interested in preserving IBM materials and so would send things to me. In the 1980s and 1990s, for example, cartoons and other humor related to IBM flowed my way, a box full to be precise. If someone came across a 50th anniversary history of IBM in Ireland or Norway, they remembered me and grabbed a copy for my collection. When I visited an IBM factory or laboratory, I would ask if they had marketing materials or employee publications that I could have for my collection; more stuff into a brief case. Such behavior became a habit, business as usual. Little did I know at the time how much that behavior was evident all over the company, especially among the ranks of American customer engineers (the folks who installed and repaired hardware), engineers and scientists, and French salesmen and systems engineers.

After I retired from IBM, all those various sources began to dry up, but I knew the day would come when I would write a history of IBM, which I published in 2019.<sup>1</sup> In 2013, I turned to eBay to see what it had and over the next seven years acquired about a banker box of materials per year just from that one source. I acquired manuals that described IBM’s key products from the 1950s forward, with particular emphasis on materials published between the 1930s and the end of the 1980s. I found about 100 postcards about NCR, where IBM’s long term CEO Thomas Watson Sr. matured into an executive, and dozens of others of IBM facilities from the 1920s through the 1990s. I found a 1928 graduation certificate from IBM’s Sales School (I graduated in 1975 and have mine too), letters signed by senior executives, congratulating employees about their achievements, a few ancient slide decks from the 1960s and 1970s, product repair manuals from the 1930s, even a welcome packet for a new IBM

employee circa 1962 that shed light on how seriously IBM took its legal responsibilities to adhere to its 1956 Consent Decree in the United States. Press photographs of Soviet premier Nikita Khrushchev’s visit to IBM’s manufacturing site in California in 1959 kept appearing for sale over the years, including a brochure Russian visitors received earlier that year when IBM exhibited a RAMAC system in Moscow at the same industrial fair where President Richard Nixon had his famous “kitchen debate” with the premier.

I noticed one day that a dealer of all manner of things had sold me two small collections of private records of deceased IBMers in California. So I reached out to her to find out how she obtained such materials. There were few records of product engineers so acquiring more of these was important to me. She explained that her role was to clean out homes of deceased people, others moving to assisted living situations, and those of hoarders. I explained that I was collecting with the intent that all my materials would be deposited at an archive, such as the Charles Babbage Institute at the University of Minnesota. She began to send me materials that would be difficult to sell on eBay, with the result that more came in: a woman’s ID badge from Poughkeepsie, New York, issued in 1944; a file containing a patent application and lab records on the design of the device in question; photographs of individuals in various offices, a secretary’s manual on how to use an IBM Selectric typewriter (considered probably the most popular typewriter in post World War II America), yet from another teaching materials from a typing class of the 1950s.

After publishing my history of IBM, a few people reached out to me, none so fascinating as the daughter of a retired South American IBM sales executive, Luis A. Lamassonne. She had “stuff” and wanted to know if I wanted these. He had self-published a memoir in English years earlier, which one could find through Internet book dealers.<sup>2</sup> But a Spanish edition had also been published which one could not find, but a member of the family had one. Lamassonne’s daughter also had a manual he published while working in Spain just before the start of the Spanish Civil War of 1936–1939, and was willing to part with that. In 1957 on the occasion of his 25th anniversary of working at IBM, his staff put together a book of letters and photographs, as was customary around the world and still practiced in the 21st century. IBMers value their binder, as do I mine. I had never seen one from Latin America, solid evidence that IBM’s corporate culture was practiced in Latin America. That was a big deal for an historian to uncover. I now have nearly a half a banker box of materials from this one family that sheds light on IBM’s role in South and Central

America, largely from the 1930s through the 1960s. This gentleman lived to over 100 years of age, IBM's oldest retiree at one point. It made my year.

A retired IBM typewriter repair customer engineer in his '80s in the United States began an e-mail dialog with me after reading a copy of my book given to him by his daughter as a Christmas present. The exchange of e-mails went on for several weeks in which he essentially wrote his memoirs of working at IBM. He now resides in one of my file cabinets with his own folder, accompanied by others from IBMers, and still others that included typewrite customer publications I acquired through eBay. Besides saving him from anonymity, we know too little about IBM's typewriter business, yet it was always a front and center concern to IBM's CEO for 40 years, Thomas Watson Sr.

Every week I spend 15 min to see what is for sale on eBay and by online second hand book dealers, acquiring odds and ends. From AbeBooks, for example, I acquired two copies of a hardbound celebratory volume published in Nazi Germany in the 1930s on the occasion of IBM opening a new manufacturing facility. It was strange seeing photographs with Nazi Party officials in their uniforms. From France, I found several IBM France annual reports from the 1960s that even IBM's corporate archives did not have. Three books written about IBM in Argentina showed up, so too seven published in Japan. About three dozen IBMers have self-published their memoirs, including two for consumption only by their families, one of which was written by a senior vice president that proved valuable when writing my book. I doubt any publication cost me more than \$20 each, from eBay between \$5 and \$20 each plus postage. So the costs of acquiring materials remained low. Once in a while on eBay I competed with an unknown buyer, but that was fine because I suspected that rival valued the material and would preserve it too. After all, the purpose of acquiring ephemera was to preserve them.

**What lessons can I draw** from these experiences? If we accept the idea that we have personal responsibilities for preserving the history of IT, our careers, our employers, and computing at large, it is not a difficult or expensive task. But it does require that you keep an eye out for materials and preserve them. It matters less whether you think something is important; let future historians figure that out. For example, when I obtained a 1944 employee ID badge, it taught me a lot. She had to pass a criminal-record check, it had to be issued by the police department (not IBM), and it included her Social Security number, birth date, and

photograph. It was as official as a driver's license. I now know why when I lived in Poughkeepsie I could cash a check by flashing my IBM ID card rather than a driver's license—the card proved I had a job, so could probably cover the check! More seriously, the card was further proof that what IBM was doing at its factory was defense-related and therefore subject to American military security processes.

Treat collecting and preserving like a hobby; keep it fun and interesting. Encourage colleagues to do the same and when their spouse declares that their garage or basement needs to be cleaned out, reach out to a university, library, archive, or fellow collector about how to dispose of the materials. Yes, even reach out to anybody on the board of editors of the *Annals*. When my basement fills, I will ship what I have to an archive; my family knows which ones to contact should I die or otherwise be incapacitated. You should establish the same understanding with your family. Too many materials end up being put out as trash. Do not assume that after you are gone that your family will have a better idea of what you did than when you worked.

On occasion, organizations like the Charles Babbage Institute and the Computer History Museum conduct research initiatives where they ask for materials dealing with a specific topic. They also interview people on their careers. Pay attention to those events—they are often reported on in the *Annals*, retiree websites and elsewhere. Volunteer to participate. Do not assume you have nothing of value; we still know so little about the history of IT that you cannot assume your materials and history are minor. You may be the least qualified person to judge your collection or career. From IBM retiree websites, I have collected several hundred obituaries and biographies. That means each of those individuals will be remembered, their stories will be the ones historians learn from. The ancient Greek philosopher Aristotle is remembered not only because he had something to say; but, because some of his writings were preserved. Scholars believe he associated with many other philosophers who have gone anonymous, hence forgotten, because nobody saved their records long enough. So, his is the experience we know. 🤔

## REFERENCES

1. J. W. Cortada, *IBM: The Rise and Fall and Reinvention of Global Icon*. Cambridge, MA, USA: MIT Press, 2019.
2. L. A. Lamassonne, *My Life With IBM*, Atlanta, GA, USA: Protea Publishing, 2000.

## DEPARTMENT: VISUAL COMPUTING: ORIGINS

# History of the Marching Cubes Algorithm

William E. Lorensen

### EDITOR'S NOTE:

The Marching Cubes paper by Bill Lorensen and Harvey Cline, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm," was published at SIGGRAPH 1987.<sup>1</sup> According to Google Scholar, their paper has 15,667 citations (as of January 17, 2020), the most highly cited paper in computer graphics. Sadly, while writing this article Bill Lorensen passed away on December 12, 2019. Origins Department Editor Chris Johnson contributed the text in italics.

### EARLY HISTORY OF MARCHING CUBES

In 1968, I graduated from Rensselaer Polytechnic Institute (RPI) with a B.S. in Mathematics. I intended to get a Ph.D. in Math and teach Mathematics. The Vietnam War changed that. I tried to find an employer who could get a job-related draft deferment. Both IBM and GE Corporate Research (CRD) gave me reasonable offers but could not get me a deferment. They did promise to hold my job if I was drafted. I had a low lottery number (108) and was sure to be drafted. The U.S. Army Maggs Research Center could get the deferment, so I took that job as a Scientific Programmer, even though the salary was less than the IBM and GE offers.

The threat of being drafted changed my career and directed me to the nascent field of computer graphics. At my Army job, I had access to an EAI analog plotter. An IBM Reproducing Punch Machine drove the plotter with a specially wired board that interpreted  $x$ ,  $y$ , pen code coordinates, one coordinate per punched card. I wrote programs to plot curves and 2-D models, punching one card per coordinate. The analog plotter tended to draw curved lines for long straight lines. Soon I was able to purchase a CalComp Incremental Plotter. The CalComp drew lines in small increments in one of eight

directions. I was able to draw more complex representations. Finally, I purchased a Lundy Vector Refresh System, and my work moved to three dimensions. I wrote my first 2-D contouring program for the Lundy. The algorithm started with a seed point and tracked that point as it moved through the 2-D scalar field.

My work for the Army involved finite element analysis pre- and postprocessing, developing grid generators and contouring algorithms. Graphics was just a means to an end. The finite element and graphics expertise landed me a job at GE Corporate Research and Development in Niskayuna, NY, USA. GE hired me to work in the computer service group that provided programming to GE scientists. I hit the road running with the Solid Mechanics Group and became an active contributor to their finite element research.

In 1984, Carl Crawford was working in GE's Medical Systems Business Group in Milwaukee. Carl worked in the Applied Sciences Lab and was a recognized expert in Computed Tomography (CT) reconstruction. Carl was also part of a company task force that was looking

---

Digital Object Identifier 10.1109/MCG.2020.2971284  
Date of current version 28 February 2020.

for applications for an exciting new GE product, the Graphicon. The GE Graphicon was a high-performance rendering engine created at GE's Flight Simulator Group. Each GE division was asked to look for applications of this technology. Carl gave a seminar in Building 37, at the downtown Schenectady GE manufacturing facility. Several of the research labs that did advance engineering work were in Building 37.

The main research campus where I worked was about 4 miles away.

Carl described the capabilities of the (not yet built) Graphicon. He described his vision for using this polygon engine for 3-D medical surface display. The current GE 3-D medical product was based on cuberilles and was optimized for implementation in the limited-memory Data General computers (only 32k memory). GE licensed the cuberille technology from the University of Pennsylvania. Carl challenged the seminar attendees to think about how they might replace the cuberille technology with polygon-based technology.


Harvey Cline and I attended the seminar. Harvey was in the Electronic Materials Lab, and I was in Computer Systems and Services, part of the central computer service group at CRD. Harv and I had been doing some work on reconstructing 3-D surfaces from interferograms. The application was the reconstruction of microscopic surfaces of electronic materials. The final step of our reconstruction algorithm used Movie.BYU's Mosaic program to generate triangles between adjacent contours. We had been looking for medical applications and saw Carl's challenge as a way to get our foot into GE Medical Systems. Carl's talk was in the morning, and Harv and I returned to my office to brainstorm. Rubick's Cube was the rage at the time, and Harv was analyzing the Cube problem using the symmetry of cubical lattices. The accepted way of generating surfaces from volume data was to create 2-D contours on each slice and then try to connect the contours with triangles. This is exactly what Mosaic did. Whenever Harv and I got together, it was an explosion of ideas. It is difficult to say who originated or how we came up with the idea. But somehow, we determined that solving the problem one cube at a time would remove a lot of the complexity. The inside/outside concept of labeling of vertices may have come from the Movie.BYU Mosaic

program. We quickly moved to the notion of solving the volume-to-surface problem one cube at a time. Harv and I started to solve the problem for each of the 256 cases. After just a few, we realized the effort was fruitless. But, using cube symmetry operations, we reduced the number of unique cases to 14. Harv and I worked out the triangulations for the basic cases (see Figure 1). I started to write code to permute the base cases into the 256 cases.

Then, I began coding the algorithm. The algorithm was straightforward to implement. I had something running the next day. The first implementation read volumes in as ASCII files. Harv put together a simple volume of a few slices, so I had some data to test the code. I used the Movie.BYU display program to produce hidden line and shaded surface models. Movie.BYU could only render up to 8192 polygons. The original Marching Cubes implementation created polygons, although I later switched to generating only triangles. Once the code was debugged, we tried the code on some medical data. We obtained a small CT dataset that included a spine. We extracted a small region of interest that isolated the spine. Since we were restricted to rendering 8192 triangles, the data had to be reasonably small. We were excited about the results.

## POST SIGGRAPH 1987

*Soon after the paper was published, it spawned new research in a number of important areas in computer graphics and visualization (e.g., computer-aided geometric design,<sup>6</sup> polygon simplification,<sup>9,27</sup> and volume rendering<sup>7</sup>). Because generation of surfaces and structures from images is important to many medical and biological applications,<sup>5,8</sup> Marching Cubes has been used in a wide array of applications including medical diagnosis, surgical planning, biomolecular graphics, and the creation of biomedical models for functional simulation. Similarly, understanding isosurfaces of three-dimensional scalar fields is essential for computational science and engineering applications such as computational fluid dynamics, weather simulation, computational mechanics, and computational geosciences. Marching Cubes is also used in computer games and virtual reality and most recently in a deep neural network based method, DeepOrganNet, to generate and visualize high-fidelity 3-D / 4-D organ geometric models from a single-view medical image in real time.<sup>20</sup>*




**FIGURE 1.** Marching Cubes case table.

In 1998, *Marching Cubes* was recognized as one of SIGGRAPH's seminal papers in *Computer Graphics* (Figure 3).<sup>3</sup>

**DECIMATION OF TRIANGLE MESHES AND MARCHING CUBES**

Jon Zarge was a member of the Software Technology Program (STP) at CRD. The STP recruited high


potential candidates with Bachelor's degrees and offered them an opportunity to get a Master's degree at a local university (usually Rensselaer Polytechnic Institute). While they worked on their degrees at Company expense, they rotated through two or three projects within CRD. Jon was assigned to our group to work on a new medical imaging modality, biomagnetism. One of the challenging problems of this project



**FIGURE 2.** Early visualization using Marching Cubes. Harvey Cline (Left), Bill Lorensen (Middle), Siegfried Ludke (Right), November 1988.

was to solve an inverse problem. From measurements obtained on the scalp of the patient, the inverse problem was to find the source of the biomagnetic field. Most researchers used idealized, spherical models of the skull and brain. The Biomag team set out to generate patient specific models. Their first choice for model generation was, of course, Marching Cubes. Marching Cubes posed two problems. First, the models generated by Marching Cubes were huge in polygon count. Second, as we found out, the models had topology problems. Jon's task was to generate models that could be analyzed. As I recall, we did not actually have the code to do the analysis.

Jon was being mentored by Will Schroeder (see Figure 6). Will laid out a high-level approach to solve the problem. The approach eliminated single vertices according to some out of plane criterion. Then, the neighborhood of the removed vertex was retriangulated. Will had a strong background in polygonal model topology. Step one of the process was to determine



**FIGURE 3.** Bill Lorensen and Harvey Cline on the occasion of the selection of Marching Cubes as one of SIGGRAPH's seminal papers in Computer Graphics.<sup>3</sup>

and encode the topology of the triangular surface. Jon wrote code to build a topological data structure to enable the vertex remove, hole-fill approach. This is when Jon brought the severity of the Marching Cubes topology problem\* to my attention. This had already been pointed out by Martin Durst's letter to the *ACM SIGGRAPH Quarterly*.<sup>4</sup> To be honest, I had tried to verify the extent of Durst's discovery but did not have the right tools. Jon's topology builder was just what I needed. While he worked on refining the initial decimation implementation, I worked on fixing the hole problem in Marching Cubes. The topology problem was easier to solve than I had anticipated. The root of the problem was my early assumption that all of the complementary cases, cases where inside/outside was reversed, were the same. Once I realized that the complementary cases needed different treatment, I

\*MC can produce a topologically inconsistent isosurface that contains holes caused by facetization ambiguity. Ambiguity analysis and disambiguation methods have been considered extensively.<sup>19</sup>




was able to resolve the ambiguity. Fortunately, my early code to generate the cases had separate inputs for the complementary cases. I generated a new case table, tested some generated models against Jon's topology checker and felt confident I had solved the problem.

Jon's implementation was in Lisp. I recall that he called the algorithm *intelligent triangle decimation*. At the time, our group was using LYMB as its development environment and delivery platform. Like many large systems, LYMB had a bit of a learning curve. The Biomag group needed a quick implementation of a decimation algorithm and Jon was not a LYMB expert. We decided to let him work out the details of decimation outside of the LYMB system. Will reimplemented Jon's stand-alone algorithm in LYMB. In the process of defining the object-oriented design, Will was able to generalize the decimation algorithm. Generalization of algorithms is a common theme followed by our group over the years.

I was a bit on the sidelines of the decimation algorithm development, concentrating on generating valid surfaces from Marching Cubes. I realized though, that this was a SIGGRAPH quality algorithm. I joined Will and Jon to start writing a description of the algorithm. I searched the literature for surface reduction techniques. I was surprised to find that there was nothing published on this topic. I also took on the task of generating examples to illustrate the power of the algorithm. Successful SIGGRAPH papers have two major components: innovative algorithms and compelling examples. I chose examples from medical imaging, industrial inspection, and terrain modeling.

SIGGRAPH 1992 was held in Chicago. Although we could not find other published work on polygon decimation, there were three papers accepted that year that dealt with reducing polygon count. Hughes Hoppe from the University of Washington and Greg Turk from Georgia Tech had the other two papers in the session. Will gave the talk to a packed room. I recall the talk was not in the main meeting room. SIGGRAPH had parallel sessions.



**FIGURE 4.** Admission tickets for the Marching Cubes Patent Wake held at *The Local Irish Pub*, Minneapolis, MN, USA, on October 25, 2005.




**FIGURE 5.** Bill Lorensen in 2019.

The decimation algorithm had a large impact on my career. The SIGGRAPH paper<sup>4</sup> is highly cited, but more than that, decimation made Marching Cubes a practical algorithm. It became a cornerstone of our model creation pipeline.

*Because isosurface extraction from 3-D volume data is so pervasive and important, there has been significant research in speeding up Marching Cubes and creating new isosurface algorithms.*<sup>9–18,27</sup> As a personal aside, in 1994, Han-Wei Shen and I wrote a paper on isosurface extraction for unstructured grids.



**FIGURE 6.** GE Research Visualization Team in 1992. Counterclockwise: Cathy Chalek, Chris Volpe, Will Schroeder, Bill Lorensen, Boris Yamron.



**FIGURE 7.** Bill Lorensen receiving the Best Software Engineering Paper from the 5th International Conference on Software Engineering in 1984.

We spent a significant amount of time trying to come up with a catchy name for the algorithm to mimic *Marching Cubes* and finally settled on *Sweeping Simplices*.<sup>12</sup> Needless to say, the name did not catch on. In 2006, Newman and Yi wrote a survey paper on the *Marching Cubes* algorithm.<sup>20</sup> The paper's aim was

to survey the development of the algorithm and its computational properties, extensions, and limitations (including the attempts to resolve its limitations). A rich body of publications related to *Marching Cubes* was included in that paper.

One day I received an envelope in the mail from Bill. When preparing to move from New York to California, Bill found the original reviews of the *Marching Cubes* paper and mailed them to me with a note saying that he thought that I would find the reviews interesting and especially a comment made by reviewer number three who commented that he thought the name *Marching Cubes* was “misleading and unhelpful.” A copy of

the original reviews of the *Marching Cubes* paper is available here.<sup>2</sup>

Given *Marching Cubes*' impact, especially in medical imaging, Bill wondered why GE did not aggressively leverage their patent on the *Marching Cubes* algorithm, but they did not. In 2005, Bill hosted a party to celebrate the expiration of the GE patent. The tickets to the party were, not surprisingly, from the *Marching Cubes* case table (see Figure 4).

According to Microsoft Academic, between 1987 and 1990, 100 papers cited the *Marching Cubes* paper. This climbed to 472 citations in the next five years and then to thousands over subsequent five-year intervals. In 2019 alone, there were 337 new citations. *Marching Cubes* lives on with large-scale parallel implementations via *VTK-m*<sup>21</sup> and *Flying Edges*,<sup>18</sup> and discrete *Marching Cubes* for segmentation masks. The algorithm is used in a variety of systems for performing visualization, modeling, and meshing such as *R*,<sup>22</sup> *Unity*,<sup>23</sup> edge group analysis,<sup>24</sup> and even weather forecasting.<sup>25</sup> Chances are many of us have directly or indirectly been impacted by the *Marching Cubes* algorithm, and it is likely that it will remain a staple of computing for the foreseeable future. Until his death, Bill maintained a *Marching Cubes* wiki.<sup>26</sup> The website is now being maintained by Will Schroeder and includes a tribute section to Bill. 🙏

## ACKNOWLEDGMENTS

The Origins Department Editors would like to thank Will Schroeder and Chuck Hansen for their help with this article.

## REFERENCES

- W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3D surface construction algorithm," *ACM Comput. Graph.*, vol. 21, no. 4, pp. 163–169, 1987. [Online]. Available: <https://dl.acm.org/doi/pdf/10.1145/37402.37422>
- "Reviews of the marching cubes paper," 1987. [Online]. Available: [http://www.sci.utah.edu/Marching\\_Cubes\\_Reviews\\_1987.pdf](http://www.sci.utah.edu/Marching_Cubes_Reviews_1987.pdf)
- R. J. Wolfe, *Seminal Graphics: Pioneering Efforts That Shaped the Field*. Los Angeles, CA, USA: ACM SIG-GRAPH, 1998.
- M. J. Dürst, "Re: Additional reference to "marching cubes," in *Proc. Int. Conf. Comput. Graph. Interactive Techn.*, vol. 22, no. 5, 1988, Art. no. 243.
- H. E. Cline, W. E. Lorensen, S. Ludke, C. R. Crawford, and B. C. Teeter, "Two algorithms for the three-dimensional reconstruction of tomograms," *Med. Phys.*, vol. 15, no. 3, pp. 320–27, 1988.
- J. Bloomenthal, "Polygonization of implicit surfaces," *Comput. Aided Geometric Des.*, vol. 5, no. 4, pp. 341–355, 1988.
- M. Levoy, "Display of surfaces from volume data," *IEEE Comput. Graph. Appl.*, vol. 8, no. 3, pp. 29–37, May 1988.
- B. A. Payne and A. W. Toga, "Surface mapping brain function on 3D models," *IEEE Comput. Graph. Appl.*, vol. 10, no. 5, pp. 33–41, Sep. 1990.
- W. J. Schroeder, J. A. Zarge, and W. E. Lorensen, "Decimation of triangle meshes," in *Proc. 19th Annu. Conf. Comput. Graph. Interactive Techn.*, vol. 26, pp. 65–70, 1992.
- J. Wilhelms and A. VanGelder, "Octrees for faster isosurface generation," *ACM Trans. Graph.*, vol. 11, no. 3, pp. 201–227, Jul. 1992.
- T. Itoh and K. Koyamada, "Isosurface generation by using extrema graphs," in *Proc. IEEE Vis.*, 1994, pp. 77–83.
- H. Shen and C.R. Johnson, "Sweeping simplices: A fast iso-surface extraction algorithm for unstructured grids," in *Proc. IEEE Vis.*, 1995, pp. 143–150.
- Y. Livnat, H. W. Shen, and C. R. Johnson, "A near optimal isosurface extraction algorithm using the span space," *IEEE Trans. Vis. Comput. Graphics*, vol. 2, no. 1, pp. 73–84. Mar 1996.
- P. Cignoni, C. Montani, E. Puppo, and R. Scopigno, "Optimal isosurface extraction from irregular volume data," in *Proc. IEEE Symp. Vol. Vis.*, 1996, pp. 31–38.
- C. L. Bajaj, V. Pascucci, and D. R. Schikore, "The contour spectrum," in *Proc. IEEE Vis.*, 1997, pp. 167–173.
- S. G. Parker, P. Shirley, Y. Livnat, C. D. Hansen, and P.-P. Sloan, "Interactive ray tracing for isosurface extraction," in *Proc. IEEE Vis.*, Oct. 1998, pp. 233–238.
- Y. Livnat and X. Tricoche, "Interactive point based isosurface extraction," in *Proc. IEEE Vis.*, 2004, pp. 457–464.
- W. Schroeder, R. Maynard, and B. Geveci, "Flying Edges: A high-performance scalable isocontouring algorithm," in *IEEE Symp. Large Data Anal. Vis.*, 2015, pp. 33–40.
- T. S. Newman and Hong Yi, "A survey of the marching cubes algorithm," *Comput. Graph.*, vol. 30, no. 5, pp. 854–879, 2006.
- Y. Wang, et al., "DeepOrganNet: On-the-Fly reconstruction and visualization of 3D / 4D lung models from single-view projections by deep deformation network," *IEEE Trans. Vis. Comput. Graphics*, vol. 26, no. 1, pp. 960–970, Jan. 2020.
- K. Moreland, et al., "VTK-m: Accelerating the visualization toolkit for massively threaded architectures," *IEEE Comput. Graph. Appl.*, vol. 36, no. 3, pp. 48–58, May/ Jun. 2016.
- D. Feng and L. Tierney, "Computing and displaying isosurfaces in R," *J. Statist. Softw.*, vol. 28, 2008, pp. 1–24, doi: 10.18637/jss.v028.i01.
- Unity, "High speed CPU-based marching cubes," 2020. [Online]. Available: <https://assetstore.unity.com/packages/tools/modeling/high-speed-cpu-based-marching-cubes-117483>
- C. A. Dietrich, C. A. C. Scheidegger, J. L. D. Comba, L. Nedel, and C. T. Silva, "Edge groups: An approach to understanding the mesh quality of marching methods," *IEEE Trans. Vis. Comput. Graphics*, vol. 14, no. 6, pp. 1651–66, Nov./Dec. 2008.
- H. Skotnes, G. Hartvigsen, and D. Johansen, "3D visualization of weather forecasts and topography," [Online]. Available: <https://munin.uit.no/handle/10037/401>, 1994.
- B. Lorensen, "Marching cubes Wiki," 2020. [Online]. Available: [marchingcubes.org](http://marchingcubes.org)
- J. Wilhelms and A. Van Gelder, "Octrees for faster isosurface generation," *Comput. Graph.*, vol. 24, no. 5, pp. 57–62, Nov. 1990.

Contact department editors Chris Johnson at [crj@sci.utah.edu](mailto:crj@sci.utah.edu), Dave Kasik at [dave.kasik@gmail.com](mailto:dave.kasik@gmail.com), or Mary C. Whitton at [mcwhitton@gmail.com](mailto:mcwhitton@gmail.com).



# 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](mailto:conferences@computer.org).

## APRIL

### 2 April

- HPCA (IEEE Int'l Symposium on High-Performance Computer Architecture), Seoul, South Korea

### 4 April

- ICST (IEEE Int'l Conf. on Software Testing, Verification and Validation), virtual

### 11 April

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

### 25 April

- VTS (IEEE VLSI Test Symposium), San Diego, USA

## MAY

### 4 May

- ICCPS (ACM/IEEE Int'l Conf. on Cyber-Physical Systems), Milano, Italy
- RTAS (IEEE Real-Time and Embedded Technology and Applications Symposium), Milano, Italy

### 9 May

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

### 15 May

- FCCM (IEEE Int'l Symposium on Field-Programmable Custom Computing Machines), New York, USA

### 16 May

- IC FEC (IEEE Int'l Conf. on Fog and Edge Computing), Messina, Italy

### 17 May

- ISORC (Int'l Symposium On Real-Time Distributed Computing), Västerås, Sweden

### 18 May

- ISCV (Int'l Conf. on Intelligent Systems and Computer Vision), Fez, Morocco
- ISMVL (IEEE Int'l Symposium on Multiple-Valued Logic), Dallas, USA

### 19 May

- AQTR (Int'l Conf. on Automation, Quality and Testing, Robotics), Cluj-Napoca, Romania
- SELSE (IEEE Workshop on Silicon Errors in Logic – System Effects), virtual

### 21 May

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

### 22 May

- ISPASS (IEEE Int'l Symposium on Performance Analysis of Systems and Software), Singapore
- SP (IEEE Symposium on Security and Privacy), San Francisco, USA

### 23 May

- ETS (IEEE European Test Symposium), Barcelona, Spain
- SEAMS (Int'l Symposium on Software Eng. for Adaptive and Self-Managing Systems), Pittsburgh, USA

### 25 May

- SERA (IEEE/ACIS Int'l Conf. on Software Eng., Management and Applications), Las Vegas, USA

### 30 May

- DCOSS (Int'l Conf. on Distributed Computing in Sensor Systems), Los Angeles, USA
- IPDPS (IEEE Int'l Parallel & Distributed Processing Symposium), Lyon, France

## JUNE

### 6 June

- EuroS&P (IEEE European Symposium on Security and Privacy), Genoa, Italy
- MDM (IEEE Int'l Conf. on Mobile Data Management), Paphos, Cyprus

### 11 June

- ISCA (ACM/IEEE Int'l Symposium on Computer Architecture), New York, USA

### 14 June

- WoWMoM (IEEE Int'l Symposium on a World of Wireless,



Mobile and Multimedia Networks), Belfast, UK

#### 19 June

- CVPR (IEEE/CVF Conf. on Computer Vision and Pattern Recognition), New Orleans, USA

#### 25 June

- CSCLOUD (IEEE Int'l Conf. on Cyber Security and Cloud Computing), Xi'an, China

#### 26 June

- ICIS (IEEE/ACIS Int'l Conf. on Computer and Information Science), Zhuhai, China

#### 27 June

- COMPSAC (IEEE Computers, Software, and Applications Conf.), Torino, Italy
- DSN (IEEE/IFIP Int'l Conf. on Dependable Systems and Networks), Baltimore, USA
- HOST (IEEE Int'l Symposium on Hardware Oriented Security and Trust), McLean, Virginia, USA

---

## JULY

#### 1 July

- ICALT (IEEE Int'l Conf. on Advanced Learning Technologies), Bucharest, Romania

#### 6 July

- ISVLSI (IEEE Computer Society Symposium on VLSI), Nicosia, Cyprus

#### 10 July

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

#### 11 July

- ICME (IEEE Int'l Conf. on Multimedia and Expo), Taipei, Taiwan

#### 21 July

- CBMS (IEEE Int'l Symposium on Computer-Based Medical Systems), Shenzhen, China

---

## AUGUST

#### 1 August

- ICCP (Int'l Conf. on Computational Photography), Pasadena, USA

#### 2 August

- MIPR (IEEE Int'l Conf. on Multimedia Information Processing and Retrieval), virtual

#### 4 August

- BCD (IEEE/ACIS Int'l Conf. on Big Data, Cloud Computing, and Data Science Eng.), Danang, Vietnam

#### 7 August

- CSF (IEEE Computer Security Foundations Symposium), Haifa, Israel

#### 9 August

- IRI (IEEE Int'l Conf. on Information Reuse and Integration for Data Science), virtual

#### 15 August

- RE (IEEE Int'l Requirements Eng. Conf.), Melbourne, Australia

---

## SEPTEMBER

#### 6 September

- CLUSTER (IEEE Int'l Conf. on Cluster Computing), Heidelberg, Germany

#### 12 September

- ARITH (IEEE Symposium on Computer Arithmetic), virtual

#### 18 September

- QCE (IEEE Quantum Week), Broomfield, Colorado, USA

#### 19 September

- AI4I (Int'l Conf. on Artificial Intelligence for Industries), Laguna Hills, USA
- AIKE (IEEE Int'l Conf. on Artificial Intelligence and Knowledge Eng.), Laguna Hills, USA
- ESEM (ACM/IEEE Int'l Symposium on Empirical Software Eng. and Measurement), Helsinki, Finland
- TransAI (Int'l Conf. on Transdisciplinary AI), Laguna Hills, USA

#### 26 September

- ASE (IEEE/ACM Int'l Conf. on Automated Software Eng.), Ann Arbor, USA



Learn more  
about IEEE  
Computer Society  
conferences

[computer.org/conferences](https://computer.org/conferences)

# Evolving Career Opportunities Need Your Skills

Explore new options—upload your resume today

[www.computer.org/jobs](http://www.computer.org/jobs)

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

Take advantage of these special resources for job seekers:



JOB ALERTS



TEMPLATES



WEBINARS



CAREER  
ADVICE



RESUMES VIEWED  
BY TOP EMPLOYERS

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



IEEE  
COMPUTER  
SOCIETY



IEEE