

06.20

Computer

CYBERTHREATS



IEEE



IEEE
COMPUTER
SOCIETY

vol. 53 no. 6

www.computer.org/computer



CALL FOR SPECIAL ISSUE PROPOSALS

Computer solicits special issue proposals from lead experts. 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 2021 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



Deadline for proposal submission: 1 July 2020



Computer



13

EIC'S MESSAGE Cyberpandemics

JEFFREY VOAS AND PHIL LAPLANTE

JUNE 2020

FEATURES

28

Reliability Inversion:
A Cautionary Tale

BEHROOZ PARHAMI

34

An Enterprise
Transformation Guide
for the Inevitable
Blockchain Disruption

MEHMET DEMIR,
OZGUR TURETKEN,
AND ATEFEH MASHATAN


44

Is Privacy
Regulation Slowing
Down Research
on Pervasive
Computing?

CLAUDIO BETTINI, SALIL KANHERE,
MARC LANGHEINRICH,
ARCHAN MISRA, AND
DELPHINE REINHARDT

ABOUT THIS ISSUE CYBERPANDEMICS

This issue features a virtual roundtable and three articles about cyberthreats.



Departments

4 Elsewhere in the CS

Membership News

Cover3 CS Connection

Cover3 IEEE Computer Society Information

COLUMNS

7 SPOTLIGHT ON TRANSACTIONS

Learn to Play Maximum
Revenue Auction

XIAOTIE DENG, TAO XIAO,
AND KEYU ZHU

9 50 & 25 YEARS AGO

ERICH NEUHOLD

11 COMPUTING THROUGH TIME

ERGUN AKLEMAN

16 VIRTUAL ROUNDTABLE Cyberthreats in 2025

JAMES BRET MICHAEL,
RICHARD KUHN, AND
JEFFREY VOAS

53 STANDARDS

IEEE SA Open: Engaging
Industry, Academia, and
Researchers in Open
Source Development

ROBBY ROBSON

57 CYBERTRUST

The Top 10 Risks of Machine
Learning Security

GARY MCGRAW, RICHIE BONETT,
VICTOR SHEPARDSON,
AND HAROLD FIGUEROA

62 THE IOT CONNECTION

Securing the Internet of
Things: An Ongoing Challenge

SEAN W. SMITH

67 REBOOTING COMPUTING

The Rise of the Quantum
Internet

MARCELLO CALEFFI,
DARYUS CHANDRA,
DANIELE CUOMO,
SHIMA HASSANPOUR,
AND ANGELA SARA CACCIAPUOTI

73 IT INNOVATION

Inference Acceleration:
Adding Brawn to the Brains

MARK CAMPBELL

77 OPEN SOURCE EXPANDED

Managing Your Open Source
Supply Chain—Why and How?

NIKOLAY HARUTYUNYAN

82 BODY OF KNOWLEDGE

An Ambassador for
Neural Networks

DAVID ALAN GRIER

Circulation: *Computer* (ISSN 0018-9162) 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. IEEE Computer Society membership includes a subscription to *Computer* magazine.

Postmaster: Send undelivered copies and address changes to *Computer*, IEEE Membership Processing Dept., 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Canadian GST #125634188. Canada Post Corporation (Canadian distribution) publications mail agreement number 40013885. Return undeliverable Canadian addresses to PO Box 122, Niagara Falls, ON L2E 6S8 Canada. Printed in USA.

EDITOR IN CHIEF

Jeffrey Voas
IEEE Fellow
j.voas@ieee.org

ASSOCIATE EDITOR IN CHIEF

Elisa Bertino
Purdue University
bertino@cs.purdue.edu
**ASSOCIATE EDITOR IN CHIEF,
COMPUTING PRACTICES**
Rohit Kapur
Synopsys
kapurfamily04@gmail.com

ASSOCIATE EDITOR IN CHIEF, PERSPECTIVES

Jean-Marc Jézéquel
University of Rennes
jean-marc.jezequel@irisa.fr
George K. Thiruvathukal
Loyola University Chicago
gkt@cs.luc.edu

2020 IEEE COMPUTER SOCIETY PRESIDENT

Leila De Floriani
University of Maryland,
College Park
deflo@umiacs.umd.edu

AREA EDITORS

CLOUD COMPUTING

Schahram Dustdar
TU Wien

COMPUTER ARCHITECTURES

David H. Albonesi
Cornell University
Erik DeBenedictis
Zettaflops, LLC

CYBER-PHYSICAL SYSTEMS

Oleg Sokolsky
University of Pennsylvania

CYBERSECURITY

Rick Kuhn
NIST

DIGITAL HEALTH

Christopher Nugent
Ulster University

EMBEDDED COMPUTING

Marilyn Wolf
University of Nebraska

HIGH-PERFORMANCE COMPUTING

Vladimir Getov
University of Westminster

INTERNET OF THINGS

Michael Beigl
Karlsruhe Institute of Technology

SECURITY AND PRIVACY

John Viega
Capsule8

SOCIAL-PHYSICAL-CYBER SYSTEMS

Mike Hinchey
University of Limerick

SOFTWARE ENGINEERING

Phil Laplante
Pennsylvania State University

VISION, VISUALIZATION, AND AUGMENTATION

Mike J. Daily
HRL Laboratories

COLUMN AND DEPARTMENT EDITORS

AFTERSHOCK

Hal Berghel
University of Nevada, Las Vegas

Robert N. Charette

ITABHI Corporation
John L. King

University of Michigan

BODY OF KNOWLEDGE

David Alan Grier
Djaghe, LLC

COMPUTER ECONOMICS

Nir Kshetri
University of North Carolina-
Greensboro

COMPUTING THROUGH TIME

Ergun Akleman
Texas A&M

CYBER-PHYSICAL SYSTEMS

Dimitrios Serpanos
University of Patras

CYBERTRUST

James Bret Michael
Naval Postgraduate School

EDUCATION

Irena Bojanova
NIST
Phil Laplante
Pennsylvania State University

THE IOT CONNECTION

Trevor Pering
Google

IT INNOVATION

Mark Campbell
Trace3

OPEN SOURCE EXPANDED

Dirk Riehle
University of Erlangen-Nuremberg

OUT OF BAND

Hal Berghel
University of Nevada, Las Vegas

REBOOTING COMPUTING

Erik DeBenedictis
Zettaflops, LLC

SPOTLIGHT ON TRANSACTIONS

Ron Vetter
University of North Carolina
Wilmington

STANDARDS

Forrest "Don" Wright
Standards Strategies, LLC

WEB EDITOR

Zeljko Obrenovic
Incision

50 & 25 YEARS AGO

Erich Neuhold
University of Vienna

ADVISORY PANEL

Doris L. Carver, Louisiana State University (EIC Emeritus)
Carl K. Chang, Iowa State University (EIC Emeritus)
Bob Colwell, Consultant
Sumi Helal, University of Florida (EIC Emeritus)
Bill Schilit, Google
Ron Vetter, University of North Carolina Wilmington (EIC Emeritus)
Alf Weaver, University of Virginia



CS PUBLICATIONS BOARD

Fabrizio Lombardi (VP of Publications), Cristiana Bolchini, Javier Bruguera, Carl K. Chang, Fred Douglass, Charles Hansen, Shi-Min Hu, Antonio Rubio, Diomidis Spinellis, Stefano Zanero, Daniel Zeng

MAGAZINE OPERATIONS COMMITTEE

Diomidis Spinellis (Chair), Lorena Barba, Irena Bojanova, Shu-Ching Chen, Gerardo Con Diaz, Lizy K. John, Marc Langheinrich, Torsten Moller, David Nicol, Ipek Ozkaya, George Pallis, VS Subrahmanian, Jeffrey Voas

COMPUTER STAFF

Senior Managing Editor

Geraldine Krolin-Taylor
g.krolin-taylor@ieee.org

Cover Design

Matthew Cooper

Peer Review Administrator

computer-ma@computer.org

Publications Portfolio Manager

Carrie Clark

Senior Advertising Coordinator

Debbie Sims

Publisher

Robin Baldwin

IEEE Computer Society Membership Director

Erik Berkowitz

IEEE Computer Society Executive Director

Melissa Russell

IEEE PUBLISHING OPERATIONS

Senior Director, Publishing Operations

Dawn M. Melley

Director, Editorial Services

Kevin Lisankie

Director, Production Services

Peter M. Tuohy

Associate Director, Information Conversion and Editorial Support

Neelam Khinvasara

Senior Art Director

Janet Dudar

Digital Object Identifier 10.1109/MC.2020.2984251

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 the IEEE by writing to the IEEE Intellectual Property Rights Office, 445 Hoes Lane, Piscataway, NJ 08854-4141 or pubs-permissions@ieee.org. Copyright © 2020 IEEE. All rights reserved. IEEE prohibits discrimination, harassment, and bullying. For more information, visit www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

ELSEWHERE IN THE CS

Computer Highlights Society Magazines

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.

Computing in SCIENCE & ENGINEERING

Jupyter Notebooks as Discovery Mechanisms for Open Science: Citation Practices in the Astronomy Community

Citing data and software is a means to give scholarly credit and facilitate access to research objects. Citation principles encourage authors to provide full descriptions of objects, with stable links, in their papers. As Jupyter notebooks (JNs) aggregate data, software, and other objects, they may facilitate or hinder citation, credit, and access to data and software. This article from the January/February 2020 issue of *Computing in Science & Engineering* reports on a study of references to JNs in astronomy over a five-year period (2014–2018). References increased rapidly, but fewer than half of the references led to JNs that could be located and opened. JNs appear better suited to supporting the research process than to providing access to research objects. The authors recommend that authors cite individual data and software objects and that they stabilize any notebooks cited in publications. Publishers should increase the number of citations allowed in papers and employ descriptive metadata-rich citation styles that facilitate credit and discovery.

Digital Object Identifier 10.1109/MC.2020.2976375
Date of current version: 4 June 2020

IEEE Annals of the History of Computing

High Noon on the Creative Frontier: Configuring Human and Machine Expertise

In 1960, CBS aired a special program, “The Thinking Machine,” which featured three Western playlets scripted by a computer programmed by MIT researchers. Nearly 60 years later, two researchers at Autodesk used a computer program to help design a chair. In this article from the October–December 2019 issue of *IEEE Annals of the History of Computing*, the author links these two seemingly discrete examples of computational creativity to highlight how digital fabrication technologies have served as an important test site for defining human and computational expertise. The author illustrates how concepts of “creativity” and “routine” were produced alongside the concepts of computational creativity during the development of digital fabrication. This dichotomy of “creative” and “routine” is not only used to determine the kinds of tasks that are appropriate for humans and computers to perform within the design and production process but is also used to render invisible the embodied craft knowledge required to substantiate these systems.

IEEE Computer Graphics AND APPLICATIONS

Aggregated Ensemble Views for Deep-Water Asteroid Impact Simulations

Simulation ensembles such as the ones simulating deep-water asteroid impacts have many facets. Their analysis, in terms of detecting spatiotemporal patterns, comparing multiple runs, and analyzing the influence of simulation parameters, requires aggregation at multiple levels. The authors of this article from the January/February 2020 issue of *IEEE Computer Graphics and Applications* propose respective visual encodings embedded in an interactive visual analysis tool.

Factual and Counterfactual Explanations for Black Box Decision Making

The rise of sophisticated machine learning models has brought accurate but obscure decision systems that hide their logic, thus undermining transparency, trust, and the adoption of artificial intelligence in socially sensitive and safety-critical contexts. The authors of this article from the November/December 2019 issue of *IEEE Intelligent Systems* introduce a local rule-based explanation method, providing faithful explanations of the decision made by a black box classifier on a specific instance. The proposed method first learns an interpretable, local classifier on a synthetic neighborhood of the instance under investigation, generated by a genetic algorithm. Then, it derives from the interpretable classifier an explanation consisting of a decision rule, which explains the factual reasons for the decision, and a set of counterfactuals, suggesting the changes in the instance features that would lead to a different outcome. Experimental results show that the proposed method outperforms existing approaches in terms of the quality of the explanations and of the accuracy in mimicking the black box.

Performance Analysis of Microservice Design Patterns

Microservice-based solutions are currently gaining momentum because they do not have the disadvantages of traditional monolithic architectures. Business interest in microservices is increasing because the microservice architecture brings a lightweight, independent, reuse-oriented, and fast service deployment approach that minimizes infrastructural risks. This approach is at an early stage of its development, and in view of this, it is important to understand the performance of its design patterns. In this article from the November/December 2019 issue of *IEEE Internet Computing*, the authors obtained performance results related to query response time, efficient hardware usage, hosting costs, and packet-loss rates for three microservice design patterns practiced in the software industry.

High-Quality Fault Resiliency in Fat Trees

Coupling regular topologies with optimized routing algorithms is key in pushing the performance of

interconnected networks of supercomputers. In this article from the January/February 2020 issue of *IEEE Micro*, the authors present Dmodc, a fast, deterministic routing algorithm for parallel generalized fat trees, which minimizes congestion risk even under massive network degradation caused by equipment failure. Dmodc computes forwarding tables using a closed-form arithmetic formula by relying on a fast preprocessing phase. This allows for the complete rerouting of networks with tens of thousands of nodes in under a second. In turn, this greatly helps centralized fabric management react to faults using high-quality routing tables and has no impact on running applications in current and future very large-scale, high-performance computing clusters.

Modification of Gradient Vector Flow Using Directional Contrast for Salient Object Detection

Scene analysis is a relevant research field because of its several applications in the area of computer vision. This article from the October–December 2019 issue of *IEEE MultiMedia* attempts to analyze scene information present in the image by augmenting salient object information with background information. The salient object is initially identified using a method called *minimum directional contrast* (MDC). The underlying assumption behind using this method for defining salient objects is that salient pixels have higher MDC than do nonsalient pixels. Finding MDC provides us with a raw salient metric. The gradient vector flow (GVF) model of image segmentation inculcates the raw saliency information. The gradient of MDC is calculated and added to the data term of the energy functional of GVF so that the contour formation utilizes not only edge formation but also saliency information. The result gives us the added background information as well as the salient object. Three public data sets are used to evaluate the results. The comparative study of the proposed method for salient object detection with other state-of-the-art methods available in the literature is presented in terms of precision, recall, and F1-score.

Design Differently: Pen and Paper for Laser Cutting

Interdisciplinary teams and studies need new approaches to design prototypes using tools that are indistinguishable from the ones they are used to. The authors of this article

ELSEWHERE IN THE CS

from the October–December 2019 issue of *IEEE Pervasive Computing* utilize a digital pen and physical paper to build a smart interface for laser cutters, giving nontechnical experienced people the possibility to rapidly, seamlessly, and collaboratively fabricate creative prototypes.

IEEE SECURITY & PRIVACY

Does Insurance Have a Future in Governing Cybersecurity?

Cyberinsurance could achieve public policy goals for cybersecurity using private-sector means. Insurers assess organizational security postures, prescribe security procedures and controls, and provide postincident services. The authors of this article from the January/February 2020 issue of *IEEE Security & Privacy* evaluate how such mechanisms impact security, identify the market dynamics restricting their effectiveness, and sketch out possible futures for cyberinsurance as governance.

THE SOFTWARE


Migrating a Software Factory to Design Thinking: Paying Attention to People and Mind-Sets

Design thinking (DT) has found its way into software engineering, promising better requirements elicitation,

customer relations, and cohesion within the development team. The authors of this article from the March/April 2020 issue of *IEEE Software* report on ProAction Technologies' migration toward DT and evaluate the process through interviews with employees and clients.

IT Professional

Sending More With Less: Crowdsourcing Integrated Transportation as a New Form of Citywide Passenger–Package Delivery System

Although much effort has been devoted by both academic and industrial communities to improve the efficiency of urban passenger and package flows, current urban transport systems still fail to balance speed and cost. To fill the gap, in this article from the January/February 2020 issue of *IT Professional*, the authors propose a novel form of transport system called *crowdsourcing integrated transportation*. It leverages the underused transport capacity, which is generated while delivering passengers to hitchhike packages so that more transportation needs can be met with fewer vehicles and drivers (that is, sending more with less). They identify the unique features of the new delivery system when compared to the traditional transport systems and discuss the key research challenges and potential solutions. They implement passenger-occupied taxis as the package carriers and evaluate their effectiveness. 

Editorial: Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *Computer* does not necessarily constitute endorsement by the IEEE or the IEEE 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 copyediting, proofreading, and formatting added by IEEE. For more information, please go to: http://www.ieee.org/publications_standards/publications/rights/paperversionpolicy.html. Permission to reprint/republish this material for commercial, advertising, or promotional purposes or for creating new collective works for resale or redistribution must be obtained from IEEE by writing to the IEEE Intellectual Property Rights Office, 445 Hoes Lane, Piscataway, NJ 08854-4141 or pubs-permissions@ieee.org. Copyright © 2020 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 is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.



Learn to Play Maximum Revenue Auction

Xiaotie Deng, Peking University

Tao Xiao, Shanghai Jiao Tong University

Keyu Zhu, Georgia Institute of Technology

This installment of Computer's series highlighting the work published in IEEE Computer Society Journals comes from IEEE Transactions on Cloud Computing.

Auctions have recently become a key protocol to allocate resources and determine prices for services provided over the Internet, in cloud computing, and for the Internet of Things. Machine learning, in another direction, has naturally been utilized to learn the underlying value distributions of customers for better mechanism designs.

Our work¹ is motivated by a real-world scenario in which the task is to allocate cloud computing resources (for example, the Amazon Elastic Compute Cloud), and it is imperative to determine how much to charge for each

service. Auctions are often repeated in cloud settings, and customers repeatedly play in the market. Learning their value distributions is expected in such a repeated game in which customers reveal more and more of their value distributions. From another perspective, Bayesian statistics heavily rely on a prior knowledge of data uncertainty to predict future events, and the optimization on decision variables affects future outcomes.

This creates a scenario in an auction in which the auctioneer tries to achieve future optimality based on its learned probability distributions for the values of buyers, such as the case of the Myerson's auction, commonly referred to as the maximum revenue auction.

In the era of big data and cloud computing, the optimum auction in the Bayesian setting would be a game of two parties through data to form a conceptual two-stage process. The auctioneer collects bidding data to learn the prior distribution of the value distribution of buyers. Based on the learned prior information, the auctioneer's optimum auction extracts the maximum optimal revenue from the participating buyers.

Digital Object Identifier 10.1109/MC.2020.2981988
Date of current version: 4 June 2020



FIGURE 1. The traditional versus practical auction scenarios.

In another direction, the buyers can report their deviated value distributions to acquire a better utility function value in the equilibrium of the game of the seller and buyers. More

As long as the auction and learning algorithm are specified, individual agents will respond to the auction protocol and auctioneer's learning strategy and submit their strategic bids.

specifically, buyers are assumed to have values over the resource, which follow certain probability distributions called the *prior*. Traditional Bayesian auction theory provides a good solution when the probability distributions are common knowledge.² In practice, however, this traditional theory faces a lot of challenges (Figure 1).


Statistical learning naturally becomes the most promising tool to learn the prior information, based on historical bids submitted by the customers. With the assumption that the bids are independent samples from the distributions, the optimal auction allows buyer value distributions to be learned. Hence, without specifying the auction, there will be no historical

data collected. As long as the auction and learning algorithm are specified, individual agents will respond to the auction protocol and auctioneer's learning strategy and submit their strategic bids, which may or may not be the true prior information. Our work takes this central issue into consideration.

As the first article in this regime, we consider the learning mechanism to be exactly Myerson's auction, the revenue-optimal auction mentioned at the beginning. We consider practical scenarios in which value distributions from agents can be represented by parameters. For the manipulation performed by individual agents,

we consider the simple and natural strategy space where the agents can manipulate over the parameters. Our work shows that it is possible to learn the revenue-optimal auction when individual priors are from power-law distribution family, while, for uniform and exponential distribution families, the learning task can be done when agents are from the same population.

The more general sponsored search auctions were studied subsequently under our model. Revenue equivalence results were developed among generalized first price (GFP), generalized second price (GSP), and Vickrey–Clarke–Groves auctions

in related hierarchical domains,³ which also provides a revenue justification for the switch from GSP to GFP⁴ for the repeated sponsored search auction. 

ACKNOWLEDGMENT

This research was supported by the National Natural Science Foundation of China under grant 61632017.

REFERENCES

1. X. Deng, T. Xiao, and K. Zhu, "Learn to play maximum revenue auction," *IEEE Trans. Cloud Comput.*, vol. 7, no. 4, pp. 1057–1067, 2019. doi: 10.1109/TCC.2017.2712142.
2. R. B. Myerson, "Optimal auction design," *Math. Oper. Res.*, vol. 6, no. 1, pp. 58–73, 1981.
3. X. Deng, T. Lin, and T. Xiao, "Private data manipulation in optimal sponsored search auction," in *Proc. Web Conf. 2020*, Apr. 2020, Taipei, Taiwan, pp. 2676–2682.
4. R. Shields, "After months of preparation, Google set to roll out first-price auctions," *Adweek*, Sept. 5, 2019. [Online]. Available: <https://www.adweek.com/programmatic/after-months-of-preparation-google-rolls-out-first-price-auctions/>

XIAOTIE DENG is a chair professor of computer science at Peking University, China. Contact him at xiaotie@pku.edu.cn.

TAO XIAO is a Ph.D. student at Shanghai Jiao Tong University, China.

KEYU ZHU is a Ph.D. student at Georgia Tech, Atlanta.

50 & 25 YEARS AGO



EDITOR ERICH NEUHOLD
University of Vienna
erich.neuhold@univie.ac.at



JUNE 1970

In the early years, *Computer* was only published bimonthly. Therefore, we will have to skip our interesting and/or informative extractions for June. The next one will appear in the an upcoming issue of *Computer*, and we hope you will eagerly wait for our next publication of this column.

JUNE 1995

www.computer.org/csdl/mags/co/1995/05/index.html

Binary Critic: Where the Big Money Is (p. 6) “The really big bucks are in consumer software products (about \$90 billion per year), not information processing (\$30–40 billion), ... Everyone knows about Bill Gates and the Dream Team, but the little guys in San Jose, Mountain View, Sunnyvale, and Palo Alto are also standing in line to autograph fat contracts with the “content kings.” (p. 7) “The trend is obvious: Mass media is sinking into oblivion, while products based on narrowcasting (for example, VCR tapes) are surging onto the market. ... The size of the VCR tape rental market far exceeds the box office theater market, and just about everything else. This “content” will quickly find its way onto long-playing CD-ROMs for your home computer.” [Editor’s note: *The move toward narrowcasting (personalization) has taken place, but TV stations, cable TV, PCs, and newer tablets and smartphones came and are still here. In that respect, the prediction toward PCs only has not been correct.*]

Computer Telephony Integration (p. 7) “Even though computer telephony integration (CTI) has been around since 1970, high hardware and software prices have limited it to specialized applications. However, as these prices fall dramatically, CTI is poised to explode into the mainstream. ... There are several key CTI applications that will start to generate interest in this trend, Glassman said. The first is screen-based telephony, which basically lets someone use a keyboard or click a mouse on an icon to have a computer dial a telephone number. The

second is call-based data selection, where an organization uses a phone company’s caller ID system to identify a caller. The phone company’s computer accesses a database and shows all relevant customer information before the call is answered.” [Editor’s note: *Despite some truth in this prediction, what was missed was the advent of the Internet and the rise of smartphones that brought multimedia communication to the end user, bypassing the CTI protocols envisioned in 1995.*]

Implications of Classical Scheduling Results for Real-Time Systems

(p. 16) “The scheduling theory literature is so vast that we can’t pretend to be comprehensive, but this article does present a minimum set of results and their implications. The set includes Jackson’s rule, Smith’s rule, McNaughton’s theorem, Liu and Layland’s rate-monotonic rule, Mok’s theorems, and Richard’s anomalies.” (p. 24) “Most multiprocessor scheduling problems are NP, but for deterministic scheduling this is not a major problem. We can use a polynomial algorithm and develop an optimal schedule if the specific problem is not NP-complete, or we can use off-line heuristic search techniques based on classical theory implications. These off-line techniques usually need to find only feasible schedules, not optimal ones. Many heuristics perform well in the average case and only deteriorate to exponential complexity in the worst (rare) case. Good design tools would allow users to provide feedback and redesign the task set to avoid the rare case.” [Editor’s note: *This article provides a very useful analysis of scheduling processes and specifies around 20 known theorems. Thus, it should help a designer to select the best-fit methods to solve a specific scheduling task. Some of these insights are certainly still useful today.*]

A Specification-Driven Architectural Design Environment

(p. 26) “In this article, we introduce an environment we’ve developed that helps designers represent, model, and explore trade-offs at the architectural abstraction level and synthesize designs at the behavioral level. The Design Analysis and Synthesis Environment (DASE) accomplishes this by supporting design capture, design space exploration, and

validation for the final design synthesis.” (p. 28) “DSL, the internal specification language in the Design Analysis and Synthesis Environment (DASE), is a Prolog metalanguage that is interpreted through a processor. Modular object-oriented design entities called modules are its primitive building blocks.” (p. 33) “Abstracting hierarchy can facilitate simulation and modeling in varying detail. For example, while a systems analyst might not focus on a model’s lower-level details, a hardware designer would be concerned with the simulation’s timing details. By defining an observation level, the DSL simulator enables dynamic alteration of the abstraction level viewed by the user. ... Module behavior can be translated to a netlist based on petri nets, which can help analyze hardware or software systems. The petri net formalism simplifies definition of asynchronous concurrent communications.” [Editor’s note: The article represents the methodology in detail and also claims, more briefly, that it has been used in the development of various forms of asynchronous transfer mode switches.]

Interrupt Processing in Concurrent Processors (p. 36) “To help designers systematically explore options for handling interrupts and help researchers compare interrupt processing strategies, we offer a taxonomy (or classification) of implementation choices. The approach we’ve developed broadly classifies interrupt-processing techniques and implementations into six phases.” (p. 38) “Interrupt-processing systems can be implemented in many ways. Some techniques, like history buffers, are general enough that the designer could make all processor interrupts precise, depending on the desired end result. With other strategies, such as adding special-purpose registers to a processor, the designer might make only one type of interrupt precise—for example, I/O interrupts.” (p. 42) “Interrupt Processing Phases: • Detect the interrupt. • Finish pending instructions. • Undo process state changes. • Save the process state. • Run the interrupt handler. • Resume the interrupted process.” (p. 45) “We can therefore conclude that, to be completely accurate, each possible processor interrupt (or type of interrupt) should be classified on an individual basis. As a whole, processors tend to resist being conveniently classified because of the kinds of special cases just described.” [Editor’s note: The article essentially analyzes the different kinds of interrupts and how to treat them. The six phases mentioned above are analyzed in detail with different alternatives of implementation explored. Unfortunately, the article concludes with a remark that the classification of existing implementations is not reasonable, as they contain too many different choices.]

Fault Injection (p. 47) “Fault injection is an effective solution to the problem of validating highly reliable computer systems. Tools such as React are facilitating its application. ... In critical applications, such as aircraft flight control, nuclear reactor monitoring, medical life support, business transaction processing, and telecommunications switching, computing resource failures can cost lives and/or money.”


(p. 48) “Taxonomy of Experiments: Fault injection experiments can be classified according to three general attributes: system abstraction, fault model and injection method, and dependability measure.” (p. 50) “Not until the mid-1980s did academia begin actively using fault injection to conduct experimental research. Initial work concentrated on understanding error propagation and analyzing the efficiency of new fault-detection mechanisms. Research has since expanded to include characterization of dependability at the system level and its relationship to work load.” (p. 53) “The studies discussed so far focused on validating existing systems, but fault injection can also evaluate the dependability of proposed designs. We have used simulated fault injection to analyze the reliability of several alternative TMR (Triple Modular Redundant) architectures.” [Editor’s note: The article analyzes a range of possible fault-injection mechanisms, investigates applications of those mechanisms, and then discusses various tools that have been used in those and other applications. Again, this is a survey article that has stayed interesting over the 25 years since its publication.]

International Survey: Virtual-Environment Research (p. 57) “A key development goal is to involve many—eventually all of our senses in acquiring information, although initially the focus has been on visual information display. Technologies for presenting information to other sensory modalities are being developed, such as acoustic displays that present 3D sound environments to a person’s sense of hearing and haptic displays that provide tactile and force feedback to a person’s sense of touch. Technologies for presenting information to other senses, such as smell and taste, are occasionally discussed but have rarely been incorporated into a VE.” (p. 65) “VIRTUAL ENVIRONMENT RESEARCH CLEARLY INVOLVES a broad range of technologies and potential applications with the intriguing potential to affect us in previously undreamed-of ways. Internationally, many researchers are independently conducting ongoing projects to develop these technologies and to incorporate them into effective VE systems. The recent emergence of commercial VE applications and the enterprises needed to support them indicates a maturing VE industry that is ready for expansion.” [Editor’s note: Interestingly, the techniques and applications discussed in this survey are not much different from what exists today. Of course, price and performance stood in the way of widespread utilization. As in other cases, the gaming industry finally led to mass-marketed availability of many of the discussed methodologies.]

Resources for Networks in Less-Industrialized Nations (p. 66) “True global connectivity is years away. Countries must first have an internal network, no matter how small, before linking electronically to the rest of the world’s networks. ... The good news is the Internet’s phenomenal growth. The bad news is that the growth is concentrated in the Northern Hemisphere: North America, Western Europe, and parts of Asia...”

COMPUTING THROUGH TIME

ERGUN AKLEMAN



DURING MEDIEVAL TIMES, PLAGUE DOCTORS WORE A BEAK-LIKE MASK DURING A PANDEMIC. THESE MASKS, FILLED WITH AROMATIC ITEMS SUCH AS HERBS, WERE SUPPOSED TO PROTECT THEM FROM PUTRID AIR, BASED ON MIASMATIC THEORY OF DISEASE. THIS THEORY HELD THAT THE ORIGIN OF EPIDEMICS WAS DUE TO A MIASMA (A NOXIOUS FORM OF BAD AIR), EMANATING FROM ROTTING ORGANIC MATTER.

Digital Object Identifier 10.1109/MC.2020.2984375
Date of current version: 4 June 2020

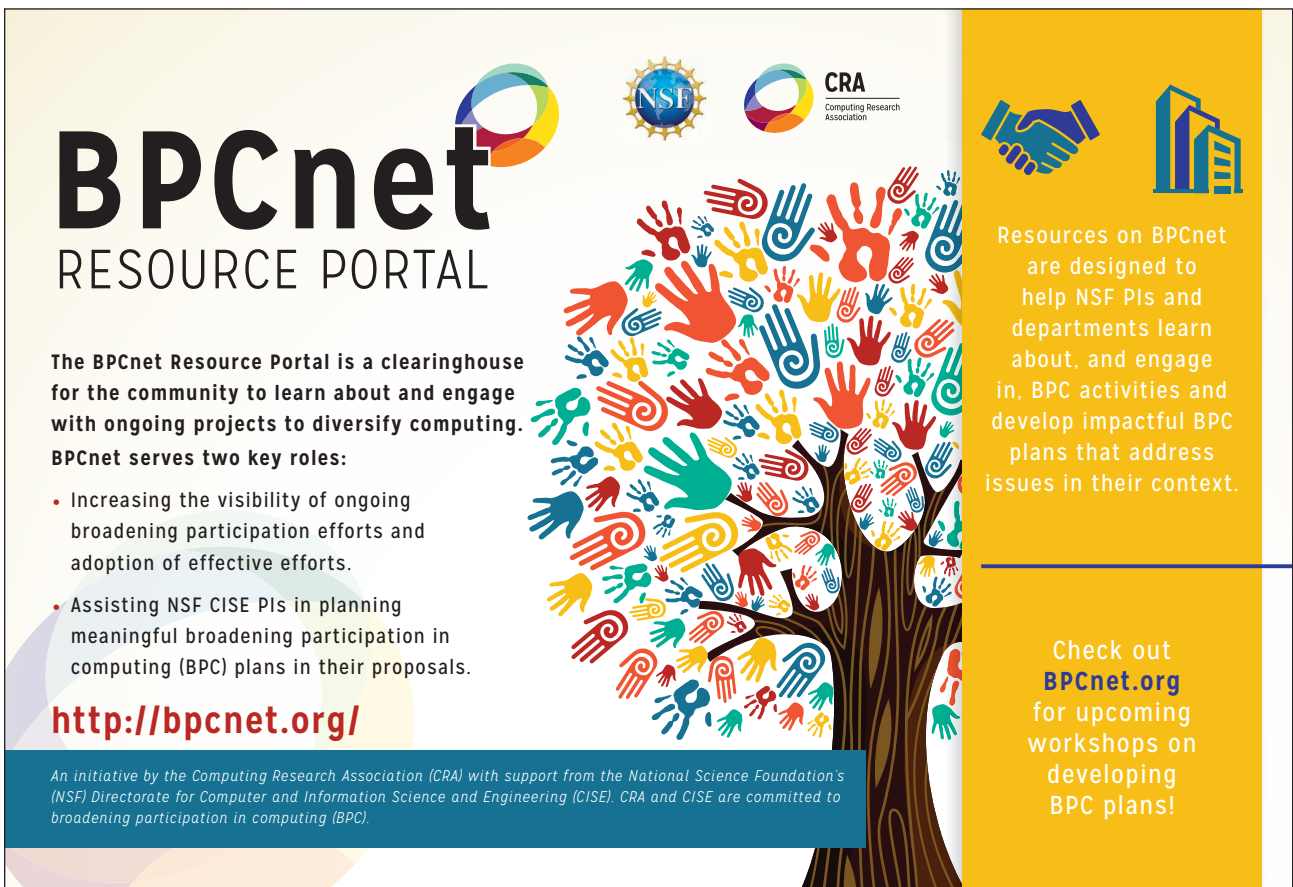
LINs are sparsely connected. Furthermore, connectivity is unevenly dispersed in industrial nations. For example, children in some Latin American schools have access to the Internet while those in US inner cities seldom do. ... Consider the evolution of the Relcom (Reliable Communications) network in the ex-Soviet Union." (p. 67) "While it has been successful, Relcom is unlike a more industrialized nation's national network. Where a typical US network uses large computers and workstations connected with dedicated high-speed communication links, Relcom uses mostly PCs and dialup telephone connections." (p. 70) "Eventually, outside experts must be supplanted by local expertise and self-sufficiency. When a network has been established, the need for people resources shifts from technical to managerial and financial skills. ... Pre-university education is an area of vast discrepancies. Ivan Illich and others have pointed out that education is often used to maintain the social status quo. Nations such as Cuba, Malaysia, and Costa Rica have ambitious computer literacy programs to reach all people - rich, poor, rural, and urban." [Editor's note: This interesting article shows us again what the situation was in 1995 and what tremendous progress has happened the world over. However, despite the fact that Wi-Fi was around

then, the article totally ignores the effect that mobile telephony has had in the development of the Internet for all.]

On-Line Monitoring: A Tutorial (p. 72) "On-line monitoring can complement formal techniques to increase application dependability. This tutorial outlines the concepts and identifies the activities that comprise event-based monitoring, describing several representative monitoring systems. ... Monitoring gathers information about a computational process as it executes and can be classified by its functionality. Dependability-Performance enhancement-Correctness checking-Security-Control-Debugging and testing-Performance evaluation. ...I focus on (first) four of the seven functional areas." (p. 77) "ON-LINE MONITORING IS INCREASINGLY SEEN AS A VIABLE means of increasing application dependability." [Editor's note: The tutorial discusses a number of monitoring systems available in 1995. They are used to increase system dependability. However, the same techniques are now used to monitor a vast number of applications, among them manufacturing, logistics, traffic control, people movement, crowd control, and so forth. Unfortunately, some of these can be used maliciously for the invasion of individuals as well as corporate privacy.]

Open Channel: The Virtual and the Paperless Office (p. 120) “Among the most promising developments of the telecommunications age is the virtual office. The potential benefits are remarkable, we hear: less traveling, better use of time, higher productivity, fewer dry cleaning bills. ... There is still one major factor, however, that could spoil the whole virtual office idea: the quality of those remote interpersonal exchanges. ... Team productivity today hinges on those inef-fable informal exchanges (sometimes known as chit-chat)

fostered by physical proximity. ... Will cyberspace ever replace such contact? ... There is another precedent here of a technology that just isn't good enough to be exploited in full. Think about WYSI-WYG, it's been around for years, yet how many of us still printout hard-copy drafts of our work, the better to review?” [Editor's note: The progress over the last 25 years has not eliminated physical office space and personal contacts (look at the amount of business travel). Also, the use of paper in office work is far from being eliminated. Just ask the paper industry.]



BPCnet
RESOURCE PORTAL

The BPCnet Resource Portal is a clearinghouse for the community to learn about and engage with ongoing projects to diversify computing. BPCnet serves two key roles:

- Increasing the visibility of ongoing broadening participation efforts and adoption of effective efforts.
- Assisting NSF CISE PIs in planning meaningful broadening participation in computing (BPC) plans in their proposals.

<http://bpcnet.org/>

An initiative by the Computing Research Association (CRA) with support from the National Science Foundation's (NSF) Directorate for Computer and Information Science and Engineering (CISE). CRA and CISE are committed to broadening participation in computing (BPC).

Resources on BPCnet are designed to help NSF PIs and departments learn about, and engage in, BPC activities and develop impactful BPC plans that address issues in their context.

Check out BPCnet.org for upcoming workshops on developing BPC plans!



Cyberpandemics

Jeffrey Voas, IEEE Fellow

Phil Laplante, Penn State

Cyberpandemics share similarities with biological pandemics, and we all have a role to play in preventing both.

Each year, Google announces what the most popular search term was from the previous year; *Disney Plus* was the top searched term on Google in the United States for 2019.¹ Given current events, we wonder if any of the following terms will become the most popular for 2020: *pandemics*, *coronavirus*, *COVID-19*? We ponder this because it appears from media reports that public acceptance of global pandemics may have become a new “normal” threat. Although global pandemics have been discussed on television and in movies for years, they have had somewhat of a science fiction angle to them even though they occurred often in past centuries.

More than 10 years ago, a few of us discussed out-of-the-box ideas, such as “what if the power grid went completely down?,” “what if the banks fail, and no one can access cash from an ATM?,” and “what if society completely or partially shuts down from a global cyber event?” In a 2009 article, we discussed various scenarios of the last idea.² The events of the coronavirus (CV) pandemic over the past few months have made us rethink that article. Note: we

are not sensationalizing the CV tragedy. Rather, we hope that revisiting the article might put the current pandemic into a perspective that can help us learn how to better deal with

future biological viruses and avoid cyberpandemics.

So, what is a *cyberpandemic*? Quite simply, it is a massive disruption of computing service that triggers second- and third-order failures of computing and noncomputing systems worldwide. Consequences could include widespread failure or malfunctioning of critical infrastructure systems and the associated major societal damage. Perpetrators of cyberpandemics could include rogue governments (or elements therein), terrorist groups, corporations and consortia (that may profit from the pandemic’s affects), malicious actors (of varied motivations), individuals, or even an accident (for example, a weaponized malware gone awry). As far as we know, a cyberpandemic has never been successfully perpetrated, but we all know that the weaponization of cyberspace is very real. A full potential of that is yet to occur.

The conditions that could lead to a cyberpandemic are similar or analogous to those for a biological pandemic: human complexity, attack multiplicity, and delayed effects. In the cyber world, human complexity is represented by people “packed too tightly in cyberspace” and the resultant complex social interactions online. Attack multiplicity means that the attack involved multiple

Digital Object Identifier 10.1109/MC.2020.2984253
Date of current version: 4 June 2020

IN THIS ISSUE

Parhami, the author of "Reliability Inversion: A Cautionary Tale," explains the notion of "reliability inversion." He claims this phenomenon occurs in practice for actual systems under realistic assumptions and points to certain system architectures that are more amenable to producing tight reliability bounds with tractable analytical models or simplified simulation-based models. An example involving centralized versus distributed reconfiguration switching in 2D processor arrays is used to support the ideas with quantitative results.

In "An Enterprise Transformation Guide for the Inevitable Blockchain Disruption," Demir et al. present a methodology termed the *blockchain technology transformation framework (BTTF)*. The article claims that this approach can inform decision makers on how blockchain fits in their processes, what data will be in their transactions, and who the participants will be. This framework builds a design map by which process

owners can analyze the suitability of blockchain technology. Through this approach, the authors believe that BTTF can provide organizations with a way to redesign their processes or identify opportunities for using smart contracts. Use case examples in supply chain and real estate are provided.

The last article in this issue is "Is Privacy Regulation Slowing Down Research on Pervasive Computing?" Bettini et al. present a study that investigated the impact of the recent evolution of personal data protection legislation on researchers in mobile and pervasive computing. The authors gathered feedback from more than 150 researchers in this field to better understand if this attitude is shared. Their findings indicate that most respondents do not feel there are any major impediments in adhering to privacy regulations, and they also found that the respondents were somewhat familiar with the latest legal developments and the majority seemed to be in favor of clear and strict privacy regulation.

Digital Object Identifier 10.1109/MC.2020.2988546
Date of current version: 4 June 2020

– Jeffrey Voas, Editor in Chief

simultaneous attacks and used more than one orthogonal (computer virus) vector mechanism. These may include one or more noncyber components in the attack. In the human analogue, this is somewhat equivalent to a virus that attacks the immune and nervous systems. In a cyberpandemic, attackers could launch the attack by using a noncyber component as a diversion, to soften the environment for the attack, or as the trigger mechanism, to signal the start of the attack or disrupt society's ability to recover. Social networks increase the likelihood of people falling into traps, can be used to create diversions for an attack (for example, flash mobs and riots), and can propagate or trigger malware. Delayed effects mean that symptoms emerge long after infection has occurred, making widespread dispersal likely and difficult to prevent.

In 2009, we introduced five potential scenarios for cyberpandemics; space in this editorial precludes reviewing them, but one such scenario (wag the dog) involved exploiting some type of natural disaster (such as a human pandemic) to distract attention from the impending threat as well as tire and weaken response agencies before the launch of a cyberpandemic. In fact, on 18 March 2020, a cyberattack was launched against the U.S. Department of Health and Human Services, apparently, to prevent the agency from responding to the CV. The origin of the attack is still unknown,³ and this attack did not reach the scale of a cyberpandemic. The point here is simple: a biological virus can greatly impact global economics and financials; however, an attack on the cyber infrastructure can also affect human health outcomes, for instance, holding

hospitals hostage via ransomware or completely shutting down the supply chains that medical professionals rely on.


As we write this editorial (in mid-March), the final toll of the CV is unknown, and we hope it will soon be conquered. But what is clear is that everyone must help fight this and future biological pandemics.

However, nations, organizations, and individuals also have a role to play in preventing cyberpandemics. The role of nations and groups of nations is clear. The role of individuals is also clear—as with biological pandemics, each person has a role to play in cyberspace by applying the principles of 1) least exposure, 2) defense in depth and separation of privileges, and 3) being aware of "dry runs" and probing activities and reporting them. Charitable, business, and professional

organizations also have a role, and we can marshal their forces to help.

We call upon members of the IEEE Computer Society to consider future research efforts on

- › telemedicine, the Internet of Things, and enabling technologies
- › big data, data analytics, and visualization
- › high-performance, cloud, fog, and edge computing (to support pandemic concerns)
- › human-computer interaction and the appropriate interfaces
- › social networking, communication protocols, and related psychosocial aspects
- › reengineering and rethinking education for virtual delivery (for example, clinicals, labs, and internships).

There are many others that we are missing in this short list. Not coincidentally, all of these are areas of interest to the IEEE Computer Society and *Computer*. Stay well. 

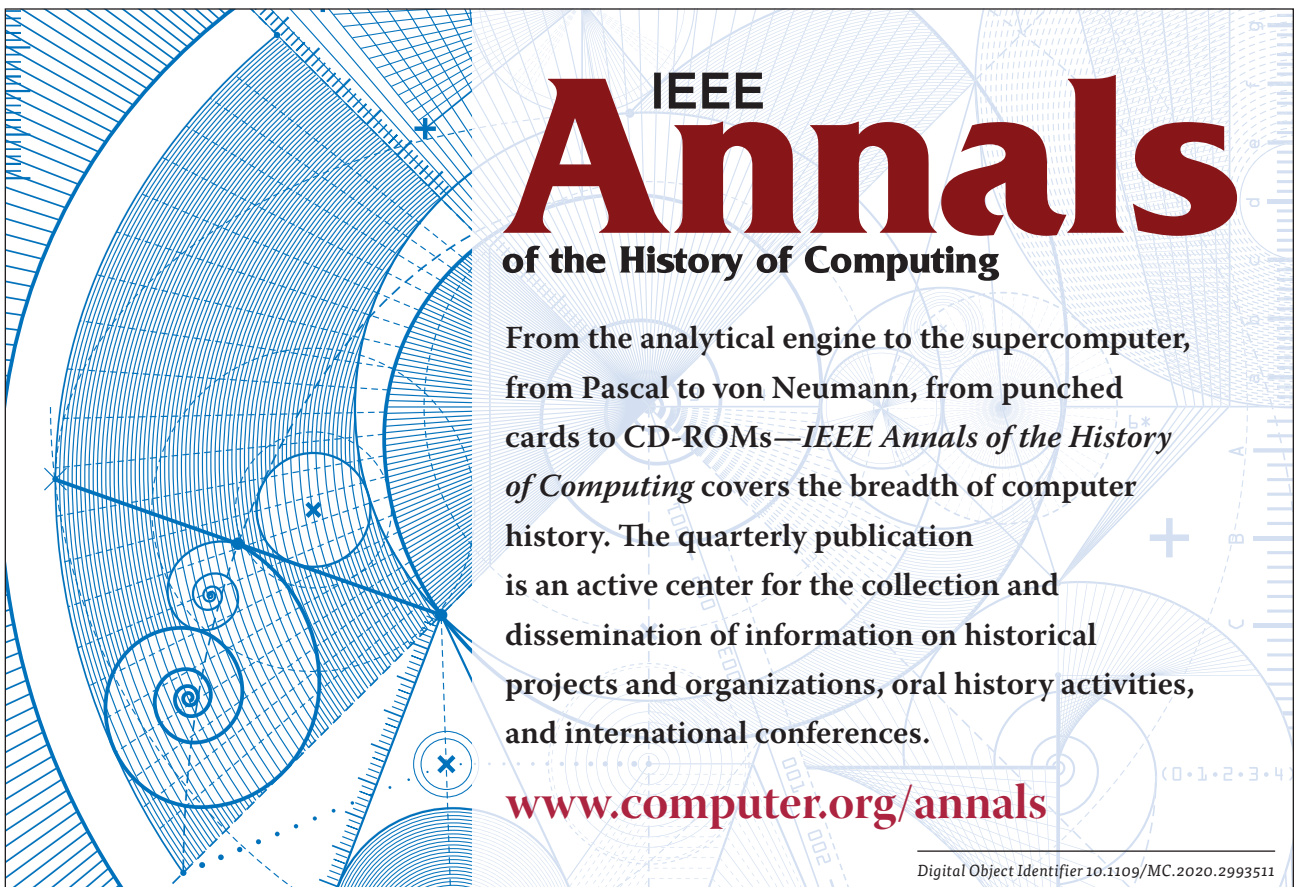
REFERENCES

1. J. Moreno, “These were the top Google searches and trends of 2019,” *Forbes*, Dec. 29, 2019. [Online]. Available: <https://www.forbes.com/sites/johanmoreno/2019/12/24/these-were-the-top-google-searches--and-trends-of-2019/#625c7ee43089>
2. P. Laplante, B. Michael, and J. Voas, “Cyberpandemics: History, inevitability, response,” *IEEE Security & Privacy*, vol. 7, no. 1, Jan./Feb. 2009, pp. 63–67. doi: 10.1109/MSP.2009.4.
3. H. Samsel, “Cyber attack hits Department of Health and Human Services amid government coronavirus

response,” *Security Today*, Mar. 18, 2020. [Online]. Available: <https://securitytoday.com/articles/2020/03/18/cyber-attack-hits-department-of-health-and-human-services-amid-government-coronavirus-response.aspx>

JEFFREY VOAS is the editor in chief of *Computer*. Contact him at j.voas@ieee.org.

PHIL LAPLANTE is a professor of software and systems engineering at Penn State and a member of the *Computer* editorial board. He is a Fellow of the IEEE. Contact him at plaplante@psu.edu.



The graphic features a complex background of blue technical drawings, including a large gear-like structure on the left, various circular patterns, and a grid of lines. The text is overlaid on the right side of the graphic.

IEEE
Annals
of the History of Computing

From the analytical engine to the supercomputer, from Pascal to von Neumann, from punched cards to CD-ROMs—*IEEE Annals of the History of Computing* covers the breadth of computer history. The quarterly publication is an active center for the collection and dissemination of information on historical projects and organizations, oral history activities, and international conferences.

www.computer.org/annals

Digital Object Identifier 10.1109/MC.2020.2993511



Cyberthreats in 2025

James Bret Michael, Naval Postgraduate School
Richard Kuhn, National Institute of Standards and Technology
Jeffrey Voas, IEEE Fellow

Computer hosts a virtual roundtable with six experts to discuss upcoming cyberthreats in 2025.

In *Computer*, virtual roundtables (VRTs) are virtual panels. We ask a series of questions to a group of experts via email to ascertain their thoughts on a topic du jour. One difference between VRTs and face-to-face panels is that no expert knows who the other experts are. That is different from an in-person panel, where answers from one panelist can affect the responses of others.

In this VRT, our topic of discussion is what the upcoming cyberthreats in 2025 might be. We could have asked about other years, such as 2030 or 2035. However, for this topic, the larger the number, the more the answers become sheer speculation. We believe that distant, futuristic speculation is of little value to the reader, given the relentless hacks that occur daily.

In this VRT, we invited six experts to respond to 12 questions. Their written responses may have undergone minor edits. However, as organizers, we attempted to keep

their words as verbatim as possible. The six experts are Jon Brickey (Mastercard), Simson Garfinkel (U.S. Census Bureau), Gary McGraw (Berryville Institute of Machine Learning), Latif Ladid (Université du Luxembourg), Bruce Potter (shmoo.com), and John Viega (Cap-

sule8). (See “Roundtable Panelists” for more information about the panel.)

It is important to note that the opinions of the experts are their own, with no input from the article editors. We hope readers who are concerned with cyberthreats and cybersecurity will find these questions and responses enlightening.

COMPUTER: What will be the prevalent types of cyberthreats in the year 2025?

JON BRICKEY: By 2025, threats will be more automated, more intelligent, more disruptive, and even destructive. Nation-state actors will push the envelope and use cyberattacks against critical infrastructure because they can’t achieve strategic effects through more traditional means. We’ll also see more of the same on the criminal actors side as they continue to take advantage of an explosion in consumer-focused fintech [financial technology] apps.

Digital Object Identifier 10.1109/MC.2020.2983529
 Date of current version: 4 June 2020

ROUNDTABLE PANELISTS

Jon Brickey is senior vice president, Cybersecurity Evangelist, for Mastercard Operations and Technology. In this role, he supports Corporate Security's mission of delivering safety and security at the speed of business. Before joining Mastercard, he served in the U.S. Army and retired as a colonel. Brickey received a Ph.D. in computer science and information systems from the University of Colorado Denver in 2010. Contact him at jon.brickey@mastercard.com.

Simson Garfinkel is the senior computer scientist for confidentiality and data access at the U.S. Census Bureau. He holds seven U.S. patents and has published more than 50 research articles in computer security and digital forensics. Garfinkel received a Ph.D. in computer science from Massachusetts Institute of Technology in 2005. He is a Fellow of the IEEE and of the Association for Computing Machinery and a member of the National Association of Science Writers. Contact him at simson.l.garfinkel@census.gov.

Latif Ladid is a senior researcher on the Faculté des Sciences, des Technologies et de Médecine, Université du Luxembourg. He is founder and president of the IPv6 Forum and board member of the 3rd Generation Partnership Project since 1999. Ladid received a post-graduate diploma in business administration studies from the Business and Management Institute of Leeds Polytechnic, United Kingdom. He received the 2002 IPv6 Forum Internet Pioneer Award and the 2016

IPv6 Life Time Achievement Award. Contact him at latif@ladid.lu.

Gary McGraw is cofounder of the Berryville Institute of Machine Learning and the author of *Software Security* (AWL, 2006) and 10 other software security books. McGraw received a dual Ph.D. in computer science and cognitive science from Indiana University. He served on the IEEE Computer Society Board of Governors and produced the monthly "Silver Bullet Security" podcast for *IEEE Security & Privacy* magazine for 13 years. Contact him at gem@garymcgraw.com.

Bruce Potter is the chief information security officer at Expel. He is responsible for cyber risk management and ensuring the secure operations of Expel's services. He cofounded Ponte Technologies. He also founded the Shmoo Group in 1996 and helps run the popular annual hacker conference, ShmooCon, in Washington, D.C. Contact him at gdead@shmoo.com.

John Viega is the chief executive officer and a cofounder of Capsule8. He coauthored the Galois counter mode, which encrypts more than two thirds of encrypted web traffic. He has coauthored many books on security, including the first book for software engineers on security. Viega received an M.S. in computer science from the University of Virginia. He has an extensive track record at both start-ups and large security companies. Contact him at john@viega.org.

SIMSON GARFINKEL: Before we consider the cyberthreats of 2025, it's important to consider the likely cyber landscape of 2025, as determined by current trends in technology evolution, policy changes, and the environment in general. Assuming the deployment of technology continues for the next five years, the year 2025 is likely to see increased deployment of partially finished, somewhat buggy, hybrid hardware/software systems

in every aspect of our technological infrastructure, with the knowledge that these systems can be patched and upgraded after they are shipped to customers. It will be common for newly purchased, newly deployed systems to immediately perform software updates on first start. This will allow for faster time to market, but the net result will be more computerization with systems that are less mature and less reliable.

NIST defines a cyberthreat as "an event or condition that has the potential for causing asset loss and the undesirable consequences or impact from such loss." Given the NIST definition and the likely technology scenarios, I suspect that we will see increased attacks with financial and political goals. On the financial side, I expect more theft and extortion. On the political side, I expect more attacks aimed at delegitimizing

governments. Those might be information operations, such as those against Facebook and Twitter that have received so much coverage, or attacks against infrastructure with the goal of delegitimizing governments by demonstrating that they cannot protect their citizens. Specifically, I am expecting threats against embedded systems, threats against the accuracy and completeness of information in consumer-facing systems, and the use of cybersystems for intelligence gathering (for example, surveillance and information theft).

Disinformation and misinformation will continue to be big problems in cybersecurity in 2025. Machine learning (ML) vulnerabilities will be a new reality by 2025. If we start working on ML security (MLsec) immediately, we may avoid some real trouble.

BRUCE POTTER: I think we are currently in a wave of commodity-grade attacks around cryptocurrency/cryptomining/ransomware that will largely have ebbed by 2025. For the most part, I expect endpoints and operating systems (OSs) to be reasonably hardened,

JOHN VIEGA: The threat landscape will evolve to keep up with the times, but, in some sense, not much will be different. There will be plenty of career criminals looking to make money via whatever is expedient (for example, cryptomining, ransomware, and so on). And then there will be nation-state attackers that are far more targeted, having vastly less day-to-day impact on businesses but a much larger geopolitical impact.

COMPUTER: Which of those prevalent cyberthreats will pose the highest risk to society?

Software-related security vulnerabilities will continue to top the most important problems for us to solve, even in 2025.

LATIF LADID: The cyberthreats will still be money oriented, but the political threats (influencing and interfering in the national political systems) will become rampant around the world. New cyberthreats will come also from the zillions of very-low-cost Internet of Things (IoT) devices shipped to the market with zero security, especially in smart cities. Obviously, 5G will enable new higher-speed attacks with massive broadcast attacks. Cryptocurrencies will be also very good candidates for hacking, including government policies to regulate them.

GARY MCGRAW: Software-related security vulnerabilities will continue to top the most important problems for us to solve, even in 2025. Although we now know what to do to make software security work, we still have a long way to go in actually doing those things.

reasonably patched, and less likely targets for adversaries. I think the focus will continue to move from access (that is, persistence on a host, OS-level access) to data compromise. Data are where the money is; data can be had in big chunks when you find them. I think attackers will continue to look for ways into data stores of interest and find ways to monetize them if they're cybercriminals or use them for intelligence operations if they are nation-states.

The other threat will be around information operations. As the Internet has democratized over the years, the ability for an adversary to leverage various platforms as part of a large-scale information operation has skyrocketed. The attack here is largely nontechnical but uses technical means for execution. Like most attacks, for a little investment, attackers get a huge return.


BRICKEY: While criminals will continue to have a negative impact on society (not just via financial crimes but also using ransomware), nation-state threats will pose the highest risk as they disrupt and destroy critical infrastructure.

GARFINKEL: Threats to democratic institutions, industrial control systems (including autonomous vehicles), and the financial sector represent the highest risk to society because they have the power to damage society in ways from which we could not readily recover.

LADID: See my response to the first question.

MCGRAW: Since software is working its way into everything we build, all aspects of society are subject to software risk. My view is that building secure software remains the most important immediate focus.

POTTER: The information operations efforts will have the largest impact. We see that today in election security. Puppet accounts and



willing media participants result in the ability for misinformation to spread rapidly around the globe. Couple that with a few targeted technical attacks (think: the reporting websites of a few swing counties in Wisconsin, Michigan, and Ohio), and, suddenly, the entire country is in an uproar. Cheap to carry out with an unbelievably large impact.

VIEGA: From a technical sense, while IoT and cloud technologies are going to be growing targets, people will always be the weakest link. The most prevalent, effective, and worrisome threats are going to be nation-state disinformation campaigns—interfering with the elections and the politics of the Western world.

I say Western world because there's a massive and obvious asymmetry that puts the Western world at a particular disadvantage—English. There are nearly 800 million nonnative English speakers in the world, spread around the world, and all fairly highly skilled. But there are fewer than 200 million nonnative Mandarin speakers and about 150 million nonnative Russian speakers. That expertise is not really concentrated in the Western world. Not being well equipped to reciprocate makes it much harder to build enough capability that would lead to some sort of mutual nonaggression agreement. For the foreseeable future, we're quite possibly stuck with subtle but tremendously effective geopolitical manipulations through the media.

COMPUTER: What types of vulnerabilities will be common in the future? Are there any technical advances that portend solving age-old problems, such as those that take advantage

of a computer's memory-management systems (such as Spectre and Meltdown)?

BRICKEY: With the burgeoning amount of code designed for the cloud and containerless computing environments as well as the lack of secure coding training/awareness, there are bound to be key vulnerabilities in the future. There will likely be major vulnerabilities discovered in cloud computing that will impact the majority of enterprises.

GARFINKEL: Spectre and Meltdown caught many organizations off guard: they were viewed as fundamentally new kinds of exploits. Given the rate of innovation, we're likely to see yet another fundamentally new exploit before 2025.

for defensive purposes is likely to result in increased use of formal methods to find data-dependent exploits.

LADID: Mitigating cyberattacks is a mix of policy and new mitigation technologies. The Finnish STRATCOM regulator is an excellent reference cybersecurity model to be adopted by the countries that can adopt this centralized cybersecurity model.¹

McGRAW: The future is here, but it is sparsely distributed. Though we have been making technical progress at the bleeding edge, older tech with known problems will still be prevalent in 2025.

POTTER: Vulnerabilities will shift to ones of process versus low-level



Puppet accounts and willing media participants result in the ability for misinformation to spread rapidly around the globe.

More concerning in my mind, though, is improved attacker fluency in exploiting existing vulnerabilities. For example, in 2010, Kris Kaspersky and Alice Chang demonstrated that remote code execution was possible by exploiting bugs in specific versions of Intel CPUs: using these bugs requires that the attacker have knowledge of specific targets, but the bugs are stealthy and the payoff can be huge, because the bugs can't be demonstrated on processors that aren't vulnerable. Like Spectre and Meltdown, these bugs exist because the silicon in the CPU doesn't faithfully implement the application binary interface. Likewise, I expect that increased use of formal methods

technical defects. For example, a continuous integration and continuous delivery pipeline that utilizes a public bucket for code or a misconfigured access control list at an integration test provider can give up all your source code without even touching your network. With the continued migration to infrastructure as a service (IaaS) and software as a service (SaaS) providers, even by huge companies, the ability to configure everything in a secure manner becomes far more important than the impact of side-channel attacks on processors like Spectre and Meltdown. Helping organizations understand their new attack surface and how to manage security on hundreds of SaaS solutions

is where the next big investment in tech is going to happen. Cloud access security brokers (CASBs) and others aren't going to cut it.

On the host, I think we will continue to see the “drip, drip, drip” of research on Spectre-like attacks. These vulnerabilities are an indictment of the last two decades of chip advancement. We made systems faster by doing everything at once and later figuring out what the right path was. Turns out, that's hard to do without leaking information. I imagine the chip manufac-

The rise of surveillance capitalism shows just how clueless most users are when it comes to their own behavioral data. This will get worse, not better.

turers will figure that out (much like we've figured out how to deal with differential power analysis and other low-level attacks), but it will come at the cost of performance and economic efficiency.

VIEGA: Processor-level problems will continue to be a ripe source of challenges due to their complexity. But most of them have been, and will continue to be, lightly exploited. Attackers follow the path of least resistance, and there will generally be much easier ways in. People will still be the weakest link, and misconfigurations that lead to issues will be common. Insofar as vulnerabilities go, the continued move away from C and C++ for applications will be a good thing, but we're unlikely to see a decline in good old input validation problems like basic command injection. They're easy to add and often hard to protect against generically.

COMPUTER: In 2025, what will the general public's expectations be for cybersecurity and online privacy?

BRICKEY: Consumers may have more control over the use of their data by 2025, so online privacy expectations will likely increase. By 2025, several U.S. states will have their own privacy laws, and there will be increasing pressure on the federal government to pass a national law. I doubt there will be much higher expectations for cybersecurity, as pervasive IoT makes it more difficult

than ever for the public to keep tabs on the security of their devices.

GARFINKEL: I expect that the general public will become increasingly fatalistic about the state of cybersecurity and resigned to the existence of exploitable bugs. I think that such attitudes are problematic, as they are likely to result in decreased cybersecurity funding over time.

LADID: Privacy has been dead for a long time, starting with the yellow pages and loyalty cards. Privacy is very big business. Privacy by design is possible only if it is integrated in the OS. As a case in point, Internet Protocol version 6 (IPv6) has a privacy protocol, unlike IPv4. The Mac address is scrambled so that footprints won't be harvested by web scanners.


McGRAW: The public seems to be completely in the dark with respect to

computer security and privacy, especially in the United States. The rise of surveillance capitalism shows just how clueless most users are when it comes to their own behavioral data. This will get worse, not better.

POTTER: I hope the members of the public at large will actually feel like they have some control over their lives and personal data online. As laws like General Data Protection Regulation (GDPR) spread around the world, consumers will feel more empowered to control their data and will have more visibility into where their data are. That said, I believe that most consumers have a general expectation of security and privacy but really aren't equipped (nor should they have to be) to understand it fully. For instance, when installing a video camera that is protected only by a username and password, a consumer should know a) that's not particularly secure and b) don't reuse passwords. But as the recent rash of Ring camera compromises have shown, many consumers don't understand that. It's on companies like Ring to require multifactor authentication to protect consumers, not on consumers to demand it.

VIEGA: People in our society want to feel safe, and if there's not an obvious daily threat, they'll continue to feel that way. I think people will know there's some risk but feel mostly safe and almost wholly unconcerned, whether it's warranted or not.

COMPUTER: What advances in cybersecurity defenses and privacy protections should we be investing in now to address the cyberthreats anticipated to be prevalent in 2025?



BRICKEY: We need to do better at data-centric cybersecurity strategies. We need to see chief privacy officers, chief information security officers, and chief risk officers working closer together for additional protections. Boards of directors will require they do this, and limited budgets will force them to find dual-purpose solutions.

GARFINKEL: Our best defense would be to build systems that have dramatically fewer vulnerabilities. We can do this by significantly reducing our reliance on software written in unsafe programming languages, increasing the use of formal methods, and isolation approaches.

For example, today, compiled programming languages like Go and Rust offer the speed of C/C++ while offering memory and type safety. JavaScript is also fast enough for a large number of applications, as evidenced by the success of node.js. As a result, new projects shouldn't be started in C/C++, and existing code should be rewritten. Sun had an effort in the 1990s to rewrite the Unix kernel in Java: it failed. Sun made two errors: Java in the 1990s didn't offer sufficient performance, and the project sought to redesign the entire OS from the kernel up. Since most cybersecurity problems are in applications and libraries, it makes more sense to start rewriting network servers, system libraries, and command-line tools, leaving the kernel for last. Until then, we should significantly increase the use of automated static analysis tools for existing code.

LADID: Wars have been stopped with peace treaties. Cybersecurity needs peace treaties around the world to first address government threats. This is the

biggest task for the next 10 years and even beyond. But we will have still a cold cyberwar for years to come.

McGRAW: Educating the public about privacy risks and the poor tradeoffs people are making when they use much advanced technology.

POTTER: Data mapping is a huge one. Understanding which data move where in an organization is key, especially as attacks continue to focus more and more on data themselves. To prop-

with embedded devices, since they're becoming ubiquitous and are generally both quite vulnerable and incredibly difficult to keep patched.

COMPUTER: Given that it is hard to predict five years out what information and computing technologies will be ubiquitous, how the systems those technologies are used in will be developed and sustained, and how those technologies will be employed by users of information and communications technology (ICT), how much



To properly implement security controls and have any hope of protecting personal information, companies need to know where their data are.

erly implement security controls and have any hope of protecting personal information, companies need to know where their data are. This is easy in small companies but incredibly difficult in large ones.

Also, a general focus on data minimization will help as well. Companies view data as assets because data feed analysis activities and can be used to better understand their products and customers. However, data are also a liability and need to be treated as such. This is an engineering discipline that needs to spread across the industry.

VIEGA: The security industry needs to invest in keeping up with new technology. Most people in the industry really don't get cloud technology, they don't understand containers, and they don't understand anything in the ecosystem of modern software (for example, service meshes). Also, the industry needs to figure out how to do better

credibility can we place in educated guesses about what the future cyberthreats will be? Are there too many unknowns?

BRICKEY: There are already too many unknowns, and we know that trends in the IoT (to make consumers' lives more convenient) will add to the complexity. There's an app for everything, and that will expand to include autonomous vehicles, robotic assistants, and artificial intelligence (AI)-enabled assistants. Consumers will expect more of this, and companies are more than willing to deliver and gain a foothold on the market.

GARFINKEL: "The future has arrived—it's just not evenly distributed yet." This quotation, frequently attributed to cyberpunk writer William Gibson, can be easily applied to the state of ICT in general. Most of the likely uses of computers and networks in 2025 are

already employed today; even the popular apps of 2025 probably exist today in some form.

For example, many people see ransomware as a relatively recent phenomenon, but the first ransomware, the so-called AIDS Trojan, was deployed in 1989; a paper describing a more sophisticated system was presented by Young and Yung at the 1996 IEEE Symposium on Security and Privacy (“And there is nothing new under the sun,” Ecclesiastes 1:9).

Most of the likely uses of computers and networks in 2025 are already employed today; even the popular apps of 2025 probably exist today in some form.

This is good news for prognosticators: predicting the future is really just an exercise in understanding the present. We can improve our forecasting ability by studying the progress of earlier IT transitions. But it’s important to pick your examples carefully. The early history of the Internet looks a lot like the early history of radio, but not at all like the early history of television. I think that this is because both radio and the Internet were conceived and initially deployed as two-way communications systems, but when they went mainstream, they both evolved into predominantly one-way systems for disseminating information.

Looking specifically to cyberthreats, I think that current trends are likely to accelerate and that systems used by governments and major corporations in 2025 will be more restrictive, more locked down, and more utilitarian than anything widely used today, while educational institutions, small

businesses, and individuals will enjoy increasingly sophisticated and powerful tools. Balancing productivity and risk will remain a difficult task for security professionals.

LADID: This is a very tough question. However, the Finnish example has shown very good results. It’s a good start but needs to invest in mitigating new cyberthreats enabled by new technologies. The speed of hackers is accelerating.

The users are employing new made-easy technologies, such as iPhones, without understanding how they function and how to protect them. I was at a university a week ago, and the university has introduced a computer course for beginners, as the new students do not even know how to use a mouse or a keyboard, let alone go into Windows, as they are used to just swiping on the iPhone screen or using face recognition.

McGRAW: Everyone is a BSer!

POTTER: There are a ton of unknowns. But we’re at a point in the industry where we have enough historical data of what global computing habits look like such that we can make more educated guesses than in the past. While we can’t predict revolutionary products that will change the world, we can predict roughly how they will communicate, the types of data they will have,

and the systems that will be involved. The push toward cloud computing by basically everyone allows us to predict vulnerabilities and controls, even if we aren’t sure of exactly the cloud service that will be deployed.

VIEGA: I think it’s safe to say attackers will generally follow the path of least resistance. What that path tends to look like will vary based on how the technical landscape evolves and how the security industry evolves. In five years, there are bound to be a few surprises, but it’s a short enough time horizon that I don’t think we’ll be living in a fundamentally different world.

COMPUTER: Cybersecurity guidance typically recommends rapid application of updates and patches, with the implicit assumption that original installs and updates are free of malware. But there are increasing reports of malware contained in code coming from the original vendor, particularly for smartphones. What are the prospects for containing this threat?

BRICKEY: I think we’ll see some improvements in this area due to more attention to supply chains and the inherent risk in third parties. There is a growing number of products and processes, like the blockchain technologies, to verify code at inception and throughout its lifecycle; however, this is not likely to be common practice in the near future.

GARFINKEL: The threat of malware and vulnerabilities delivered with vendor software is nothing new: the 1988 Morris Worm spread, in part, by using a “trap door” created for debugging that allowed remote code execution. Ken Thompson, in his Turing

Award lecture, discussed how another back door could be put into the C compiler, although it is unclear if an infected C compiler was ever distributed by AT&T.

It is conceivable that the risk of vendor-supplied malware and back doors might be addressed by using formal methods and proof-carrying code. But this could only work for a very restrictive set of programs whose behavior can be formally specified. Here, it is useful to remember that the definition of a trusted system is one that can violate your security policy: this applies to both software and vendors alike.

LADID: This is incorrect. Microsoft spends a lot of effort to patch its security software holes, and people don't even pay attention to updating their OSs. Same for Apple. Obviously, software is not a 100% science, so it will be always attackable, and the updates should be made mandatory, a tough message to get into 5 billion simple users. Automating updates as Apple does during the night is a good solution, but the device has to be connected and powered during the night. I cannot get this done myself.

McGRAW: There is no avoiding the patch, unless we build perfect software to begin with. Since that's not possible, we'll always be stuck with patching.

POTTER: Pay more money for your products. When things are cheap/free, you are the product. We're seeing this in low-end Android smartphones because the vendors either have poor cybersecurity practices or they are finding other ways to monetize their platform. When companies are properly compensated

for the products they make, they have more resources to dedicate to building a higher-assurance product.

VIEGA: Not good. I think back to a story Steve Bellovin told me a long time ago, about a person at Bell Labs who was responsible for a bug so egregious and easy to exploit that he was personally pretty convinced it was an intentional back door. But there was no way to be sure that it wasn't just a stupid error.

Now, take into consideration the trend of systems getting larger and larger based on the number of lines of code, primarily because projects are shifting from "mostly proprietary code and a little bit of open source" to "mostly open source and a little bit of proprietary code." (I've heard people claim that we've gone from 10% of new code being open source to 90% over the last 10–15 years, and that feels right.) There's a lot that can (and will) continue to go wrong here.

COMPUTER: Adversarial images are a popular research topic, but we have not seen real-world exploitation. As autonomous systems become prevalent, how concerned should we be about this cyberthreat?

BRICKEY: We should be concerned about this, but mostly on the lower end of the impact scale—as a nuisance to society. It could be a larger threat as our military systems are deployed around the world and as adversaries share knowledge about vulnerabilities and exploits.

GARFINKEL: Adversarial images, like putting a sticker on a stop sign to make a convolutional neural network think that it's a "speed limit 55" sign, makes

for a popular research topic because the results are both dramatic and easy to understand. We haven't seen them in the wild because the cross section of autonomous cars and miscreants is currently quite small. However, the availability of high-power laser pointers has resulted in several pilots being dazzled while they were landing commercial passenger aircraft, so I think that we should be very concerned about all sorts of malicious threats that autonomous systems may face from miscreants.

LADID: I have two European Union (EU) projects: 5G-DRIVE and 5G-MOBIX, where we are addressing autonomous networks, and, indeed, the security, privacy, and cybersecurity issues are not yet comprehended properly, as we need large-scale deployment to assess the new risks coming from the 5G ecosystem around moving vehicles but also from all of the functionalities coming from the vehicles, the multi-access edge computing side, the backbone, and so on.

It's a huge task, as different standards bodies have to join forces as their standards merge in these scenarios, and, obviously, these bodies do not really understand each other due to different terminologies and ontologies at different layers of the stacks. This will be up to the most innovative mobile operators to create a special consortium for each vertical (driving, health, energy, and so on) with the corresponding vertical industries. China Mobile has set up four verticals with a dedicated company for each one to look at the entire scenario and do massive trials.

As a case in point, for instance, in Kuwait, the main mobile operator ZAIN is one of the first 5G deployers, as

the spectrum was given free of charge to its mobile operators, and ZAIN has created a company to use drones to control the oil fields. It's worth a couple million dollars.

These verticals will be, for sure, the biggest topics in this decade for cybersecurity, though industry might get it right but could hide their issues like anyone else.

McGRAW: MLsec is far more involved than adversarial images. ML systems are riddled with risks. At the Berkeley Institute of Machine Learning, we have identified 78 risks that should be carefully considered and mitigated when engineering is building and deploying ML. ML is used for far more than simply autonomous systems, for what it's worth.

POTTER: I think we're still a long way from autonomous systems being part of our daily lives due to overt threats, like adversarial images, and inadvertent threats, like "reality is more complicated than a lab environment." In closed environments, I think we can remove the human element pretty easily. But in open-world environments, we're still miles away from being able to let these systems run free. Once adversarial images aren't a concern (that is, the technology is good enough to recognize the difference between a pineapple and an owl, or a piece of electrical tape doesn't change a 35 mi/h zone into an 85 mi/h zone), then we can have autonomous systems everywhere.

VIEGA: For image processing, if I want to break onto your phone, I can still use a passcode, hold the phone up to your face, or use a zero-day to just break onto the phone. These options

are generally easier, and we can expect people to follow the path of least resistance. That's a testament to the quality of Face ID for raising the bar.

Beyond authentication, I haven't seen ML techniques provide the most effective lines of defense. If they did, then I'm sure adversarial learning would have more of a real-world impact. In general, the industry is heavily leveraging algorithms, but without enough regard to the quality of the inputs to those algorithms.

COMPUTER: Governments and corporations provide extensive advice on cybersecurity. Are average users becoming more aware of cyberthreats and of how to protect themselves? Are they following the security experts' advice?

BRICKEY: Average users are becoming more aware, but it's difficult for people to keep pace with changes in technology. Unfortunately, as more and more people experience a cyberattack of some sort, they are getting more familiar with the concept. Groups like the AARP, Security Advisor Alliance, and a number of other organizations (such as the Global Cybersecurity Alliance) are spreading the word to increase awareness; however, the attack surface is increasing quicker than our ability to train people at scale.

GARFINKEL: These are questions of fact that need to be answered with research. For example, a 2019 survey of about 4,800 participants by InternetNZ found that 10% of respondents were "a little or not at all concerned about security," a finding that the authors found "worrisome."² A U.S. survey of 2,000 students at a large public university and its branch campuses located

in the U.S. Pacific Northwest garnered a 24.9% response rate and found that 57% had personally experienced a cyberthreat, 77% have antivirus software installed on their computer, 55% had made an online purchase while logged onto a public Wi-Fi network, and 37% had shared their passwords with another person. Perhaps most significantly, "32% of students agreed that the steps needed to protect their online security and privacy is too overwhelming to think about."³


So, some people are following expert advice, and others aren't. Ideally, we would build systems that make following advice less important. Experts should also offer better advice. For example, there should be no security risk when using a modern web browser on a public Wi-Fi network.

LADID: Computer emergency response teams (CERTs) do an excellent job. However, looking at the cybersecurity map (<https://globalsecuritymap.com/>), I am not sure CERTs are efficient in their tasks.

There is a huge lack of skilled cybersecurity experts, even of security experts. Security is a very expensive investment with no immediate return on investment; hence, it is delayed or left until a disaster happens, and then the chief executive officer understands its value, but then he or she will still see it as a penalty.

McGRAW: No. Users remain blithely ignorant, and too much generic computer security advice is inconsistent and even incorrect.

POTTER: Eh. I think people continue to think about cybersecurity wrong, and that's OK. Properly designed systems shouldn't require experts' advice



to be used securely. They should just work that way. Plus, people overpersonalize threats and believe that at every Starbucks there's a criminal trying to steal their credit card and that every public charging cable is trying to hijack their phone. In reality, *these threats don't even exist*, but it's top of mind for many consumers. They've heard messages about wireless snooping and plug jacking and think it's happening everywhere. It's actually happening nowhere. So, let users do their thing, and let's build them more secure systems.

VIEGA: No—should we expect them to? I think the security industry is in a bubble, and it expects security to be at the top of every company's and every person's priority list. Often, it's so far down the list that it takes regulatory compliance to drive significant spending.

As for users, if they live their lives in so much fear that they worry about cyberthreats day to day, then the industry on the whole has failed them because almost nobody wants to live that way.

COMPUTER: Does cybersecurity guidance generally keep pace with new threats, or if not, to what degree does it lag?

BRICKEY: I would say it lags significantly. Larger companies and governments do a pretty good job of keeping pace, but the smaller the enterprise, the less likely that it can keep up.

GARFINKEL: In my experience, cybersecurity guidance keeps pace with new threats but not with cybersecurity research. For example, many organizations that received guidance were fast to install software patches against the

Meltdown and Spectre side-channel vulnerabilities, but those same organizations have not revised their policies to allow password managers or remove the requirement for quarterly password changes. Another example: NIST 800-63B, "Digital Identity Guide-

really comes in waves. As new technology is developed, we deploy first and then think about security later. Look at Kubernetes as an example. There are a lot of companies moving to Kubernetes, but there's nearly zero accepted guidance on how to do it

Generally, the industry is incredibly reactive, always being asked to secure the projects and technologies other people already ramped up.

lines," revised June 2017, states, "Verifiers SHOULD permit claimants to use 'paste' functionality when entering a memorized secret. This facilitates the use of password managers, which are widely used, and in many cases, increases the likelihood that users will choose stronger memorized secrets." Nevertheless, many financial and government websites disable the paste button on critical forms.

LADID: The one who has the skills is first, and, obviously, it's on the side of the attacker. New AI-based cybersecurity simulate attacks on the networks to constantly fix the holes in the networks and predict the attacks. Investing in prevention, similar to emergency and safety services, is a way forward. But we do not have too many security software companies with worldwide span, apart from one or two.

MCGRAW: Depends on who we're talking about. If we want to make a dent in the security problem, we need to focus on educating and informing engineers and operators of advanced systems.

POTTER: For enterprises, I think guidance for the secure use of systems

securely. Most companies are figuring it out as they go and are trying to put in appropriate controls. But we're a long way from where (for example) we are with enterprise endpoints. You want to know how to securely deploy and operate laptops? WOOT! There are mature products, a huge body of knowledge, and lots of good processes. Technologies like Kubernetes will get there eventually, but security guidance will always lag.

VIEGA: Generally, the industry is incredibly reactive, always being asked to secure the projects and technologies other people already ramped up. It also doesn't help that the industry was originally stocked with network people, and the world is much more software driven. I'd say the security industry is generally three to five years behind the rest of the tech industry, and I don't see that changing quickly.

COMPUTER: Cyberthreats arise from the vast amount of consumer data collected and stored, an attractive target for cybercriminals and other malicious actors. But new privacy laws, such as the EU's GDPR, give the public more control over their data. How

will evolving privacy rules affect the cybersecurity landscape?

BRICKEY: They should help consumers with their data, but they won't help security much. The only benefit for consumers is that companies will be held financially liable with these new laws, but it will increase the cost of doing business, and that will be passed to the consumers.

GARFINKEL: Some new privacy laws come with heavy fines and statutory damages for the misuse of consumer data or for data breaches that result from data-handling practices that are inconsistent with industry standard practices. Such laws may cause companies to take cybersecurity more seriously. However, these laws may also have the perverse side effect of encouraging hackers to break into corporate data banks so that the companies can be extorted.

Overall, it's important for security professionals to remember that the main purpose of these privacy laws is to restrict what authorized users can do with the data that they already have in their possession. As such, I expect these laws to stimulate interest in access controls and auditing—two areas that have not received much attention in recent years.

LADID: Not sure GDPR is resolving any problem; on the contrary, it is now overabused by even serious websites to get people to sign off on the cookies to have access. That was not the intended purpose of this bureaucratic idea. Don't get lawyers involved in technologies they don't understand; there are a thousand loopholes.

McGRAW: They won't matter much at all, sadly.

POTTER: I think as these laws spread around the world, the impact will be profound. Companies have largely been able to do anything they want with personal data, assuming they keep them secured. Users are starting to have a say, and it's going to require companies to have much better operational control than they've had in the past. For instance, requesting a company to delete all their data about me is likely a very manual process for most U.S.-based companies. Once this right is codified into law, companies will have to make investments in automating that process. The side benefit is that they will centralize data and get better control over it, which will lead to better hygiene and security in general.

VIEGA: It will drive more spending, but in a way that continues to put an oversized burden on development organizations. I think the world can absorb it and benefit from somewhat better control over their data and somewhat less risk. But I haven't seen enough evidence yet to say for sure.

COMPUTER: Are there any promising new security or privacy paradigms for addressing cybersecurity threats?

BRICKEY: The most promising aspect of cybersecurity is developing a better workforce and training them to solve problems and think like our adversaries. No specific tool is going to help, but our workforce—armed with good technology and processes—can go a long way.

GARFINKEL: Two technologies that we have been exploring at the U.S. Census Bureau are differential privacy and secure multiparty computation. Both of these offer new approaches for

protecting the privacy of confidential information about individuals while allowing that information to be used in a way that benefits the public good. These technologies have rapidly matured in recent years and are now ready to be deployed at scale to solve a wide number of problems.

LADID: The vendors (all industries and their shops) are exploiting privacy to get their customers to buy more of their stuff. Go to a marketing conference and talk about respecting privacy, you will be looked at as communist against the cherished capitalism. I had that experience, so I now avoid marketing conferences.

McGRAW: We need to focus some of our existing ideas around architectural risk analysis on ML.

POTTER: I think the time for formal methods may finally be at hand. The work that Google and others are doing on formal methods is showing promise of working at a scale we haven't seen yet. Combined with the push toward more SaaS offerings, this means that code will be more centralized when compared to code in legacy enterprises. It means we may finally have a chance of verifying code and helping protect a large number of people with a small tech investment.

VIEGA: In terms of tackling the biggest problems we have, such as our social media problem, I don't see any technology emerging over the next few years that I expect will be a magic wand, if that's the question.

Security is like the nuclear arms race—there's a feedback loop pushing both the defenders and attackers to get better. Things are generally improving,

and as an industry we've made some huge improvements that are worth celebrating. But a definitive solution to any one problem just shifts the playing field; it doesn't stop the game.

There are several takeaways from this VRT, one of them being that, by the year 2025, one of the leading types of cyberthreats will be the misappropriation and use of data to inflict geopolitical instability. The disturbances attackers can effect will become increasingly amplified as society becomes ever more dependent on being always connected through the cyberphysical systems that comprise the IoT. These disturbances may be further amplified by the rapid adoption of ML and higher levels of machine intelligence into the full gamut of systems upon which society relies for peace and stability.

DISCLAIMER

The views and conclusions contained herein are those of the panelists and authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of their employers. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes, notwithstanding any copyright annotations thereon. Certain commercial entities, equipment, or materials may be identified in this document to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

Another takeaway is that the cat-and-mouse game between the defender and attacker will continue to be asymmetric; that is, the maxim will continue to hold that the effort expended by the defender is much greater than that of the attacker. DevOps, as practiced in 2025, could play into the hands of attackers unless the security community shores up the protection of our software-development processes. In addition, the prevailing view among the participants in this roundtable is that software quality will continue to be poor in terms of reliability and security. There is hope that developers will adopt good security and software engineering practices, such as formal methods and secure programming languages, but the view is that economics, not legislation or policy, will drive software quality.

A third takeaway is that the complexity of cybersystems and the information infrastructure on which they operate will continue to grow, with guidance on security and privacy practices continuing to lag well behind the adoption of new information technology. Furthermore, as with today, in the year 2025, users cannot be expected to understand the detailed technical aspects of the risks that cybersystems and information infrastructure pose to them. However, security and privacy experts will need to do a better job of hardening systems and infrastructure, making them wear their security and privacy policies on their sleeves. We thank the participants of this VRT for giving us insight into the set of cyberthreats and attendant challenges we can expect to encounter in the year 2025. ■

REFERENCES

1. L. Ladid and J. Armin, "Whitepaper: The Finnish electronic

communications regulator TRAF-ICOM—A cybersecurity reference model for Europe," European Union Systemic Analyser in Network Threats Consortium, Athens, Greece, White Paper, Feb. 2019. [Online]. Available: https://project-saint.eu/sites/default/files/whitepaper_1_ul_cybe_final_1.pdf

2. S. S. Tirumala, M. R. Valluri, and G. Babu, "A survey on cybersecurity awareness concerns, practices and conceptual measures," in *Proc. 2019 Int. Conf. Computer Communication and Informatics (ICCCI)*, Coimbatore, India, pp. 1-6. doi: 10.1109/ICCCI.2019.8821951. [Online]. Available: <https://ieeexplore.ieee.org/document/8821951>
3. D. Sarathchandra, K. Haltinner, and N. Lichtenberg, "College students' cybersecurity risk perceptions, awareness, and practices," in *Proc. 2016 Cybersecurity Symp. (CYBERSEC)*, Coeur d'Alene, ID, pp. 68-73. doi: 10.1109/CYBERSEC.2016.018. [Online]. Available: <https://ieeexplore.ieee.org/document/7942427>

JAMES BRET MICHAEL is a professor in the Naval Postgraduate School's Computer Science and Electrical and Computer Engineering departments. Contact him at bmichael@nps.edu.

RICHARD KUHN is a computer scientist in the Computer Security Division at the National Institute of Standards. Contact him at kuhn@nist.gov.

JEFFREY VOAS is the editor in chief of *Computer*. Contact him at j.voas@ieee.org.



Reliability Inversion: A Cautionary Tale

Behrooz Parhami, University of California, Santa Barbara

Reliability analysis is often based on worst-case assumptions to produce guaranteed lower bounds on system survival probability. Reliability engineers make lower bounds as tight as possible, but sometimes system structure is unfriendly to the derivation of tight bounds. Unfortunately, loose reliability lower bounds make it difficult to compare design alternatives or to select among competing systems.

Reliability inversion is a new concept being introduced in this article for the first time. Briefly, it leads to a less reliable system being deemed more reliable because of uncertainties in reliability modeling. We will define the idea in greater detail in the next section. Uncertainty in

reliability estimates makes the selection of the most reliable design or system a challenging task, regardless of how the uncertainty is represented: probability, possibility, fuzzy, rough sets, intervals, and the like.¹ The greater the uncertainties, the harder the comparison. When reliability modeling leads to large uncertainties, we might say that the system is not (easily) modelable.

Besides well-known “ilities” (reliability, availability, and other attributes described by words ending in

Digital Object Identifier 10.1109/MC.2019.2958907
Date of current version: 4 June 2020

“ility”), dependable system operation is also contingent on lesser known “ilities” (performability, testability, serviceability, and so on). We propose *modelability* as a new addition to this group of terms. When used qualitatively, the term refers to the ease of accurate reliability modeling. Similar to testability and a number of other “ilities,” which were first introduced as qualitative notions and later quantified, we hope that modelability can someday advance to the quantitative domain.

Modelability is of the same nature as (design for) analyzability, also known as design for analysis,² itself predated by concepts such as design for manufacturing (manufacturability). Analyzability requires honoring certain design constraints that allow the use of simpler tools for analysis. In the domain of electronic circuits, design for packageability³ is quite similar. Both notions constrain the design process, which may seem to lead to higher costs and longer design times. However, somewhat counterintuitively, the end result is often economy and shorter time to market.

RELIABILITY INVERSION DEFINED

The exact reliability of a system is often unknowable. If we had hundreds of identical copies of a system and could run them for decades, observing system failures, we could ascertain the actual reliability with high confidence. A large number of copies and long running times would be needed because, at typically high system reliabilities, failures are extremely rare; so to obtain statistically valid results, extensive data collection is required. An alternative is to make simple, pessimistic assumptions about

subsystems and their interactions, in an analytic or simulation model, to derive a lower bound on reliability. Models do not completely eliminate the need for experimentation as model parameters may be derived, and models themselves tuned, based on experimental observations.

The actual system reliability could be much better than a model-based lower bound. We see in Figure 1 that even though System A is more reliable than System B (if we somehow knew the actual reliabilities), the model-based lower bounds ascribe a higher

reliability to System B. We thus have no choice but to recommend System B over System A as being more reliable. This situation is what we call *reliability inversion*, in analogy to the similarly disruptive phenomenon of priority inversion in real-time task scheduling⁴ that wreaked havoc during the Mars Pathfinder mission of the late 1990s.⁵

Reliability is, of course, a function of time. Generally, one cannot say that a system is always more reliable than another one. One system may be more reliable for short mission times, while another fares better for long mission durations. So, let us enter the time factor into the notion of reliability inversion. The actual and modeled reliabilities of Systems A and B are depicted in Figure 2. With regard to unknowable actual reliabilities, System A is better for short mission times, whereas System B does better over the long run. With regard to model-based bounds, however, System B is uniformly better and would be the preferred choice in all cases.

We may call a system for which the guaranteed lower bound is very close to actual reliability a *highly modelable* system. Conversely, a system has poor modelability when the bound is much

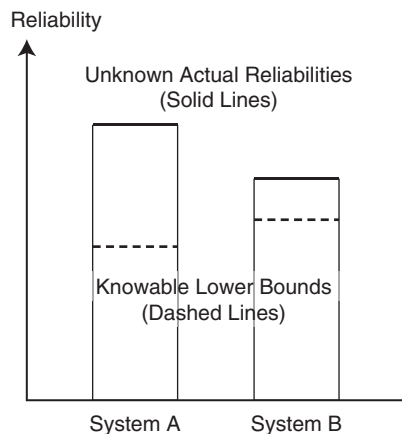


FIGURE 1. Reliability inversion.




FIGURE 2. True reliability versus modeled lower bound.

lower than the actual reliability. Of the two systems depicted in Figure 2, System B has better modelability than System A, although its actual reliability is worse for short mission times. If we were to choose System A or B for a particular critical application, we would choose B because we have no way of knowing the true reliabilities. All we have to go by are the bounds provided by reliability models, and the bound for System B is uniformly better than that of A.

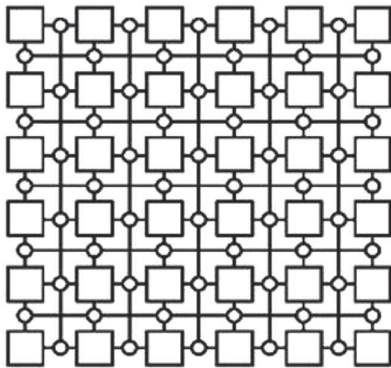


FIGURE 3. A square 5×5 array of PEs with a spare row (bottom) and a spare column (right).




FIGURE 4. The array of Figure 1, configured to salvage a 5×5 healthy array from a 6×6 injured one.

It may be argued that reliability inversion is a blessing in disguise. Because models are imperfect, in the sense of not taking all failure causes and mechanisms into account, perhaps the wider gap between the lower bound and the actual reliability can provide a safety margin to guard against unpredictable or overlooked failure causes and mechanisms. However, best practices in reliable system design and tenets of safety engineering require us to provide deliberate and predictable safety margins, rather than rely on a margin materializing by happenstance. While it is true that playing too close to the edge may be dangerous, especially in highly complex systems,⁶ we prefer to distance ourselves from the edge deliberately, rather than haphazardly.

RECONFIGURABLE PROCESSOR ARRAYS

In this section, we introduce a class of reconfigurable processor arrays for use as examples to demonstrate reliability inversion. In particular, a special case of redundancy and reconfiguration in which an $n \times n$ mesh or grid of processing elements (PEs) is augmented with one spare row and one spare column, for a redundancy ratio of $(2n + 1)/n^2 = O(1/n)$, along with embedded switches that allow processors to change their row or column neighbors when nodes malfunction. This constitutes a good example to pursue, in view of its extensive assessment and documentation.^{7,8}

In the references just cited, and the examples we will draw upon, reconfiguration is performed to return a processor array with malfunctioning nodes to its initial healthy configuration to be able to execute the original $n \times n$ mesh algorithms without modification. Specifically, we are not considering the

kind of reconfiguration that extends the computational power of the array (in a complexity-theory sense), allowing it to achieve significant speedup in performing certain computations via dynamic adaptation.⁹

To make the examples even more concrete, we will consider a 5×5 guest array within a 6×6 host array, that is, one with a spare row (at the bottom) and a spare column (on the right), as depicted in Figure 3. Originally, the nodes in the topmost five rows and the leftmost five columns are active, with the configuration changing as nodes malfunction. When a PE becomes unusable, it can be dealt with in various ways. It can be bypassed in its respective row or column, and/or it can be configured out by downward shifting the rows or rightward shifting the columns (Figure 4).

We will not discuss the details of the switching mechanisms and algorithms that affect reconfiguration,¹⁰ mentioning only that any double-PE malfunction can be tolerated through reconfiguration, but there are worst-case patterns of three unusable PEs that exceed the scheme's reconfigurability.¹¹ As can be seen in Figure 3, we have 60 switches, arranged on tracks between PE rows/columns, to allow salvaging a 5×5 guest array from a 6×6 host. More generally, given an $n \times n$ original array embedded in an $(n + 1) \times (n + 1)$ augmented array, the number of switches required is $2n(n + 1)$, that is, linear in the number of PEs.

CENTRALIZED VERSUS DISTRIBUTED SWITCHING

To demonstrate that reliability inversion is not just a theoretical curiosity, we show that it can occur in actual systems under realistic conditions. We consider the reconfiguration scheme

depicted in Figures 3 and 4 as an example, focusing on a 5×5 guest network embedded in a 6×6 host array. The system remains functional after reconfiguration if all of the switches work and if 34 of the 36 PEs are functional. Let the PE failure rate be λ and the switch failure rate be σ . Then,

$$\text{Module/PE reliability} = r = e^{-\lambda t} \quad (1)$$

$$\text{Overall switching reliability} = e^{-(60\sigma)t} \quad (2)$$

$$\begin{aligned} \text{System reliability} \\ = e^{-(60\sigma)t} R_{34\text{-out-of-36}}(r), \end{aligned} \quad (3)$$

where $R_{k\text{-out-of-}n}(r)$ is the k -out-of- n reliability for modules of uniform reliability r . Computationally,

$$\begin{aligned} R_{34\text{-out-of-36}}(r) &= r^{36} + 36r^{35}(1-r) \\ &\quad + (36 \times 35 / 2)r^{34}(1-r)^2 \\ &= r^{34}[r^2 + 36r(1-r) \\ &\quad + 630(1-r)^2] \\ &= r^{34}[595r^2 - 1224r \\ &\quad + 630] \\ &= r^{34}[1 + (1-r)(629 \\ &\quad - 595r)]. \end{aligned} \quad (4)$$

Substituting (4) into (3) and using $\sigma = 0.01\lambda$, we get the reliability plot shown as a gray line in Figure 5.

We next consider a reconfiguration scheme based on the use of multiplexers (muxes) within PEs, so that each PE can select its north/above and west/left neighbors from among three possibilities, as shown in Figure 6(b). Again, we delete some details that demonstrate the equivalence of the two schemes with regard to reconfigurability. Now, each PE becomes a tad more complex, increasing its failure rate to $\lambda + \alpha\sigma$, where σ is the failure




FIGURE 5. System unreliability for a reconfigurable array of PEs as a function of λt for one PE.




FIGURE 6. (a) External (centralized) switches can be replaced by (b) muxes within PEs (distributed).

rate of the original track switches and α is the distribution overhead, representing the increase in switch hardware complexity as a result of the distribution process. We now have the following reliability equations:

$$\text{Module/PE reliability} = r' = e^{-(\lambda + \alpha\sigma)t} \quad (5)$$

$$\text{System reliability} = R_{34\text{-out-of-36}}(r'). \quad (6)$$

In our numerical example, we take $\alpha = 2$ as a reasonable pessimistic value, given the presence of $60/36 \cong 1.67$ switches per PE in the centralized scheme, with a 2×2 switch built from two 2-to-1 muxes. The distributed scheme needs two three-input muxes per PE.

The resulting unreliability curve is shown as the heavy black line in Figure 5. We note that the reliability advantage of the distributed scheme declines as λt increases. This is because for large λt values, PE malfunctions will dominate, making switching differences less relevant. If we extend the curves for even larger values of λt , say, up to one, unreliabilities will approach one, rendering the systems both indistinguishable and practically useless.

DEMONSTRATING RELIABILITY INVERSION

We see that for λt values in the range of practical interest, distributed switching offers uniformly higher reliability lower bound. Of course, as noted earlier, this does not mean that the reliability of the distributed scheme is always higher, only that it lends itself to the derivation of tighter bounds. To complete our demonstration of potential reliability inversion in a practical setting, we need to show that, under some reasonable assumptions, the centralized system may in fact have higher reliability, despite its poorer reliability lower bound.

Consider modeling the centralized switches in greater detail, rather than lumping all switching hardware together into a hard core modeled by (2). A lot of extra work would be required in this case as switch failures and failure interactions depend on both the switch architecture and implementation technology. Let us assume an implementation technology for which a switch can be assumed not to fail unless we attempt to change its state. If, on average, only six operational switches are needed for correct reconfiguration (with high probability, reconfiguration entails bypassing

a single PE), then the pertinent system reliability equation is

$$\text{System reliability} = e^{-(6\sigma)t} \times R_{34\text{-out-of-36}}(r). \quad (7)$$

Equation (7) does not yield a reliability lower bound, so it cannot be used for system comparisons with certifiable outcome. However, it suffices for the purpose of demonstrating that centralized switching can have higher reliability than the distributed scheme under certain conditions. A plot of (7) is shown as the dotted gray line in Figure 5. We see that the dotted line (possibly) goes below the heavy black line beginning at $\lambda t = 10^{-2}$. We can verify that this is indeed the case by looking a few data points based on (3) and (7) (Table 1).

To provide an intuitive feel for our conclusions, we note that the reliability bound for centralized reconfiguration is not tight because we had to proceed with the highly pessimistic assumption that the entire switching network forms a critical core. We had no choice here as which switches will need to be reprogrammed for a particular pattern of PE malfunctions is unknown. In the distributed scheme, on the other hand, switches are integrated into the PEs; thus, as long as 34 of the 36 PE-switch modules are functional, we can successfully reconfigure the system. We do not care about the health of the switching mechanism any more than

we care about PE health. The system has no single point of failure.


Even though we considered only a relatively small example, the difference between reliabilities of the centralized and distributed schemes only grows as we enlarge the array. So, the results do scale up to very large PE arrays of practical interest. As mentioned previously, larger arrays will show greater benefits for distributed reconfiguration in terms of the differences between the lower bounds. They will also amplify the fairly small inversion appearing in Table 1.

In this article, we have tried to raise awareness of the notions of reliability inversion and modelability, using a concrete example for experimental validation of the abstract ideas. Even though work on reconfiguration schemes and algorithms for degradable processor arrays has continued unabated since papers previously cited,¹²⁻¹⁵ such variations, extensions, and improvements do not affect the formulation of reliability inversion. Design and reliability modeling considerations for reconfigurable 2D processor arrays with centralized and distributed switching will be taken up in a companion article.¹⁶

Besides modelability benefits, distributed reconfiguration of 2D processor arrays also leads to a more regular and modular design, hence providing greater packageability as well as suitability for realization as very large scale integrated circuits. This is an important side benefit that is similar to those cited for design for analyzability.²

The perils discussed here in connection with reliability inversion can be summed up by the following maxim: A benefit that is not observable to us,

TABLE 1. Reliability inversion data points.				
λt	0.010	0.020	0.050	0.100
Equation (3)	0.994	0.964	0.735	0.308
Equation (7)	0.994	0.965	0.742	0.319

because models don't show it, is no benefit at all. 

REFERENCES

1. E. Zio and N. Pedroni, "Methods for representing uncertainty: A literature review," French Foundation for an Industrial Safety Culture, Toulouse, Mar. 2013. Accessed on: Apr. 15, 2020. [Online]. Available: <https://www.foncsi.org/fr/publications/cahiers-securite-industrielle/literature-review-uncertainty-representation/CSI-uncertainty-representation.pdf>
2. R. Suri and M. Shimizu, "Design for analysis: A new strategy to improve the design process," *Res. Eng. Design*, vol. 1, no. 2, pp. 105–120, 1989. doi: 10.1007/BF01580204.
3. P. H. Dehkordi and D. W. Bouldin, "Design for packageability—Early consideration of packaging from a VLSI designer's viewpoint," *Computer* vol. 26, no. 4, pp. 76–81, Apr. 1993. doi: 10.1109/2.206519.
4. D. Locke, L. Sha, R. Rajikumar, J. Lehoczky, and G. Burns, "Priority inversion and its control: An experimental investigation," *ACM SIGADA Ada Lett.*, vol. VIII, no. 7, pp. 39–42, 1988. doi: 10.1145/59368.59374.
5. G. Reeves, "What really happened on Mars," *Risks Dig.*, vol. 19, no. 54, 1998. Accessed on: Apr. 15, 2020. [Online]. Available: https://www.cs.unc.edu/~anderson/teach/comp790/papers/mars_pathfinder_long_version.html
6. S. Dekker, *Drift Into Failure: From Hunting Broken Components to Understanding Complex Systems*. Boca Raton, FL: CRC, 2011.
7. M. Chean and J. A. B. Fortes, "A taxonomy of reconfiguration techniques for fault-tolerant processor arrays," *Computer* vol. 23, no. 1, pp. 55–69, Jan. 1990. doi: 10.1109/2.48799.
8. M. Sami and R. Stefanelli, "Reconfigurable architectures for VLSI processing arrays," *Proc. IEEE*, vol. 74, no. 5, pp. 712–722, 1986. doi: 10.1109/PROC.1986.13533.
9. Y. Ben-Asher, D. Peleg, R. Ramaswami, and A. Schuster, "The power of reconfiguration," in *Proc. Int. Colloquium Automata, Languages, and Programming*, 1991, pp. 139–150. doi: 10.1007/3-540-54233-7_130.
10. M. Fukushi, and S. Horiguchi, "Reconfiguration algorithm for degradable processor arrays based on row and column rerouting," in *Proc. 19th IEEE Int. Symp. Defect and Fault Tolerance in VLSI Systems*, 2004, pp. 496–504. doi: 10.1109/DFTVS.2004.1347875.
11. B. Parhami, "Dependable computing: A multilevel approach," Nov. 19, 2019. Accessed on: Apr. 15, 2020. [Online]. Available: https://www.ece.ucsb.edu/~parhami/text_dep_comp.htm
12. G. Jiang, J. Wu, and J. J. Sun, "Efficient reconfiguration algorithms for communication-aware three-dimensional processor arrays," *Parallel Comput.*, vol. 39, no. 9, pp. 490–503, 2013. doi: 10.1016/j.parco.2013.04.005.
13. J. Wu, T. Srikanthan, G. Jiang, and K. K. Wang, "Constructing sub-arrays with short interconnects from degradable VLSI arrays," *IEEE Trans. Parallel Distrib. Syst.*, vol. 25, no. 4, pp. 929–938, 2013. doi: 10.1109/TPDS.2013.114.
14. J. Wu, N. Liu, S.-K. Lam, and G. Jiang, "Shortest partial path first algorithm for reconfigurable processor array with faults," in *Proc. IEEE Trustcom/BigDataSE/ISPA*, 2016, pp. 1198–1203. doi: 10.1109/TrustCom.2016.0194.
15. J. Qian, W. Cao, J. Hu, J. Zhang, Z. Xu, and Z. Z. Zhou, "Satisfiability-based method for reconfiguring power efficient VLSI array," *IEICE Electron. Express*, vol. 13, no. 23, pp. 1–11, 2016. doi: 10.1587/elex.13.20160930.
16. B. Parhami, "Reliability and modularity advantages of distributed switching for reconfigurable 2D processor arrays," unpublished.

ABOUT THE AUTHOR

BEHROOZ PARHAMI is a professor of electrical and computer engineering at the University of California, Santa Barbara. His research interests include computer arithmetic, parallel processing, and dependable computing. He has served on the editorial boards of several journals, including *IEEE Transactions on Computers* and *IEEE Transactions on Sustainable Computing*. He is a Life Fellow of the IEEE. Contact him at parhami@ece.ucsb.edu.



IEEE COMPUTER SOCIETY
DIGITAL LIBRARY

Access all your IEEE Computer Society subscriptions at
computer.org
/mysubscriptions



An Enterprise Transformation Guide for the Inevitable Blockchain Disruption

Mehmet Demir, Ozgur Turetken, and Atefeh Mashatan, Ryerson University

Blockchain technology presents significant potential along with risk. The blockchain technology transformation framework informs decision makers on how blockchain fits in their processes, what data will be in the transactions, and who the participants will be, allowing users to lower risks and increase their chances for a successful blockchain solution.

Blockchain technology offers great potential to both revolutionize and disrupt businesses. Industry is taking notice of this potential, as evidenced by numerous bootcamps, courses, and seminars about this technology. However, some executives have been caught in a position where they are informed about the concept but not equipped with a set of critical questions to ask to leverage the potential of blockchains. Technology professionals are knowledgeable, yet not enough substantial business problems have been solved with blockchains. At times, unorganized effort is spent on practice projects in sandbox

environments, with little to learn from the cumulative experience of the community. Instead, what is available is a collection of stories about projects with no convincing evidence of their business benefits. This situation, combined with the concerns about the technology due to issues of privacy, security, performance, and capacity, makes it imperative to organize the thinking on blockchain-based innovation.

We believe a good start is to identify a set of critical questions to decide whether and how a blockchain-based solution could work for a particular organization. Reluctance to adopt disruptive technologies may be a significant competitive disadvantage for an organization, whereas proactive planning can be a significant advantage. Understanding where and how blockchain

technology will disrupt existing processes is beneficial.

We propose a framework through which enterprises can determine if and how they can transform their business processes to be supported by blockchain technology. We also provide key questions to provide insight into how blockchain technology might be helpful.

New blockchain-based business models should benefit all involved stakeholders. Increased involvement generally enhances the reliability and resistance of these systems. Marketing this paradigm to classically trained individuals is a managerial challenge.

Due to the nascency of the technology, widely accepted industry standards do not yet exist, and organizations are defining their own access rights, data structures, and allowable transactions.¹ This lack of standards is another managerial challenge, which blockchain technology transformation framework (BTTF), can help alleviate. BTTF provides a guideline for standard-defining activities to help organizations form a complete set of definitions in their blockchain solution. By following BTTF, executives can also find out whether blockchain is the right solution for their business challenges. A well-designed blockchain solution based on BTTF increases understandability for stakeholders and demonstrates business benefits to decision makers, limiting speculations.

There has been a great deal of work done in both academia and industry to enable and improve blockchain solutions. With the number of blockchain projects in the industry doubling every year, industry players have formed 108 consortia to collaborate on solutions for their respective industries.²

In this article, we formalize the best practices that practitioners are using across different industries. While we guide readers through the initial process of deciding whether or not blockchain technology is a viable solution to their problem, this exercise must be followed by other considerations that are more specific to each scenario in blockchain technology adoption.

BLOCKCHAIN TECHNOLOGY AS A DISRUPTION VEHICLE

Business impacts

Without a compelling reason, businesses would not just switch to blockchain technology. To evaluate potential benefits, the following features of blockchain implementations need to be analyzed for their impacts on a particular organization and specific business scenario, as some of these features may be positive for some organizations and negative for others. For example, transparency may concern stakeholders of a certain organization due to its impact on privacy and liabilities, whereas another organization may consider it an asset.

Auditability and traceability. Auditing is essential and very manually intensive. In the absence of trust, auditors spend considerable time and resources to cross-check the validity of data. Blockchains solve this problem by keeping the complete history of transactions and providing traceability guaranteed cryptographic methods. An auditor can easily verify the veracity of transactions based on the events on the blockchain.

Transparency. Having the state and outcome of a business process be transparent to the stakeholders increases their trust in the system and improves service experience. It assures

all participants of the integrity of the system and the processes. Blockchains can deliver this when the transactions are occurring on a network open to all participants. The value proposition is at its highest when it brings transparency to lengthy processes such as supply chain management.

Trust. When processes involve applications owned by different parties, disputes arise over what exactly has caused an incident to occur. When parties rely on their own copy of the records, reconciliation becomes a major part of a resolution. Blockchain technology can enable the participants to have the same copy of the records, leading to a quick and cost-effective resolution with higher confidence.

Permanency. Information is power, and there may be intentions to not share it or only share what supports a specific cause. In business-to-business communication, omission can be used for the purpose of hiding mistakes or failures. Communication platforms migrated to blockchains have the advantage of maintaining the original truth through this tamper-evident mode of communication. Blockchains enforce the availability, integrity, and permanency of the complete truth.

Eliminating system dependencies and intermediaries. Blockchains can remove the need for a separately maintained book of record, a central authority, or an intermediary through its decentralized architecture, which also removes the risk of a single point of failure. New blockchain-based systems can effectively complete transactions such as cross-border money transfer in minutes without any intermediaries.

Event-driven automation. Smart contracts have made event-driven automation possible. Coupled with the trust provided by blockchain technology,

smart contracts can simplify complex business processes by alleviating the need for manual interventions without compromising the integrity or quality of the overall process.

Blockchain-enabled features

Below is a list of features that blockchain technology helps to improve. When one considers a benefit such as transparency, he or she should question whether it would add value, eliminate a weakness, provide an advantage, or preclude a threat from competitors. An alternative approach would be investigating whether the corresponding question is a common question in the business process.

- › *Process tracking—Who does what?* Blockchains are very good at recording business events and communicating those to all participants. Such event communication and persistence make blockchains ideal for process tracking.
- › *Sensitive records—Who can access what?* Sensitive records can be protected with cryptography. The ownership of records can be transferred on digitally signed transactions. Encryption can protect the necessary authorization tokens while access and permissions can be traced and audited.
- › *Identity management—Who is who?* Trading partners can share identity-related information on blockchains, for example, verifying the information about a customer and placing the customer's public credentials on the blockchain with a flag indicating that this is a verified customer.
- › *Digital asset ownership—Who has what?* Cryptocurrencies showed

that ownership transfer can securely occur on blockchains without an intermediary.

- › *Voting—Who approves what?* Voting is very similar to digital assets from an ownership perspective. The ownership of the vote, that is, the ability to send or assign the vote, would be given to the user at the beginning of the process. Businesses can model complex processes with smart contracts combined with voting.
- › *Product traceability—Where is what?* Tracing the order, transportation, and subsequent delivery of the products in a supply chain can be handled on blockchains. Blockchains would inform partners of the events, and the status would be shared on the ledger. Order, payment, transportation, and delivery events can be managed by smart contracts.
- › *Intermediary and settlement agencies—Why the middleman?* When there is a distributed ledger and all participants trust the accuracy of the data, a middleman is not needed.

Challenges ahead of mainstream implementations

Some challenges with the widespread implementation of blockchain technology are easier to resolve, while others may take considerable amount of time and coordination among industry stakeholders. Some of these roadblocks are of technical nature while others are business related.

Technology challenges: A promising technology at its infancy. Unlike many other technologies that were first developed and matured in academia,

blockchain technology has not gone through academic due diligence, which makes it susceptible to a variety of issues.

- › *Software issues:* Each active participant needs the blockchain network-specific software for issuing transactions with consensus. Such software is developed in open source platforms and encapsulates the rules of the network that may change with maturity. There can be small changes updating some of the rules slowly,³ material changes where the network should be upgraded to a new version,⁴ or even emergency changes⁵ to prevent a high-risk issue. To publish the updated software and manipulate the behavior of the blockchain network, there are two well-known choices: soft and hard forks.
- › *Technical integration challenges:* Introducing blockchain technology in an established enterprise requires adopters and connectors between legacy systems and the blockchain. The architectural differences may make the integration near impossible.⁶ Blockchain adoption could mean a major revamp or a total development from scratch due to incompatibility.
- › *Scalability and performance:* Due to the decentralized architecture and consensus mechanisms, transaction verification takes some time on a blockchain. This can be easily tolerated in many cases, such as a supply chain, while it may be a roadblock in others, such as stock trading.⁷
- › *Cybersecurity:* There are several ways blockchains are secured.

Cryptographic methods secure the interactions by preventing forgery of blocks or preventing nodes trying to tip the consensus. The system is strong and solid as a whole but is vulnerable at its nodes. If participants do not have adequate security at their ends, blockchains are open to malicious activity through impersonated clients. If hackers access the private key of a participant, they can issue bogus transactions. The anonymity provided by the blockchain empowers hackers in this case. For example, BitCoin had reputational problems when one of the exchanges got hacked and bitcoins were stolen.⁸ This exchange went bankrupt and public trust toward blockchains got a hit.

Business challenges. The nascent nature of blockchain technology will be more concerning to executives who look at technology merely as a business enabler.

- › *Talent shortage:* Blockchain-focused technical skills are not yet taught in standard higher education curricula; therefore, the industry does not yet have a sizable pool of experts who can implement robust blockchain implementations. As a result, besides cryptocurrencies, blockchain instances are mostly proof-of-concept implementations with only 5–10% moving to production.⁹
- › *Cost-benefit analysis:* Most of the early adopters of blockchain were thinking about participation in blockchain ecosystems as an investment rather than as a traditional business use with an expected return on investment. The upfront cost of blockchain implementation is high; it includes new infrastructure and a capable team, so existing revenues can be negatively impacted. A big initial investment and loss of existing revenues are justifiable in the presence of sizable benefits. However, some costs or benefits are not easily measurable, hence making the adoption decision difficult. Unlike operational efficiency, it is not easy to assign a dollar value to trust or reputational risks.
- › *Governance:* The health and sustainability of business interactions are guaranteed through defined rules and responsibilities. An intermediary system can manage interactions and maintain service-level agreements. An authority can define rules and enforce accountability. In a decentralized architecture, however, we lose intermediaries and authorities and have to opt for decentralized governance in the form of consensus mechanisms or a regulatory body,¹ which does not define a single owner for the governing rules and can result in volatility and uncertainty.
- › *Uncertain regulatory status:* Laws tend to catch up slowly with new technology such as blockchains.¹ Current major players such as banks, insurance companies, government agencies, and law firms that are highly regulated are waiting for clear rules for widespread adoption; hence, there is considerable effort toward legislation. Most concerns are about illegal activities such as money laundering. For governments and revenue agencies, money flow and related tax implications are still a concern.
- › *Cultural adaption:* Business owners are used to solving their problems with systems by sharing minimal information and concentrating on divided responsibilities. In blockchains, sharing information makes it more secure. This change, which not only distributes power but also reduces the control of former authorities, would likely threaten some potential participants. Attracting participants is important for the success of the blockchain.¹ Trusting a system with a greater number of participants, rather than one with centralized authority, is a new concept which requires a culture change.
- › *Reluctance to change:* Fear of unknown technology and its possible shortcomings can cause concern. “If not broken, why fix it?” has been the motto of many. Meanwhile, resistance from third parties, such as trusted intermediaries who may lose their relevance, adds to the overall reluctance.

BLOCKCHAIN TECHNOLOGY TRANSFORMATION FRAMEWORK

Many blockchain research initiatives focus on applying blockchain technology to a specific scenario or industry. It is common to see use cases described for an industry and decide suitability with the end state, such as the final solution by following a flowchart.¹⁰ These frameworks can be more narrowly focused on current technologies and problems, instead of the transformation of business and discovery of opportunities.

TABLE 1. A list of framework questions.

Question	Action
Who? Participation	Identify independent collaborators in the process. Decide if anonymous participation is allowed. Decide who can approve or govern this process. Identify how participants can benefit from a trustable distributed ledger and transparency.
What? Tokenization	Identify digital assets used in the transactions. Find out how these assets are currently represented and stored by each participant. List the sensitivity of the attributes toward transparency. Find out the book of record process dependencies. Find out what information current intermediaries request and provide.
Where? Network and interaction	Identify how participants interact with each other and how this would change with peer-to-peer networking. Identify which interactions utilize which tokens. Identify each participant's role for each token.
Why? Trust injection	Identify current trust issues. Find out quality issues with current service. Decide how to provide trust (select from below) <ul style="list-style-type: none"> » Extended communication » Data sharing, process tracking » Tamper-resistant transaction history, logs, audit trails » Fraud prevention » Transparency and censor resistance.
When? Events and automation	Identify events in the system. Identify which events can trigger transactions that can be conducted automatically. Identify which interactions have a contractual nature.

This end-state focus also ignores which methodology is followed. For research purposes, focusing on a specific aspect of the problem is natural; but in industry, lack of a methodology can produce cookie-cutter applications, which might be unsuitable, or clones of what already exists, such as the creation of hundreds of digital coins after one or two successful ones. Unsuitable applications of blockchain technology will not provide the desired benefits to the user.

We propose a structured solution (transformation) framework for organizations to redesign their processes or

identify opportunities for using smart contracts. The introduction of a new trust model influences the number of collaborators. With the help of our framework, processes designed to communicate with a minimum number of external systems or partners can be redesigned to have many more collaboration partners.

BTFP presents five key questions to analyze the participation, tokenization, interaction, trust injection, and events/automation characteristics of the target business process. A detailed analysis of these characteristics reveals whether the business process is suitable for improvement with

blockchain technology. Each characteristic is analyzed with further questions. By answering questions in each area, organizations discover the suitability of blockchain technology for their processes.

There are two types of questions in this framework. The first type requires identification of one or more items. For these questions, the number of identified items is an indication of better suitability. For example, by answering the question of who, one identifies independent collaborators. The existence of several independent collaborators, the ability to add more, or the expectation of having more all increase the ability of a blockchain solution to improve the business process. On the other hand, if there is only one collaborator, or there is a cluster of collaborators all managed by one entity thus removing any independent decisions, a blockchain solution may not bring much value.

The second type of question focuses on decisions that enable the future direction or an existing constraint to become an input to the blockchain-based transformation. Having discussions to provide these decisions helps process owners understand the alternatives they have with blockchain technology and the consequences of using blockchains. The ability to have a clear decision shows the strong possibility of improvement, while not being able to decide indicates the possibility of future issues. For example, whether anonymous participation is allowed or not is necessary to decide the type of blockchain. The ability to decide on these items indicates a clear direction. If there are challenges to make such decisions, this could be an indication of the problem domain being too large for a single solution.

Table 1 shows the five key questions in BTFP. To understand the suitability of a

potential blockchain solution, analysis in all these five areas is necessary. Below are the descriptions of each question and its analysis process to guide process owners using BTTF. We start with understanding who (participants), continue with what (tokenization of assets and information), and then where, which reveal the details of the interaction network to understand how who and what are interacting. Why is a question to discover the issues to solve and the benefits to gain. When helps to understand events in the system that helps us to use blockchain technology with its smart contracts and automation tools.

Who? Participants

The redesign process starts with the analysis of existing actors to identify the participants involved in the process. Introducing blockchains will revolutionize the communication, interaction, and collaboration between these participants. Participants in the old process may have new roles in the new process. Depending on the overall business goals, there may also be new participants. Existing participants can remain only if they are independent collaborators in the network. For the participants in the new blockchain-led design, the next step is to decide whether every participant in the process can approve and govern. A higher number of participants justifies the use of blockchains.

What? Tokenization

What goes into an entry in the ledger (a token) is fundamental to the usage and benefit of the blockchain. Among the most common types of tokens are digital assets. Therefore, tokenization should start with identifying digital assets with attributes such as ownership and identifiers. If tokens do not emerge as a result

of this analysis, the next step can be to find out whether there are entities in the process that multiple systems are interested in. A token can be created from such an entity. If the process benefits from all transactions related to this entity being on the distributed ledger, it can be marked as a token. If there are existing books of record systems or intermediaries, they can be excluded in favor of similar functionalities over the blockchain. Analyzing the request and response structure may reveal the detail of the peer-to-peer communication over the intermediary, and this communication structure can be used to define new tokens.

Where? Interaction network

To operate on the peer-to-peer distributed network structure, each participant needs to be able to connect with several others. An important design target is to eliminate dependencies on a specific group of nodes and remove any single point of failure.

Why? Trust injection

The most valuable feature of blockchains is the trust provided to normally untrusting participants. At this point in the process design, all previous findings should be validated considering trust requirements. Existing trust issues should be listed and prioritized. If a process with the identified participants, tokens, and interactions requires trust, the use of blockchains would be justified. Each trust requirement should be matched with a particular blockchain feature.

When? Automation events

This step reveals the events that can be detected in the redesigned process for previously identified participants, tokens, and interactions. For each event, actions would be identified. If an action would

automatically trigger a transaction, smart contracts are relevant. Smart contracts would initiate new transactions when predefined events are realized in blockchains. Many legacy processes do not have an event-based approach to automated transaction execution. Therefore, identifying automated transaction sources can be an extended discovery effort. Automation may lead to cost savings. The identification of these savings is important, as it helps to justify the new blockchain implementation.

USE CASE 1: SUPPLY CHAIN, GLOBAL TRADE

Most international supply chains are difficult to track. Products and goods change several hands as they pass from manufacturers to consumers. Building a foundation of trust is hard considering the variety of trading partners. The current need for such trust is mostly fulfilled by intermediaries and legal contracts. The additional costs of acquiring trust and process traceability are very significant. For example, documentation and follow-up costs for a container shipment are more than double the cost of the physical shipment.¹¹

In the supply chain industry, there are existing blockchain solutions for the food supply chain,¹² mining,¹³ and diamond tracking,¹⁴ among others. We present a generic solution to a simplified use case of an international supply chain process to demonstrate the concept, steps, and value of BTTF in this context.

The analysis in Table 2 shows that the target supply chain use case is a good candidate for improvement with blockchain technology. There are plenty of independent collaborators. Participants have motives and benefits from the implementation. There

TABLE 2. Framework responses: supply chain, global trade.

Who?	What?	Where?	Why?	When?
<p>Independent collaborators</p> <ul style="list-style-type: none"> » Factories » Land transportation providers » Freight forwarders » Custom brokers » Governments » Ports » Ocean carriers » Insurance » Retail businesses <p>Are anonymous participants allowed?</p> <ul style="list-style-type: none"> » No. This process requires participants to have identities and permissions. <p>Can any participant approve or govern the steps of this process?</p> <ul style="list-style-type: none"> » Each participant has a major role in the governance of this process. <p>What are the DLT benefits?</p> <ul style="list-style-type: none"> » Each party can have a clear view of the process details and see incoming shipments and provide his or her process details downstream. 	<p>Digital assets used in the transactions</p> <ul style="list-style-type: none"> » Shipment » Export certificates » Container <p>How is this information currently represented and stored by each participant?</p> <ul style="list-style-type: none"> » Product details » Shipment status (OK, lost, damaged) » Ownership of the shipment » Documentation » Approvals <p>What are the book of record process dependencies?</p> <ul style="list-style-type: none"> » Factory is the book of record on the content. » Land transportation providers and ocean carriers are the book of record on location and destination. » Ports are the book of record for departure, arrival. <p>What is the information current intermediaries request and provide?</p> <ul style="list-style-type: none"> » Freight forwarders coordinate the movement of goods to their destination and handle the necessary paperwork. » They would request product and destination information and provide an estimate of arrival. 	<p>How do participants interact with each other?</p> <ul style="list-style-type: none"> » Participants currently interact with several media of communication including online, email, fax, and paper. There are several business-to-business custom integrations. <p>How do the interactions change with peer-to-peer networking?</p> <ul style="list-style-type: none"> » There would be great transparency if every stakeholder can access others and receive information from all. They can create more successful plans with more information <p>Which interactions need which tokens?</p> <ul style="list-style-type: none"> » Factories, land transportation providers, customs brokers, governments, and retail businesses need product information token. » Customs brokers, ports, ocean carriers, and retail businesses need container token. <p>What are each participant's roles?</p> <ul style="list-style-type: none"> » Since there are several equal contributors, the roles can be distributed uniformly. Participants can all form nodes to create new blocks. Due to the business volume, they all have a stake in the health of this blockchain. Privacy concerns between competing businesses should be handled at the token level to hide details that should not be shared. 	<p>Current trust issues</p> <ul style="list-style-type: none"> » Participants can hide the issues and defer responsibilities due to insufficient, imprecise, corruptible, forgettable, and not provable information. <p>What are the quality issues with current service?</p> <ul style="list-style-type: none"> » It is not clear where the shipment is, why it is late, or whose mistake delayed the arrival. <p>How is trust provided?</p> <ul style="list-style-type: none"> » Extended communication » Data sharing » Process tracking » Tamper-resistant transaction history » Audit trails » Fraud prevention » Transparency 	<p>Events in the system</p> <ul style="list-style-type: none"> » Several handover events where a participant delivers the item, and another receives it » Government approval » Customs clearance » Loss and damage <p>Which events can trigger transactions that can be handled automatically?</p> <ul style="list-style-type: none"> » Several payment and acknowledgment transactions can be automated with the delivery events, for example, custom duties to be paid when the item is in the port. <p>Which of these interactions are contractual in nature?</p> <ul style="list-style-type: none"> » From factory to the retail store, many mini transactions and payments can be coded in smart contracts and executed by sensor events, for example, 30% of the payment to be paid to the factory at the time the shipment leaves the factory and 20% to be paid when shipment is in the ocean carrier.

are several well-defined tokens present in the process. There are numerous ways that the collaborators will benefit from the new token and interaction models. The current trust and quality issues are well listed. Almost all possible ways of injecting blockchain-related trust into the new process model are confirmed. Several smart contract opportunities including a partial payment automation are identified. Our framework has been followed well in the above example, and the process is a good candidate for improvement through blockchain technology.

USE CASE 2: REAL ESTATE SALE PROCESS

Multiparty agreements, such as a real estate sale process, require information to be shared between the seller, the buyer, their lawyers, their banks, their spouses, insurance companies, the power utility, the gas company, city utilities, land registry, and government revenue taxation agencies, which are traditionally done by sharing information between two parties at a time. Smart contracts can execute the sale, transfer responsibilities, change the ownership, and transfer the money. Such a system under the close monitoring of so many stakeholders would be more trustworthy than one where each stakeholder keeps his/her own records with partial information.

The analysis in Table 3 shows that the target real estate use case is a good candidate for improvement with blockchain technology as well. There are plenty of independent collaborators. Participants benefit from the implementation. Most have clear duties and responsibilities tied to the success of the collaborated process. There are several well-defined tokens present. Ownership-related information is a good


token. With the old and new interactions, there are many ways that collaborators will benefit from the new token and interaction models. Currently, there is established trust in the system, which is based on the parties' past experience. Execution seems orderly, but transparency is limited, and operational redundancy is very high. Almost all possible blockchain-related trust injection is confirmed to inject trust and efficiency into the new process model. The majority of transactions can be automated with smart contracts. Our framework has been followed carefully, and the process is a very good candidate to be improved by the application of blockchain technology.

Applying blockchain technology without a multidimensional assessment of the business process may result in an unnatural application of blockchain that does not provide the desired benefits. To prevent this problem, we introduced a prescriptive approach for transforming business processes. BTTF is a structured way of assessing whether business processes can be improved with blockchain technology.

BTTF applies to any existing or new business process. Besides our use case examples in supply chain and real estate, other industries such as finance, government, insurance, and energy are well-known application areas that can benefit from applying BTTF. Employing BTTF for more sensitive business processes, such as those in health care, would reveal the critical compatibility issues between the process and blockchain technology.

A limitation of BTTF is the manual nature of the analysis. Our research will continue on this topic, and we will

develop a tool to automate the planning and execution of BTTF-based analysis. This tool will help in understanding the details of the analysis questions, evaluating the answers, informing users about the impact of their choices, identifying possible conflicts, generating ideas on the opportunities, as well as comparing the analysis of different processes. The comparison ability can also improve our framework with the possibility of an empirical assessment of the framework.

BTTF is a guide for a successful beginning to the thought process toward blockchain adoption. There are many more considerations that may be necessary as adopters progress in their implementations, such as joining or forming new industry partnerships and consortia. 

REFERENCES

1. M. C. Lacity, "Addressing key challenges to making enterprise blockchain applications a reality," *MIS Quart. Exec.*, vol. 17, no. 3, pp. 201-222, 2018.
2. "Solutions for blockchain consortia," ESG Intelligence, New Delhi, India. Accessed on: Oct. 23, 2019. [Online]. Available: <https://esg-intelligence.com/contact/>
3. P. Traugott, "Ethereum (ETH) working on a lot of small updates while we wait for the big ones to be announced," Captain Altcoin, Bosnia and Herzegovina, July 29, 2018. [Online]. Available: <https://captainaltcoin.com/ethereum-eth-working-on-a-lot-of-small-updates-while-we-wait-for-the-big-ones-to-be-announced/>
4. H. Partz, "Privacy altcoin Zcash announces first network update, 'not expected' to be a fork," *Cointelegraph*, Mar. 3, 2018. [Online]. Available: <https://cointelegraph.com/news/>

TABLE 3. Framework responses: real estate sale process.

Who?	What?	Where?	Why?	When?
<p>Independent collaborators</p> <ul style="list-style-type: none"> » The seller, the buyer, their lawyers, their banks, their spouses, insurance companies, the power utility, the gas company, city utilities, land registry, and government revenue taxation agencies 	<p>Digital assets used in the transactions</p> <ul style="list-style-type: none"> » Real estate ownership, insurance, power and gas contracts, city utility services, mortgage application, and mortgage <p>How is this information currently represented and stored by each participant?</p> <ul style="list-style-type: none"> » The buyer and seller have paper documents. » The lawyer has paper and scanned documents. » The bank has mortgage agreement. » The insurance company has insurance contract details. » The power and gas company has service contract details. » The land registry has the title information. 	<p>How do participants interact with each other?</p> <ul style="list-style-type: none"> » Currently, there is one-to-one interaction between participants. There is process-based centralization around lawyers. The buyer interacts with her lawyer, and her lawyer interacts with most other process stakeholders. <p>How do the interactions change with peer-to-peer networking?</p> <ul style="list-style-type: none"> » There would be great transparency if every stakeholder can access others and receive information from all. As the needed information is already shared, the number of interactions between the participants would decrease significantly. <p>Which interactions need which tokens?</p> <ul style="list-style-type: none"> » Lawyer—land registry—real estate ownership » Lawyer—insurance—title insurance contract » Buyer—insurance—home insurance contract » Buyer—power and gas companies—power and gas contract » Lawyer—city utility services—city utility services contract » Buyer—bank—mortgage » Bank—lawyer—mortgage <p>What are each participant's roles?</p> <ul style="list-style-type: none"> » Mortgage: buyer signs, bank provides funds, lawyer attaches the fund to closing agreement 	<p>Current trust issues</p> <ul style="list-style-type: none"> » The buyer does not see the process state. His or her relation is solely based on trust in the lawyer and his or her perception of the reliability of the process. » The seller does not know when the funds would be deposited. » The buyer is not sure if the ownership is transferred. » The buyer is not sure if the utility contract is in effect. » The bank is not sure if the insurance is valid and not cancelled. <p>What are the equality issues with the current service?</p> <ul style="list-style-type: none"> » Why is there a gap in information held by different public agencies and the banks? » Why do buyers need to provide the same information again and again? <p>How is trust provided?</p> <ul style="list-style-type: none"> » Extended communication: The buyer and seller can receive the extended communication events and know about the status. Income tax authorities are notified from the sale immediately. » A tamper-resistant transaction history: All steps would be on the blockchain so there is no dispute. For example, there is no dispute that the city utility bill starts from the closing day. There is no dispute that the power bill starts from the closing day. » Tamper-evident logs and audit trails: In case of a dispute, the events that happened are in the blockchain. » Fraud prevention: Every detail is shared on the blockchain about the sale. For instance, a lawyer trying to commit fraud would be obvious as all steps are on the blockchain. 	<p>Events in the system</p> <ul style="list-style-type: none"> » Sale agreement » Closing » Mortgage funding <p>Which events can trigger transactions that can be handled automatically?</p> <ul style="list-style-type: none"> » Closing event can trigger several transactions: land transfer, utility, gas, and power bills activation, insurance starts, money transfers, and so forth. <p>Which of these interactions are contractual in nature?</p> <ul style="list-style-type: none"> » Closing, land transfer, utility, power, gas, and several of these interactions are contracts. » Smart contracts can execute the sale, transfer responsibilities, and change the ownership, and transfer the money.
<p>Are anonymous participants allowed?</p> <ul style="list-style-type: none"> » No. This process requires participants to have identities and permissions. Another alternative is to have the incorporated entities as identified while individuals remain anonymous. Trusted entities such as lawyers can introduce the anonymous entities to the system to enable their anonymity, while validating they are real. 	<p>What are the book of record process dependencies?</p> <ul style="list-style-type: none"> » Land registry is the book of record for title. <p>What information do current intermediaries request and provide?</p> <ul style="list-style-type: none"> » There are not many intermediaries in this process. » The main process issue is the number of interactions between entities and risks associated with it. To understand the entities in the process, we need to look at every one-to-one process. 	<p>What information do current intermediaries request and provide?</p> <ul style="list-style-type: none"> » There are not many intermediaries in this process. » The main process issue is the number of interactions between entities and risks associated with it. To understand the entities in the process, we need to look at every one-to-one process. 	<p>What are the DLT benefits?</p> <ul style="list-style-type: none"> » Sellers and buyers can benefit with the ability to access to a repository to observe the state of the process. Lawyers can communicate the details, manage the signatures, and close the deal transparently. 	

privacy-altcoin-zcash-announces-first-network-update-not-expected-to-be-a-fork

5. Cointelegraph, "Ethereum's Istanbul hard fork." Accessed on: Apr. 19, 2020. [Online]. Available: <https://magazine.cointelegraph.com/ethereum-hard-fork-istanbul-2019/>
6. "Is blockchain right for you? Bridging the legacy gap," Veriday, Mississauga, Canada, Nov. 23, 2017. [Online]. Available: <https://www.veriday.com/blog/blockchain-right-bridging-legacy-gap/>
7. B. Marr, "The 5 big problems with blockchain everyone should be aware of," *Forbes*, Feb. 19, 2018. [Online]. Available: <https://www.forbes.com/sites/bernardmarr/2018/02/19/the-5-big-problems-with-blockchain-everyone-should-be-aware-of/#2b9e879e1670>
8. R. McMillan, "The inside story of Mt. Gox, Bitcoin's \$460 million disaster," *Wired*, Mar. 3, 2014. Accessed on: Feb. 9, 2018. [Online]. Available: <https://www.wired.com/2014/03/bitcoin-exchange/>
9. S. Gupta and T. Mondal, "HfS blueprint report: Enterprise blockchain services," HfS Research, Bengaluru, India, Nov. 2017. [Online]. Available: <https://www.hfsresearch.com/blueprint-reports/hfs-blueprint-enterprise-blockchain-services>
10. A. B. Pedersen, M. Risius, and R. Beck, "A ten-step decision path to determine when to use blockchain technologies," *MIS Quart. Exec.*, vol. 18, no. 2, pp. 99–115, 2019. doi: 10.17705/2msqe.00010.
11. "Cross border supply chain solution," IBM, Armonk, NY. Accessed

ABOUT THE AUTHORS

MEHMET DEMIR is with the Computer Science Department at Ryerson University, Toronto. His main research interests focus on blockchain technology, Internet of Things, and artificial intelligence applications. Demir received an M.B.A. from the Schulich School of Business at York University, Toronto. He is a Member of the IEEE. Contact him at mehmet.demir@ryerson.ca.

OZGUR TURETKEN is a professor and associate dean of research at the Ted Rogers School of Management at Ryerson University, Toronto. His research focuses on applied text analytics and decision support. Turetken received a Ph.D. in management science and information systems from Oklahoma State University, Stillwater. He was an associate editor of *AIS Transactions on Human-Computer Interaction*, the European Conference on Information Systems, and the International Conference on Information Systems. He is a Member of the IEEE and the AIS. Contact him at turetken@ryerson.ca

ATEFEH MASHATAN is a professor of information technology management at the Ted Rogers School of Management and the founder and director of the Cybersecurity Research Lab at Ryerson University, Toronto. Her current research focus is on solving industry problems through cutting edge information security and cyber risk mitigation strategies and solutions. Mashatan received a Ph.D. in combinatorics and optimization from the University of Waterloo, Ontario, Canada. She was recognized by *SC Magazine* in 2019 as one of the top five Women of Influence in Security globally. She is a Member of the IEEE. Contact her at amashatan@ryerson.ca.

12. D. McQuade, "Blockchain-traced seafood: Helping historic New England fisheries thrive," IBM, Armonk, NY, Oct. 17, 2019. [Online]. Available: <https://www.ibm.com/blockchain/use-cases/>
13. P. Rizzo, "World's largest mining company to use blockchain for supply chain management," *Supply Chain News*, Oct. 2, 2016. [Online]. Available: http://www.supplychain247.com/article/worlds_largest_mining_company_to_use_blockchain_for_supply_chain
14. B. Marr, "How blockchain could end the trade in blood diamonds: An incredible use case everyone should read," *Forbes*, Mar. 14, 2018. [Online]. Available: <https://www.forbes.com/sites/bernardmarr/2018/03/14/how-blockchain-could-end-the-trade-in-blood-diamonds-an-incredible-use-case-everyone-should-read/#38347884387d>



Is Privacy Regulation Slowing Down Research on Pervasive Computing?

Claudio Bettini, Università degli Studi di Milano

Salil Kanhere, University of New South Wales

Marc Langheinrich, Università della Svizzera Italiana

Archan Misra, Singapore Management University

Delphine Reinhardt, University of Göttingen

Privacy legislation has often been identified as a roadblock for advanced context-aware applications. The feedback collected from more than 150 researchers in pervasive computing reveals a different attitude. Has pervasive computing's privacy challenge been solved?

Reading about large-scale data privacy violations has become commonplace. The year 2018 alone saw almost 1 billion user accounts involuntarily disclosed or hacked, including customers of restaurant chains, retailers, and hotels as well as major

online services (Facebook, Uber); in July 2019, more than 2 billion log entries from the Internet of Things (IoT) management platform Orvibo were stolen, containing user accounts, passwords, and even recorded smart camera conversations. Mobile and pervasive technology and services have a role in this scenario, since they introduce new devices and communication protocols (all with plenty of room for unforeseen vulnerabilities), novel types of

personal data, and a new scale for the amount of data being collected.¹¹

Driven, in part, by these large-scale privacy violations and increasing public concern, regulators in most countries have taken action in revising legal requirements related to personal data protection. In 2016, the European Union (EU) approved its new General Data Protection Regulation (GDPR), which went into effect in May 2018; in the United States, California passed a new data privacy law in June 2018 that, in many ways, resembles GDPR [the so-called California Consumer Privacy Act of 2018 (CCPA)]; and Japan has taken similar steps with the amended Act on the Protection of Personal Information.

As a consequence, companies have begun to invest a significant amount of money and resources in ensuring the continuing compliance of their systems and products with the changing legal landscape, including adjusting the design, production, and test processes of their products. This will inevitably have an impact on the type of services that will be offered, their cost, and the timing of their appearance on the market. Although the value of investing in security is usually well understood, investing in privacy is often seen only as a cost, especially by small and medium-sized enterprises.

This seemingly unavoidable tradeoff (more privacy means fewer services), in principle, also applies to research in this space: although there are usually ample exceptions for research, legislation does have a significant impact once prototypes move into commercialization. Similarly, increased privacy awareness has also heightened the bar for getting ethical clearance, both at an institutional level and within the wider research community.

We are not the first to investigate how privacy affects pervasive research.

One of the first efforts in this space came from Langheinrich and Lahlou in 2003,⁹ who found a high level of nonconcern among researchers. Although both the legal landscape and research practices have significantly changed since then, Bednar et al.¹ found similar levels of disinterest more than 15 years later while interviewing six senior software engineers. In 2019, Spiekermann et al.¹⁶ surveyed 124 engineers and found that, although most considered privacy important, few enjoyed including it within their systems. These last two studies also found that many engineers struggle with the organizational environment: they face a lack of time and autonomy that are necessary for building ethical systems. A similar effort by Szekely¹⁷ focused on IT professionals, although their survey targeted surveillance issues and was limited to two countries, Hungary and The Netherlands, highlighting a higher level of awareness in the second country.

Our work explicitly focuses on researchers. Specifically, we wanted to understand how well the research community in pervasive computing understands current privacy legislation, in what way does it affect their work, and to what extent their attitudes toward ethical decisions in this space vary. Within the context of a privacy-focused panel at the IEEE 17th International Conference on Pervasive Computing and Communications (PerCom 2019), we conducted a brief survey among active researchers in this field to inform the panelists' discussions. This article summarizes their responses and comments on the implications.

THE SURVEY

Settings and Sample

Our online questionnaire of 10 questions (see "Survey Questions and Possible

Answers") was distributed through the PerCom conference mailing list. The survey took about 5 min to complete, and no rewards were provided to the participants. A total of 154 researchers on mobile and pervasive computing from both academic and industrial institutions responded. When asked about their personal attitudes regarding the (digital) sharing of personal data, 49% indicated belonging to the category "quite concerned" of those who "want to know exactly who gets [their] data and what they do with it."

The second most represented category was "quite liberal," with 28%, followed by "very concerned," with 18%. The remaining 5% of the participants indicated being "very liberal," that is, "enjoying sharing to a large audience including location, pictures, video." When the central two categories (quite concerned and quite liberal) were mapped to a pragmatist approach to privacy, these results fell roughly in line with the distribution of Westin's privacy categories (pragmatist, unconcerned, fundamentalist) found in prior surveys.⁸ It also means that the large majority of researchers is concerned about the protection of their personal data when acting as users of digital services.

To better situate these concerns within a concrete pervasive computing setting, we asked participants to select up to two types of sensor data (from an overlapping list of four examples) that they thought should not be collected or retained by municipal authorities in a smart city application:

1. indoor location in public spaces (for example, shopping centers)
2. outdoor (cellular) location data
3. location data via video analytics

SURVEY QUESTIONS AND POSSIBLE ANSWERS

1. What is your attitude about sharing your personal digital data?
 - Very liberal (I enjoy sharing to a large audience including location, pictures, video)
 - Quite liberal
 - Quite concerned (I want to know exactly who gets my data and what they do with it)
 - Very concerned (I protect my email and mobile number, I do not post on socials, I usually deny consent).

2. Do you think about privacy issues when you design a new mobile/pervasive solution?
 - Always
 - Yes, but only if it handles very sensitive information
 - No, only later after the prototype is ready and needs to be tested/deployed
 - Never.

3. How much do you know about the data protection legislation?
 - I regularly follow the updates to the data protection legislations
 - I know the basic principles of current data protection legislations
 - I know the basic principles but only for the one in my country
 - Very little or nothing.

4. Do you know where (countries) the EU GDPR applies?
 - Worldwide to all citizens
 - All of Europe
 - Only the EU
 - Wherever EU citizens are served.

5. Which of these types of context/personal sensor data do you think applications run by city/municipal authorities should *not* be authorized to collect or retain (due to privacy concerns)? (Pick at most two.)
 - Indoor location in public spaces
 - Outdoor cellular location data
 - Person identity (from video analytics) in public spaces
 - Publicly crawlable social media data
 - They can collect all of the mentioned data if they declare the purpose and I believe it is useful.

6. What do you think about the obligation of incorporating "privacy by design" in products/systems?
 - I agree; it saves later costs for compliance and leads to better products
 - I disagree; this requires special expertise and it is not always needed
 - I do not know what "privacy by design" is or what it involves.

7. Have you ever tried to include some privacy protection technique in your solutions?
 - Yes, anonymization
 - Yes, obfuscation or statistical perturbation (including differential privacy)
 - Yes, cryptography and/or blockchain-based techniques
 - Yes, more than one of the techniques mentioned in other answers
 - No (possibly only basic security measures like access control).

8. If you answered Yes to the previous question, have you ever encountered difficulties when including privacy protection techniques in your solutions?
 - None; it was always straightforward
 - It was difficult to translate and apply the legal requirements into my solution
 - It was difficult to choose among the different solutions to best protect privacy
 - It was difficult to ensure that the implemented solution was the most efficient one.

9. Did you experience obstacles to your work in pervasive computing due to the privacy regulation?
 - Yes, and I had to abandon my project
 - Yes, and I had to invest significant resources (time/money) for complying
 - Yes, but I easily found a solution to comply
 - No
 - No, but I suspect I might later on
 - Other (free text).

10. Do you think that a strict privacy regulation will have a positive effect on the adoption of pervasive systems?
 - Yes, it would provide better guarantees for privacy-concerned users
 - No
 - Not sure.

4. scraping of their publicly available social media data.

Alternatively, participants could select “all of the above.”

Although approximately 25% (one in four) of the participants were happy to have all of these data collected, 69% would not support video-based tracking. Interestingly enough, only 31% considered indoor location sharing to be problematic, even though the survey did not specify the tracking method [such as via Bluetooth Low Energy (BLE) tags or, in fact, video analytics]. Outdoor (cellular) location tracking saw similar levels of concern (33%), whereas sharing social media data had only 18% of participants concerned.

This example illustrates that the type of sensor plays an important role in people’s privacy concerns. Even though the end result of, say, location

tracking via a cellular signal may be identical to a location trace obtained using video (or depending on the density of cameras, video information could even be much more detailed), our participants reacted strongly to the idea of video-based tracking. Also, the fact that most participants did not worry about having their social media streams monitored seems to suggest that they are either not aware of the significant risks that such a practice entails (see, for example, Rosenblum¹⁴) or consider any and all public data as lost from their control.⁶

Regulation Awareness and Knowledge

We first asked our participants how much they estimate knowing about general data protection legislation (Figure 1). Overall, 14% of our respondents reported that they regularly follow

updates to data protection legislation. In addition, 42% indicated that they know the basic principles of current data protection legislation; an additional 35% stated that they know those basic principles only for the laws applying to their country. Combined, this still means that more than 91% of our respondents felt that they have a good grasp on current national privacy legislation—an astonishingly high value. Only 9% indicated knowing very little or nothing about privacy law. Obviously, such self-assessment is no proof of actual knowledge, and we did not rigorously test the actual legal understanding of our participants.

However, a follow-up question asked participants to indicate in which countries the EU GDPR would apply. Almost half of our respondents (47%) correctly answered that it applies wherever EU citizens are served (the marketplace rule, Article 3(2) of GDPR); 36%

incorrectly assumed that it would apply only within the EU, while 9% believed GDPR would apply for all countries in Europe. The remaining 8% believed that it would apply worldwide to all citizens. Although this is clearly a low bar, understanding the scope requirement is an important aspect of today's legal

privacy landscape, as laws increasingly use this approach. For example, California's CCPA similarly applies to data controllers that "do business in the State of California," irrespective of their physical establishment. The answers to this question seem to indicate that, in practice, there is still a gap

between legislative intent and practitioner awareness.

System Design Approach and Privacy by Design

Since our sample was composed of researchers working in the area of pervasive computing, we were interested



FIGURE 1. The participants' knowledge, experiences, and beliefs about privacy regulations.

in knowing how privacy legislation has affected them in their development of pervasive computing technologies and applications. A very positive result is that a large majority of our respondents either indicated that they “always” think about privacy issues when developing new solutions (35%) or that they do so if the system “would handle very sensitive information” (52%). However, 12% would not consider privacy in the design phase but only in a later phase when the prototype would need to be tested/deployed, and 1% indicated that they never think about privacy issues.

Given these numbers, it is not surprising that a majority of our sample (59%) agreed that the obligation imposed by GDPR of incorporating “privacy by design” in products or systems is a good thing, as it “saves costs later for compliance and leads to better products.” However, 27% of our participants indicated that they did not know what “privacy by design” meant or what it involved. The remaining 14% were against this obligation, indicating that it required “special expertise” and “would not always be needed.”

Applying Privacy Protection and Potential Obstacles

Moving from theory to the practice of privacy protection, we were interested in knowing which privacy protection techniques our participants already implemented in their systems. When presented with several key categories of privacy protection methods, a large group of participants (39%) reported using only basic security measures, such as access control. The remaining used the following techniques: anonymization (25%), cryptography and/or blockchain-based techniques (9%), or obfuscation or statistical

perturbation including differential privacy (6%). About one out of five participants (21%) declared using more than one of these techniques.

Among the participants who had already included privacy-preserving techniques in their solutions ($n = 97$, multiple choices possible), only 21% indicated that they did not encounter any difficulties. In contrast, 50%

previously abandoned a project due to privacy legislation.

The last question of the survey collected opinions in favor of or against strong privacy laws. A majority of our participants (51%) indicated that a strict privacy regulation should have a positive effect on the adoption of pervasive systems, providing better guarantees for privacy-concerned users and

WHEN PRESENTED WITH SEVERAL KEY CATEGORIES OF PRIVACY PROTECTION METHODS, A LARGE GROUP OF PARTICIPANTS (39%) REPORTED USING ONLY BASIC SECURITY MEASURES, SUCH AS ACCESS CONTROL.

reported that it was difficult to ensure that the implemented solution was the most efficient one. In addition, 33% found it difficult to translate and apply the legal requirements to their solution, while 30% had difficulties choosing among the different solutions to best protect privacy.

Impact of Regulation on Adoption

Regarding the more general issue of the impact of regulation on the design and deployment of pervasive solutions, 43% of our participants declared that they did not (yet) face obstacles due to privacy regulation. Some participants had to deal with privacy issues but easily found a solution to make their system legally compliant (20%). However, a high number of them (32%) had to invest significant resources in terms of time and money to comply with privacy legislation. A few participants (5%) indicated that they had

developers. A good number of participants disagreed (20%), while a third (29%) were not sure either way.

DISCUSSION

We see three key issues emerging from our survey: 1) the benefits of strong privacy legislation for research, 2) the lack of guidance for implementing privacy by design, and 3) fundamental challenges in today’s privacy regimes. We briefly discuss these in the following sections.

Benefits of Legislation

As our participants indicated, strict regulation can have tangible benefits, not only for users (who may more readily adopt a pervasive service) but also for service providers. In fact, having clear legal guidelines has been beneficial for several of the authors of this article. Although it has become commonplace to obtain ethical clearance

from an institutional review board prior to running a particular study, universities increasingly include their legal departments when it comes to authorizing field deployments. For example, when one of the authors (Archan Misra) attempted to deploy smart services (see, for example, Kandappu et al.⁷) across his university's campus, legal services were actually grateful for the concrete legal guidance offered by Singapore's 2012 privacy law (the 2012 Personal Data Protection Act).

Having concrete rules and practical guidelines will be essential to ensure that privacy laws reduce uncertainty for researchers rather than increase it (see also the "Privacy-by-Design Guidance" section). Similarly, if achieving legal compliance can be possible only if service quality is reduced (for example, by removing or limiting a system's personalization capabilities), users may not perceive any benefits of such legal protection and simply vote with their feet by using more complete but less privacy-friendly systems. It might be time to reconsider the principal approach to privacy legislation, which, in many jurisdictions, attempts by default to minimize the collection and procession of personal data (see also the "Fundamental Challenges" section).

Privacy-by-Design Guidance

Although fewer than a third of our participants indicated that they did not know what "privacy by design" meant, this still is a significant number, especially since our respondents can all be considered technology experts. This points to the obvious gap between GDPR's well-meant inclusion of this principle (Article 25: "Data Protection by Design and by Default") and the

lack of concrete technical guidance on how to implement this principle. Specific privacy-enhancing technologies for mobile and pervasive computing, often inspired by solutions in databases, have been investigated for almost two decades, leading to a very rich set of methods—recent surveys can be found in Bettini and Riboni,² Christin,³ and Gkoulalas-Divanis and Bettini.⁵

However, this wealth of methods does not mean that it is any easier to apply the right method in the right context. Researchers need more guidance on how to incorporate these technologies and procedures in their systems during design. Although several attempts have been made to formulate more concrete methodologies in this space (see, for instance, Chapter 5 in Langheinrich and Schaub¹²), the huge variety of pervasive computing systems render the idea of a simple how-to that could be followed in each and every project infeasible. Instead of refining and extending our vast array of methods for protecting personal data, we need more research into the practical application (that is, integration) of such techniques into pervasive systems.

Fundamental Challenges

About one-third of our respondents indicated they had to "invest significant resources for complying" or even had to abandon their project due to the challenges posed by privacy laws. Legal scholars have long since challenged the suitability of current privacy laws for today's technology landscape. Even though both GDPR and CCPA were created with social media firmly in mind, they nevertheless still trace their roots back to the privacy laws of the 1970s—when

data were still stored on punch cards. Not only was the amount of stored and processed data minuscule compared to today's volume (according to the World Economic Forum, the amount of digitally stored data will reach 44 zetabytes in 2020—that's 40 times more bytes than stars in the universe⁴), the key stakeholders were predominantly governments, and most of the captured data were still manually entered. Among the key criticisms are the predominant focus on personal rights versus overarching social benefits,¹³ the belief that data can actually be anonymized (and that it matters),¹⁰ and the fundamental limits of meaningful notice and choice.¹⁵

Veil,¹⁸ in particular, criticizes GDPR's one-size-fits-all approach, which may make it easy for lawmakers to craft legislation but which imposes many of the same obligations (Veil counts no fewer than 68) on both large international companies (irrespective of whether they operate a social media site or, say, manufacture escalators), a plumber, or a small church club. Although large legal frameworks that apply across many regulatory contexts are appealing, the vast differences in data collection motivations and corresponding privacy risks may require a much more individualized approach.

The main goal of our study was to investigate, discuss, and better understand the impact of the recent evolution of personal data protection legislation on the work of the research community on mobile and pervasive computing and possibly identify critical issues that deserve more attention from researchers and/or by regulators. Among several insights,

ABOUT THE AUTHORS

CLAUDIO BETTINI is a full professor in the Computer Science Department at Università degli Studi di Milano, where he leads the EveryWare laboratory. His research interests include the areas of data management in mobile and pervasive computing, data privacy and security, context awareness and context reasoning, and temporal and spatiotemporal data management. Bettini received a Ph.D. in computer science from the University of Milan in 1993. He is a member of the steering committee of the IEEE International Conference on Pervasive Computing and Communications, and he has been associate editor of the *Pervasive and Mobile Computing Journal*, *The VLDB Journal*, and *IEEE Transactions on Knowledge and Data Engineering*. He is Senior Member of the IEEE. Contact him at claudio.bettini@unimi.it.

SALIL KANHERE is a professor in the School of Computer Science and Engineering at the University of New South Wales, Sydney, Australia. His research interests include the Internet of Things, pervasive computing, the blockchain, crowdsourcing, data analytics, and privacy and security. Kanhere received a Ph.D. in electrical engineering from Drexel University, Philadelphia, Pennsylvania. He regularly serves on the organizing committee of a number of IEEE and ACM international conferences. He is on the editorial board of Elsevier's pervasive and mobile computing and computer communications. He is a Senior Member of both the IEEE and ACM. Contact him at salil.kanhere@unsw.edu.au.

MARC LANGHEINRICH is a full professor with the Faculty of Informatics at the Università della Svizzera Italiana in Lugano, Switzerland. His research interests include privacy in mobile and pervasive computing systems, particularly with a view toward social compatibility. Langheinrich received a Ph.D. from ETH Zürich, Switzerland. He is a member of the steering committee of the ACM International Joint Conference on Pervasive and Ubiquitous Computing conference series and chairs the Internet of Things (IoT) Conference steering committee. He has been a general chair or program chair for most major conferences in the field, including UbiComp, IEEE International Conference on Pervasive Computing and Communications, Pervasive, and the IoT Conference and currently serves as the editor-in-chief of *IEEE Pervasive Computing*. Contact him at langheinrich@ieee.org.

ARCHAN MISRA is a professor and the associate dean of research in the School of Information Systems at Singapore Management University. His research interests include problems spanning wireless networking, mobile and pervasive computing, and urban sensing. Misra received a Ph.D. in electrical and computer engineering from the University of Maryland at College Park in May 2000. He chaired the IEEE Computer Society's Technical Committee on Computer Communications from 2005 to 2007. Contact him at archanm@smu.edu.sg.

DELPHINE REINHARDT is a full professor and head of the Computer Security and Privacy group at the University of Göttingen. Her current research interests include privacy, usability, and ubiquitous computing. Reinhardt received a doctoral degree in computer science (with distinction) on privacy in participatory sensing in 2013 at Technische Universität Darmstadt. Contact her at reinhardt@cs.uni-goettingen.de.

three partially unexpected key observations emerge from our survey.


- › Most of the researchers believe that they are well aware of their legal obligations with respect to privacy.

- › Only very few researchers mention insurmountable obstacles due to privacy regulation—the majority of respondents either did not face obstacles or found a privacy-preserving solution.

- › A majority of researchers seems to be in favor of clear and strict privacy regulation.

Although we believe that our survey population is a sufficiently large and diverse sample of the target research

community, we are also well aware of the limitations of our study: 1) the sample is not uniformly distributed over geographical areas (slightly skewed to the EU and Asia), and 2) it includes mostly academic researchers. Regarding this latter aspect, our results can be complemented by the results of related studies focused on the engineering and IT professional population,^{1,16,17} as briefly discussed in the introduction.

Overall, the observations arising from our study are surprising and certainly deserve further investigation to understand, for example, the degree to which researchers are actually informed about privacy regulations applying to their specific research and how compliant their assumed solutions are. Our study also highlights the need for enhancing the interface between privacy regulation and privacy-preserving solutions in terms of both publicly available case studies and application-specific practical guidelines. This goal can be achieved only by a multidisciplinary joint effort including legal, technical, and social expertise. 

REFERENCES

1. K. Bednar, S. Spiekermann-Hoff, and M. Langheinrich, "Engineering privacy by design: Are engineers ready to live up to the challenge?" *Inf. Soc.*, vol. 35, no. 3, pp. 122–142, 2018. doi: 10.1080/01972243.2019.1583296.
2. C. Bettini and D. Riboni, "Privacy protection in pervasive systems: State of the art and technical challenges," *Pervasive Mobile Comput.*, vol. 17, pp. 159–174, Feb. 2015. doi: 10.1016/j.pmcj.2014.09.010.
3. D. Christin, "Privacy in mobile participatory sensing: Current trends and future challenges," *J. Syst. Softw.*, vol. 116, pp. 57–68, June 2016. doi: 10.1016/j.jss.2015.03.067.
4. J. Desjardins, "How much data is generated each day?" World Economic Forum, Cologne, Switzerland, Apr. 2019. [Online]. Available: <https://www.weforum.org/agenda/2019/04/how-much-data-is-generated-each-day-cf4bdddf29f/>
5. A. Gkoulalas-Divanis and C. Bettini, Eds., *Handbook of Mobile Data Privacy*. Berlin: Springer-Verlag, 2018.
6. E. Hargittai and A. Marwick, "What can I really do? Explaining the privacy paradox with online apathy," *Int. J. Commun.*, vol. 10, pp. 3737–3757, July 27, 2016.
7. T. Kandappu et al., "Campus-scale mobile crowd-tasking: Deployment and behavioral insights," in *Proc. 19th ACM Conf. Computer-Supported Cooperative Work & Social Computing (CSCW '16)*, New York, 2016, pp. 800–812. doi: 10.1145/2818048.2819995.
8. P. Kumaraguru and L. F. Cranor, "Privacy indexes: A survey of Westin's studies," Carnegie Mellon Univ., Pittsburgh, PA, Tech. Rep. CMU-ISRI-05-138, 2005.
9. M. Langheinrich and S. Lahlou, "A troubadour approach to privacy," *Disappearing Computer Initiative*, Ambient Agoras Rep. no. 15.3.1, pp. 2–29, 2003. [Online]. Available: <https://uc.inf.usi.ch/pubs/2003%20-%20Langheinrich%20-%20Ambient%20Agora%20D15.3.1%20-%20Troubadour%20Approach%20to%20Privacy%20-%20released.pdf>
10. J. Lane, V. Stodden, S. Bender, and H. F. Nissenbaum, *Privacy, Big Data, and the Public Good: Frameworks for Engagement*. New York: Cambridge Univ. Press, 2015.
11. M. Langheinrich, "Privacy by design: Principles of privacy-aware ubiquitous systems," in *Proc. 3rd Int. Conf. Ubiquitous Computing (UbiComp 2001)* (LNCS, vol. 2201), G. D. Abowd, B. Brumitt, and S. A. Shafer, Eds. Berlin: Springer-Verlag, Sept. 2001, pp. 273–291.
12. M. Langheinrich and F. Schaub, *Privacy in Mobile and Pervasive Computing* (Synthesis Lectures on Mobile and Pervasive Computing), M. Satyanarayanan, Ed. San Rafael, CA: Morgan & Claypool, 2018.
13. H. F. Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford Univ. Press, 2009.
14. D. Rosenblum, "What anyone can know: The privacy risks of social networking sites," *IEEE Security Privacy*, vol. 5, no. 3, pp. 40–49, May 2007. doi: 10.1109/MSP.2007.75.
15. D. J. Solove, "Privacy self-management and the consent dilemma," *Harvard Law Rev.*, vol. 126, no. 7, pp. 1880–1903, 2013.
16. S. Spiekermann, J. Korunovska, and M. Langheinrich, "Inside the organization: Why privacy and security engineering is a challenge for engineers," *Proc. IEEE*, vol. 107, no. 3, pp. 600–615, Mar. 2019. doi: 10.1109/JPROC.2018.2866769.
17. I. Szekely, "What do IT professionals think about surveillance?" in *Internet and Surveillance. The Challenge of Web 2.0 and Social Media*, C. Fuchs, K. Boersma, A. Albrechtslund, and M. Sandoval, Eds. Evanston, IL: Routledge, 2011, ch. 10, pp. 198–219.
18. W. Veil, "The GDPR: The emperor's new clothes—On the structural shortcomings of both the old and the new data protection law," *Neue Zeitschrift für Verwaltungsrecht*, vol. 10, no. 2018, pp. 686–696, 2018.



IEEE SA Open: Engaging Industry, Academia, and Researchers in Open Source Development

Robby Robson, Eduworks Corporation

The use of open source software has expanded exponentially. The IEEE Standards Association (IEEE SA) and its Corporate Advisory Group have worked to enable open source development for IEEE standards. Recently, IEEE SA announced IEEE SA Open, a new platform for open source.

On 4 March 2020, the IEEE Standards Association (IEEE SA) officially announced IEEE SA

Open (<https://standards.ieee.org/opensource>), a new IEEE platform for developing and disseminating open source software, hardware, and documentation meant to address typical challenges encountered by open source developers, such as attracting relevant participation and issues related to intellectual property.¹ Triggered by market demand for the inclusion of open source in standards, IEEE SA Open is evolving into a growing set of services that leverages the strengths of the IEEE as a global technical and professional community and can support all types of open source projects, from student and hobbyist projects, to open source code associated with scientific publications, to industry initiatives. In addition, the IEEE Foundation's new Collaborative

Open Innovation Lab allows for donations to collaborative open innovative projects overseen by IEEE SA. This article outlines the platform; its operation and philosophy; its uses

to date; and its potential significance for open standards, open science, and the software industry.

THE PLATFORM

IEEE SA Open is conceived as a freely available open collaboration and innovation platform. It supports development and operations (DevOps) with managed enterprise instances of GitLab and added components such as Mattermost (for team messaging), Amazon Elastic Kubernetes Service (for container orchestration and continuous integration), PlantUML (for creating UML diagrams from plain text), project wikis and webpages, and the ability to link to a GitHub account. This provides a fully featured DevOps

would any other public DevOps platform or repository, keeping their work private or making it public as they wish. However, several other levels of engagement with IEEE SA Open are available.

The first level beyond an individual project is a group project, meant to support broader community or entrepreneurial efforts. Groups must be approved by the IEEE SA Open community manager and have an approved open source license. The currently approved licenses are Apache 2.0, BSD 3-Clause, and the CERN Open Hardware version 1.2, although others may be considered upon request. As mentioned previously, there is no cost and no membership requirement.

The IEEE Standards Association and its Corporate Advisory Group have worked to enable open source development for IEEE standards.

and project management environment with components that are familiar to the open source community. IEEE SA Open also offers added services. Many of these, such as support from a community management team, dependency and security scans, and team messaging, are freely available to all projects. Others, such as the ability to create IEEE-approved releases and fiscal sponsorship, are available only to projects with an approved governance structure. Further capabilities, such as peer review, are still in the process of being explored and defined.

Anyone can use IEEE SA Open by signing up for a free IEEE account, accepting the terms of use,² and signing in. A simple default contributor license agreement is included in the terms of use, so starting an individual project requires little more than clicking a new project button and being assigned a repository. There is no cost, and IEEE membership is not required. Students, researchers, and entrepreneurs wishing to develop code or post data can use the IEEE SA Open platform just as they

The next level beyond an individual project is an official IEEE open source project. These projects must designate project leads and maintainers and have a documented governance process, which can range from community based and democratic to one controlled by a committee or a single individual or company, and must be approved by the IEEE SA Open Source Committee (OSCom) described below. Official IEEE projects have the ability to create the IEEE-approved releases, subject to reviews intended to avoid copyright, security, and other issues. Official IEEE project leads and maintainers must be IEEE Members, but contributors need not be. These projects have more governance responsibilities and formality but also receive more services and the advantages of using the IEEE brand.

OPEN SOURCE AND STANDARDS

Finally, there is open source incorporated into standards. IEEE SA produces

some of the most well-known and widely adopted standards in the world, so in developing IEEE SA Open, it was necessary to understand how open source could add value to standards while maintaining their integrity. To understand how this could work in practice, the IEEE SA invited applications to serve as pilot projects. This resulted in 13 projects in areas ranging from health to hardware, four of which are described in Table 1. As illustrated by the examples, there are multiple ways that open source can appear in standards. The most straightforward are as normative or informative references, for example, a standard might point to open source packages that must be used to conform to the standard (normative) or to reference implementations that help understand and apply the standard (informative). However, even in the small number of pilot projects, open source was used to generate data that appear in a standard, enabling the results to be verified and replicated by anyone; GitLab was used to maintain documents and schema referenced in standards; and continuous integration was used to containerize software, thereby supporting rapid deployment and adoption.

OPEN COMMUNITY-BASED GOVERNANCE

IEEE SA has extensive experience governing the open consensus processes used in standards development. The policies and procedures used are carefully documented and defined, maintained, and applied by volunteers from IEEE SA. This same philosophy of governance is reflected in IEEE SA Open. Official IEEE open source projects are approved by an OSCom that consists of IEEE volunteers with open source experience and has been tasked with providing guidance, oversight, and lifecycle management support for IEEE open source projects. The policy and procedures for OSCom, together with other resources, such as a maintainer's manual, are available on <https://opensource.ieee.org>. OSCom is exploring other

TABLE 1. The notable features of four IEE SA pilot projects.

Project	Notable features
1076-2019— <i>IEEE Standard for VHDL Language Reference Manual</i> . VHSIC Hardware Description Language (VHDL) is a formal notation intended for use in all phases of the creation of electronic systems that are both machine and human readable. It supports the development, verification, synthesis, and testing of hardware designs; the communication of hardware design data; and the maintenance, modification, and procurement of hardware.	This standard, published in 2019, includes normative references to open source packages. The standard does not reference a specific version, meaning that the latest release is the version of record. In addition, continuous integration is used to automatically test new versions of VHDL defined by the standard using a VHDL simulator that is part of a separate open source project on GitHub (https://opensource.ieee.org/vasg/Packages).
P1451.1.4— <i>Standard for a Smart Transducer Interface for Sensors, Actuators, and Devices—eXtensible Messaging and Presence Protocol (XMPP) for Networked Device Communication</i> . This standard defines a method for transporting IEEE 1451 messages over a network using XMPP. It is intended to support the Internet of Things.	The P1451.1.4 working group has created an open source project for XMPP interface descriptions that are expected to be referenced normatively by version in the standard (https://gitlab.com/IEEE-SA/XMPPI/IoT/-/tree/master/).
P370— <i>IEEE Draft Electrical Characterization of Printed Circuit Board and Related Interconnects at Frequencies up to 50 GHz</i> . This standard, which is currently under development, provides recommended practices for ensuring the quality of measured data for high-frequency electrical interconnects at bandwidths up to 50 GHz.	The P370 working group has created an open source project that contains the code used to generate data included in its draft standard. Although the standard has not yet been finalized and published, the code in this repository is being adopted by industry (https://gitlab.com/IEEE-SA/ElecChar/P370).
P9724.1— <i>xAPI Work Group</i> . xAPI, or experience an application programming interface (API), is an open specification published by the U.S. Advanced Distributed Learning initiative. Learning management systems, simulations, serious games, intelligent tutoring systems, and other systems that deliver education and training can use xAPI to report and record the results of student and learner activities.	The current xAPI spec is available on GitHub at https://github.com/adlnet/xAPI-Spec and, together with its predecessor (called SCORM), is a de facto standard supported by most commercial learning management systems. The xAPI work group is porting API specifications and other materials, written in markdown language, to IEEE SA Open, with the intent to reference them normatively and informatively in a standard.

SCORM: shared content object reference model.

ways to carry out its business as open source projects.

Another example of open community-based governance is the concept of peer review, which, at the time of this writing, has yet to be fully defined and implemented. Open source projects incorporated into standards are governed by standards working groups and have built-in peer review, first by the working group and then through a consensus balloting process that invites comments from interested IEEE SA members (more than 7,000 people and 300 corporations) and the public. In contrast to peer-review processes used for publications, research grants, and other purposes, this type of peer review is completely open, that is, everyone with an interest can see every comment or suggestion made. OSCom envisions tapping into the broader IEEE membership and their


professional networks to provide open peer review to open source projects not associated with standards development, including projects developing open source hardware designs, where flaws can lead to millions of dollars of retooling once the designs are translated into production.

from the code itself to the services supported by the code and the data these services generate. The efforts that eventually led to IEEE SA Open were initiated by corporate IEEE SA members who recognized this trend and believed that the IEEE could serve as a neutral and trusted provider of

IEEE SA has extensive experience governing the open consensus processes used in standards development.

The use of open source software in industry is growing by leaps and bounds, in part because software-as-a-service delivery models and artificial intelligence (AI) have shifted the core value of software

managed open source services. This remains an important aspect, but IEEE SA also has the potential to serve the greater good by providing opportunities to inspect, replicate, and validate software and hardware designs

used in manufacturing, AI, the Internet of Things, health care, autonomous vehicles, smart cities, and many other areas of science and engineering. All IEEE and IEEE Computer Society members are invited to visit IEEE SA Open, to consider moving or starting open source projects to it and express interest in helping the project evolve by sending an email to the open source community manager team at opensource@ieee.org. 


REFERENCES

1. IEEE Standards Association, "News and events: Press

releases," 2020. [Online]. Available: <https://standards.ieee.org/news/2020/silona-bonewald-ieee-sa-open-first-executive-director.html>

2. IEEE Standards Association, "Terms of use," 2020. [Online]. Available: <https://opensource.ieee.org/community/cla/terms-of-use/-/blob/master/terms-of-use.md>


ROBBY ROBSON is chair of the IEEE Standards Association (SA) Open Source Committee, a member of the IEEE SA Board of Governors, and chair of the IEEE Computer Society Standards Activity Board Standards Committee. He is the cofounder and chief executive officer of Eduworks Corporation and is currently overseeing projects in open source software development for competency management and talent analytics. Contact him at robby@computer.org.



www.computer.org/cga

IEEE *Computer Graphics and Applications* bridges the theory and practice of computer graphics. Subscribe to CG&A and

- stay current on the latest tools and applications and gain invaluable practical and research knowledge,
- discover cutting-edge applications and learn more about the latest techniques, and
- benefit from CG&A's active and connected editorial board.



The Top 10 Risks of Machine Learning Security

Gary McGraw, Richie Bonett, Victor Shepardson, and Harold Figueroa,
Berryville Institute of Machine Learning

Our recent architectural risk analysis of machine learning systems identified 78 particular risks associated with nine specific components found in most machine learning systems. In this article, we describe and discuss the 10 most important security risks of those 78.

At the Berryville Institute of Machine Learning (BIML), we are interested in “building security in” to machine learning (ML) systems from a security engineering perspective. This means understanding how ML systems are designed for security, teasing out possible security engineering risks, and making such risks explicit. We are also interested in the impact of including an ML system as part of a larger design. Our basic motivating question is how do we secure ML systems proactively while we are designing and

building them? Toward that end, we completed and published an architectural risk analysis (ARA) as an important first step in our mission to help engineers and researchers secure ML systems.¹ In this article, we briefly describe the top 10 of those 78 risks.

ML systems come in a variety of shapes and sizes; frankly, each possible ML design deserves its specific ARA. In our report, we describe a

generic ML system in terms of its constituent components and work through that generic system, ferreting out risks. The idea driving us is that risks that apply to this generic ML system will almost certainly apply in any specific ML system. By starting with our ARA, an ML system engineer concerned with security can get a jump start on determining risks in his or her specific system.

Figure 1 shows how we choose to represent a generic ML system. We describe the following nine basic components that align with various steps in setting up, training, and fielding an ML system: 1) raw data in the world, 2) data set assembly, 3) data sets, 4) learning algorithm, 5) evaluation, 6) inputs, 7) model, 8) inference algorithm, and 9) outputs.

Digital Object Identifier 10.1109/MC.2020.2984868
Date of current version: 4 June 2020

Note that in our generic model, both processes and collections are treated as components. Processes—that is, components 2, 4, 5, and 8—are represented by ovals, whereas things and collections of things—that is, components 1, 3, 6, 7, and 9—are represented as rectangles. On the BIML website, we have published the “BIML Interactive ML Risk Framework,” which details the risks associated with each component.

TOP 10 SECURITY RISKS OF ML

After identifying risks in each component, we considered the system as a whole and identified what we believe are the top 10 ML security risks. These threats come in two relatively distinct flavors, both equally valid: some are

associated with the intentional actions of an attacker, while others are associated with an intrinsic design flaw. Such flaws emerge when engineers with good intentions screw things up. Of course, attackers can also go after intrinsic design flaws, complicating the situation. The top 10 ML security risks are briefly introduced and discussed here.

1) Adversarial examples

Probably the most commonly discussed attacks against ML have come to be known as *adversarial examples*. The basic idea is to fool an ML system by providing malicious input, often involving very small perturbations that cause the system to make a false prediction or categorization. Although coverage and resulting attention might

be disproportionately large, swamping out other important ML risks, adversarial examples are very much real.

One of the most important categories of computer security risks is malicious input. The ML version has come to be known as *adversarial examples*. While important, these examples have received so much attention that they drown out all other risks in most people’s imaginations.²

2) Data poisoning

Data play an outsized role in the security of an ML system. That’s because an ML system learns to do what it does directly from data. If an attacker can intentionally manipulate the data being used by an ML system in a coordinated fashion, the entire system can be compromised. Data poisoning attacks require special attention. In particular, ML engineers should consider what fraction of the training data an attacker can control and to what extent.

The first three components in our generic model (raw data in the world, data set assembly, and data sets) are subject to poisoning attacks in which an assailant intentionally manipulates data in any or all of the three first components, possibly in a coordinated fashion, to cause ML training to go awry. In some sense, this risk is related both to data sensitivity and to the fact that the data themselves carry so much of the water in an ML system. Data poisoning attacks require special attention. In particular, ML engineers should consider what fraction of the training data an attacker can control and to what extent.³

3) Online system manipulation

An ML system is said to be online when it continues to learn during operational use, modifying its behavior over time. In this case, a clever attacker can nudge the still-learning system in the wrong direction on purpose through system input and slowly re-train the ML system to do the incorrect thing. Note that such an attack

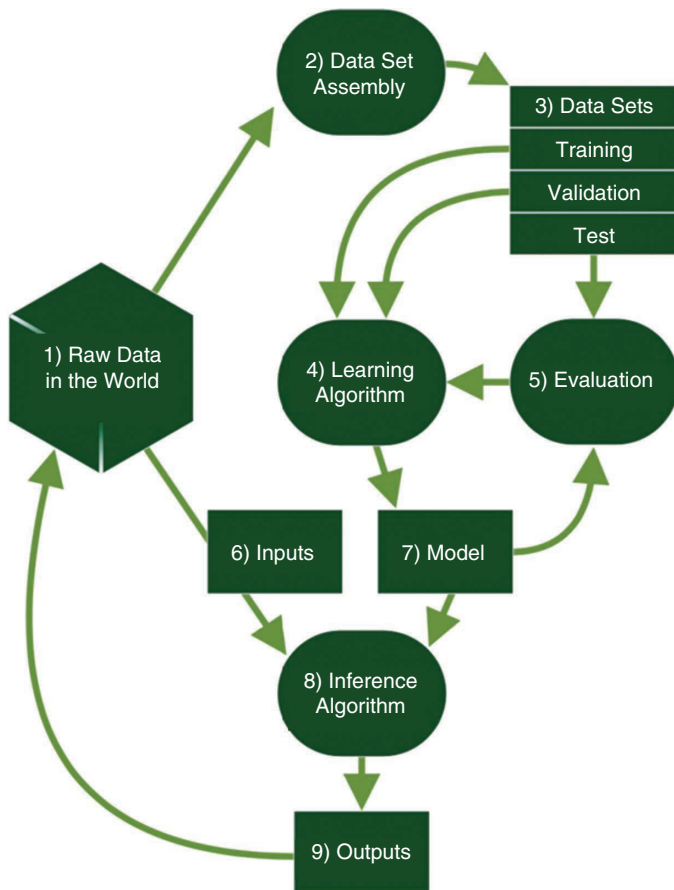


FIGURE 1. The components of a generic ML system. The arrows represent information flow.

can be both subtle and reasonably easy to carry out. This risk is complex, demanding that ML engineers consider data provenance, algorithm choice, and system operations to properly address it.

An online learning system that continues to adjust its learning during operations may drift from its intended operational use case. Skillful assailants can shift an online learning system in the wrong direction on purpose. A fielded model operating in an online system (that is, still learning) can be pushed past its boundaries. An attacker may be able to carry this out quite easily. Real-time data set manipulations can be particularly tricky in an online network where an attacker can slowly retrain the ML system to do the wrong thing by intentionally shifting the overall data set.

4) Transfer learning attack

In many cases in the real world, ML systems are constructed by taking advantage of an already-trained base model that is then finely tuned to carry out a more specific task. A data transfer attack takes place when the base system is compromised or otherwise unsuitable, making unanticipated behavior defined by the attacker possible.

Many ML systems are constructed by tuning an already trained base model so that its somewhat generic capabilities are perfected with a round of specialized training. A transfer attack presents an important risk in this situation. In cases in which the pretrained model is widely available, an attacker may be able to devise attacks using it, which will be robust enough to succeed against your (unavailable to the attacker) tuned task-specific model. You should also consider whether the ML system you are refining could possibly be a Trojan that includes sneaky behavior that is unanticipated.⁴

ML systems are reused intentionally in transfer situations. The risk of transfer outside of intended use

applies. Groups posting models for transfer would do well to precisely describe exactly what their systems do and how they control the risks in this document. A model transfer leads to the possibility that what is being reused may be a Trojaned (or otherwise damaged) version of the model being sought.

public data sources that may be manipulated or poisoned and online models.

Data sources may not be trustworthy, suitable, and reliable. How might an attacker tamper with or otherwise poison raw input data? What happens if input drifts, changes, or disappears?⁷

You should also consider whether the ML system you are refining could possibly be a Trojan that includes sneaky behavior that is unanticipated.

5) Data confidentiality

Data protection is difficult enough without throwing ML into the mix. One unique challenge in ML is protecting sensitive or confidential data that, through training, are built right into a model. Subtle but effective extraction attacks against an ML system's data are an important category of risk.

Preserving data confidentiality in an ML system is more challenging than in a standard computing situation because an ML system that is trained up on confidential or sensitive data will have some aspects of those data built right into it through training. Attacks to extract sensitive and confidential information from ML systems (indirectly through normal use) are well known.⁵ Note that even subsymbolic feature extraction may be useful since that can be used to hone adversarial input attacks.⁶

6) Data trustworthiness

Because data play an outside role in ML security, considering data provenance and integrity is essential. Are the data suitable and of high enough quality to support ML? Are sensors reliable? How is data integrity preserved? Understanding the nature of ML system data sources (during both training and execution) is of critical importance. Data-borne risks are particularly tricky when it comes to

7) Reproducibility

When science and engineering are sloppy, everyone suffers. Unfortunately, because of inherent inscrutability and the hyper-rapid growth of the field, ML system results are often underreported, poorly described, and otherwise impossible to reproduce. When a system cannot be reproduced and nobody notices, bad things can happen.

Results that cannot be reproduced may lead to overconfidence in a particular ML system to perform as desired. Often, critical details are missing from the description of a reported model. Also, results tend to be very fragile; running a training process on a different graphics processing unit (even one that is supposed to be identical in specifications) can often produce dramatically different results. In academic work, there is often a tendency to tweak the authors' system until it outperforms the baseline (which does not benefit from similar tweaking), resulting in misleading conclusions that make people think a particular idea is good when it was not actually improving over a simpler, earlier method.

8) Overfitting

ML systems are regularly very powerful. Sometimes they can be too powerful for their own good. When an ML system "memorizes" its training data set, it will not generalize to new data

and is said to be overfitting. Overfit models are particularly easy to attack. Keep in mind that overfitting is possible in concert with online system manipulation and may happen while a system is running.

A sufficiently powerful machine is capable of learning its training data set so well that it essentially builds a lookup table. The unfortunate side effect of “perfect” learning like this is an inability to generalize outside of the training set. Overfit models can

Encoding the integrity issues noted can be both introduced and exacerbated during preprocessing. Does the preprocessing step itself introduce security problems? Bias in raw data processing can impact ethical and moral implications. Normalization of Unicode to ASCII may introduce problems when encoding, for example, improper Spanish, losing diacritics and accent marks.

Metadata may help or hurt an ML model. Make note of metadata in-

impact the larger system in which the ML subsystem is encompassed. There are many ways to do this kind of thing. Probably the most common attack would be to interpose between the output stream and the receiver. Because models are sometimes opaque, unverified output may simply be used with little scrutiny, meaning that an interposing attacker may have an easy time hiding in plain sight.

A sufficiently powerful machine is capable of learning its training data set so well that it essentially builds a lookup table.

be quite easy to attack through input since adversarial examples need to be only a short distance away from training examples in input space. Note that generative models can suffer from overfitting too, but the phenomenon may be much more difficult to notice.

9) Encoding integrity

Data are often encoded, filtered, rerepresented, and otherwise processed before use in an ML system (in most cases by a human engineering group). Encoding integrity issues can bias a model in interesting and disturbing ways. For example, encodings that include metadata may allow an ML model to solve a categorization problem by overemphasizing the metadata and ignoring the real issue.

Raw data may not be representative of the problem you are trying to solve with ML. Is your sampling capability lossy? Are there ethical or moral implications built into your raw data (for example, racist or xenophobic implications can be trained right into some facial recognition systems if data sets are poorly designed)?⁸

cluded in a raw input data set; it may be a hazardous feature that appears useful on the face of it but actually degrades generalization. Metadata may also be open to tampering attacks that can confuse an ML model. More information is not always helpful, and metadata may harbor spurious correlations. Consider this example: we might hope to boost the performance of our image classifier by including exchangeable image file data from the camera. But what if it turns out that our training data images of dogs are all high-resolution stock photos, but our images of cats are mostly Facebook memes? Our model will probably make decisions based on metadata rather than content.

10) Output integrity

If an attacker can interpose between an ML system and the world, a direct attack on output may be possible. The inscrutability of ML operations (that is, not really understanding how they do what they do) may make an output integrity attack that much easier since an anomaly may be more difficult to detect.

Imagine that an attacker tweaks the output stream directly. This will

This document presents only 10 of the 78 specific risks associated with a generic ML system identified in a basic ARA by BIML.¹ Our risk analysis results are meant to help ML systems engineers in securing their particular ML systems.

In our view, ML systems engineers can devise and field a more secure ML system by carefully considering risks while designing, implementing, and fielding their specific ML system. In security, the devil is in the details, and we attempt to provide as much detail as possible regarding ML security risks and some basic security controls. **■**

REFERENCES

1. G. McGraw, H. Figueroa, V. Shepardon, and R. Bonett, “An architectural risk analysis of machine learning systems: Toward more secure machine learning,” Berryville Institute of Machine Learning, Clarke County, VA. Accessed on: Mar. 23, 2020. [Online]. Available: <https://berryvilleiml.com/results/ara.pdf>
2. X. Yu, P. He, Q. Zhu, and X. Li, “Adversarial examples: Attacks and defenses for deep learning,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 9, pp. 2805–2824, 2019. doi: 10.1109/TNNLS.2018.2886017.
3. S. Alfeld, X. Zhu, and P. Barford, “Data poisoning attacks against autoregressive models,” in *Proc. 30th AAAI Conf. Artificial Intelligence*, Phoenix, AZ, Feb. 2016. pp. 1452–1458. doi: 10.5555/3016100.3016102. [Online]. Available: <https://www.aaai.org/ocs/index.php/AAAI/>

AAAI16/paper/view/12049

4. G. McGraw, R. Bonett, H. Figueroa, and V. Shepardson, "Securing engineering for machine learning," *Computer*, vol. 52, no. 8, pp. 54–57, 2019. doi: 10.1109/MC.2019.2909955.
5. R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *Proc. 2017 IEEE Symp. Security Privacy*, pp. 3–18. doi: 10.1109/SP.2017.41.
6. N. Papernot, A Marauder's map of security and privacy in machine learning. 2018. [Online]. Available: arXiv:1811.01134
7. M. Barreno, B. Nelson, A. D. Joseph, and J. D. Tygar, "The security of machine learning," *Mach. Learn.*, vol.

81, no. 2, pp. 121–148, Nov. 2010. doi: 10.1007/s10994-010-5188-5.

8. P. Phillips, P. Jonathon, F. Jiang, A. Narvekar, J. Ayyad, and A. J. O'Toole,

"An other-race effect for face recognition algorithms," *ACM Trans. Appl. Percept.*, vol. 8, no. 2, p. 14, 2011. doi: 10.1145/1870076.1870082.

GARY MCGRAW is a cofounder of the Berryville Institute of Machine Learning. Contact him at gem@garymcgraw.com.

RICHIE BONETT is with the College of William & Mary and the Berryville Institute of Machine Learning. Contact him at richiebonett@gmail.com.

VICTOR SHEPARDSON is with Ntrepid and the Berryville Institute of Machine Learning. Contact him at victor.shepardson@gmail.com.

HAROLD FIGUEROA is with Ntrepid and the Berryville Institute of Machine Learning. Contact him at harold.figueroa@gmail.com.



CALL FOR ARTICLES

IT Professional seeks original submissions on technology solutions for the enterprise. Topics include

- emerging technologies,
- cloud computing,
- Web 2.0 and services,
- cybersecurity,
- mobile computing,
- green IT,
- RFID,
- social software,
- data management and mining,
- systems integration,
- communication networks,
- datacenter operations,
- IT asset management, and
- health information technology.

We welcome articles accompanied by web-based demos. For more information, see our author guidelines at www.computer.org/itpro/author.htm.

WWW.COMPUTER.ORG/ITPRO

Digital Object Identifier 10.1109/MC.2020.2993513





FROM THE EDITOR

Rather than always devising and creating new Internet of Things (IoT) solutions and systems, sometimes it's good to take a step back and consider the entire landscape. Often, technology interacts with social and economic systems in surprising ways, requiring a clear understanding to constrict effective solutions—particularly with security, where it's so easy to get it *wrong* if not everything is done right. In this article, several case studies of IoT security shed light on the inherent complexity involved with continuing to imbed technology into our everyday lives. —Trevor Pering

identifying various dimensions where the IoC showed blunders and the IoT, if it repeats these blunders, creates much more risk because it penetrates so much more. That book was published a bit over two years ago. Let's use its framework and look at what's happened since then.

INTERACTION WITH THE PHYSICAL

Due its permeation, IoT compromises can disrupt the physical world in surprising ways (see Figure 1). We have seen that repeatedly in the last year. In April 2019, *The Washington Post* reported that “in Australia, hacked Lime scooters spew racism and profanity.”

The video is straight out of a goofy, low budget horror movie: A row of bright-green Lime scooters, parked neatly on a sidewalk, have come to life, unleashing a filthy rush of human speech.

In June 2019, *The Verge* reported that “smart ovens have been turning on overnight and preheating to 400 degrees.” In August 2019, *ZDNet* ran the interesting headline: “Employees connect nuclear plant to the Internet so they can mine cryptocurrency.”

The attack surface of smart infrastructure can also include traditional IoC devices on the back end. In August 2019, *Ars Technica* reported on

the discovery of an open database of 28 million records—including plain text credentials—used for smart building security. We've also seen more evidence that the nation-state angle comes into play: the same month, *Ars Technica* stated that “hackers working for the Russian government have been using printers, video decoders, and other so-called Internet-of-things devices as a beachhead to penetrate targeted computer networks.”

However, the potential surprising interaction between the IoT and the physical world can also go in the other

direction. In October 2018, *Vice* reported that “a freshly installed MRI machine appeared to disable every iOS device in the hospital.” As MRI machines use strong magnets, it was tempting to conclude that the magnetic field was to blame, except that “would have disabled all electronic devices, not just iPhones.” It turns out that the machine cooled the magnet with helium, which it then vented—and helium disrupts the microelectromechanical systems (MEMS)-based clocks in newer Apple devices. (The older, standard technology of quartz crystals is immune.)

In May 2019, according to *The New York Times*, in a town in Ohio, “garage door openers and car key fobs mysteriously stopped working.” The problem turned out to be that “a homemade battery-operated device” one resident used to turn off a light was “persistently putting out a 315 megahertz signal,” inadvertently disrupting the key fobs. In November 2019, *The New York Times* stated that researchers could compromise home IoT smart speakers via a laser.



FIGURE 1. Due to its permeation throughout the physical world, security compromises in the IoT can disrupt the physical world in surprising ways.

Technology is also removing IoT protections that depend on physical distance. In recent years, colleagues in the electric power sector have worried that drones may bypass traditional physical defenses—one can no longer depend on a strong fence around an IT system to protect against a weak password at its interface. In August 2019, *Wired* reported DefCon researchers demonstrating that drones can bypass security in smart TVs.

SHELF LIVES

The security of the IoC follows in part from its “penetrate and patch” paradigm. Software holes are endemic, but frequent repair can mitigate the problem. I previously lamented that the long life of physical things does not match the short security shelf life of software—a problem exacerbated by the somewhat short shelf life of many IT vendors.

We have seen that problem continue to manifest itself. In July 2019, *Wired* reported on a set of serious vulnerabilities discovered in VxWorks, an operating system (OS) used in more than 2 billion IoT and industrial control systems (ICSs) devices. And in October, *Wired* noted that, due to the complexities of software development and corporate acquisition, the flawed code had in fact made its way into devices that initially had no apparent connection to VxWorks. In August 2019, *ZDNet* reported on an “unpatchable security flaw” discovered in Xilinx system-on-a-chip boards used in ICSs. Obsolete code still persists. In June 2019, *Tedium* reported that the New York City subway system still runs OS/2. On 14 January 2020, Microsoft announced that Windows 7 reached its official end of life, ending support and updates (<https://www.microsoft.com/en-us/microsoft-365/windows/end-of-windows-7-support>). Will we see any IoT issues arising from forgotten embedded OS/2 or Windows 7?

SMART INFRASTRUCTURE

Previously, I examined how traditional IoC-style bugs may hamper

visions for IoT-enabled infrastructure. The universe has provided more data points here as well. In April 2018, according to *Bleeping Computer*, “officials from the city of Innsbruck in Austria have shut down a local ski lift after two security researchers found its control panel open wide on the Internet, and allowing anyone to take control of the ski lift’s operational settings.” In August 2018, *Help Net Security* reported on the discovery of vulnerabilities in smart irrigation systems. “By simultaneously applying a distributed attack that exploits such vulnerabilities, a botnet of 1,355 smart irrigation systems can empty an urban water tower in an hour and a botnet of 23,866 smart irrigation systems can empty flood water reservoir overnight.” In March 2019, *MIT Technology Review* reported on the discovery of the Triton malware, giving adversaries remote access, in widespread ICSs, including installations such as petrochemical plants with a significant physical impact. In the same month, *Star Tribune* reported on U.S. Department of Homeland Security (DHS) warnings of security vulnerabilities in 750,000 defibrillators.

SMART VEHICLES

Problems continue here. In September 2018, *SoraNews24* reported that some smart cars would perceive the logo of a Japanese ramen chain as a “Do Not Enter” sign, with confusing results. In April 2019, *IFLScience* reported that hackers “managed to cause a Tesla’s self-driving feature to swerve off course and into the wrong lane, using just a few stickers they placed on the road.”

Problems are not confined to smart vehicles on the ground. *ZDNet* (in December 2018) catalogued a lengthy list of cybersecurity incidents (for example, incapacitation due to ransomware) affecting ships at sea. It is rumored that the summer 2019 incident of Iran seizing a British tanker was enabled because Iran used GPS spoofing to trick the tanker into unwittingly entering Iranian waters, where such

seizure would be legal. The potential for terrorists using cyber means to trick a container ship into loading its cargo in a dangerous weight configuration is also a worry.

In November 2018, the U.S. Federal Aviation Administration issued a Special Airworthiness Information Bulletin warning of a certain flight display system “repeatedly resetting itself in-flight,” resulting in “temporary loss of all flight display information”; however, “pulling the ADS-B circuit breaker as described in section 3.2.6 of the Flight Manual supplement has been demonstrated to resolve the issue.” In July 2019, *The Register* reported that the EU Aviation Safety Agency urges rebooting of Airbus A350 airplanes after 149 h of operation “to prevent ‘partial or total loss of some avionics systems or functions.’” The British sitcom *The IT Crowd* may be off the air, but its dictum of “Have you tried turning it off and on again?” lives on.

The recent trouble with the Boeing 737 Max, much covered by the mainstream media, provides much darker data points.

DESIGN ANTIPATTERNS FOR INSECURITY

We have seen the traditional blunder of hard-coded passwords continue in the IoT. In March 2018, the DHS warned of “default or hard-coded credentials in multiple GE Healthcare products.” In February 2019, *TechCrunch* reported the discovery of “thousands of individually exposed Internet-connected industrial refrigerators that can be easily remotely instructed to defrost” (potentially leading to various physical consequences)—due to hardcoded passwords. In March 2019, *ZDNet* reported that French thieves were able to steal over 26,000 gallons of gasoline “because some gas station managers didn’t change the gas pump’s default lock code from the standard 0000.” In September 2019, *ZDNet* reported that more than a half million GPS trackers were deployed with a default password

of “123456.” In January 2020, *ZDNet* reported that a hacker had published “a massive list of Telnet credentials for more than 515,000 servers, home routers, and IoT (Internet of Things) smart devices.” Although the term *hacker* has a pejorative connotation, it’s not clear the hacker is the one deserving criticism. The credentials were either factory defaults or easily guessed—and the devices were still running Telnet, exposing credentials over open and unauthenticated channels.

We have also seen more subtle bugs in protocol implementations cause trouble. Common Vulnerabilities and Exposures (CVE)-2018-10933 documented an amusing and dangerous bug in libssh implementations: because of shared structure in client-side and server-side code, a rogue client could gain access to authenticated connection simply by declaring `SSH2_MSG_USERAUTH_SUCCESS`. *Securolytics* then reported discovering this vulnerability in critical IoT devices. In other protocol blunders, *Ars Technica* reported that “Nike’s self-lacing sneakers turn into bricks after faulty firmware update;” *TechCrunch* reported that the Amazon Ring doorbell, as part of some handshaking when the device is configured, will broadcast the home’s Wi-Fi credentials in plain text; *Vice* reported on IoT communication protocol bugs enabling adversarial access to Nest cameras.

PUBLIC-KEY INFRASTRUCTURE

The numbers of devices and organizations in IoT visions suggests the need for public-key cryptography; yet, the potential population size suggests we be aware of public-key infrastructure (PKI) scalability trouble. We haven’t seen this trouble per se, but we have seen PKI bugs worth paying attention to. In March 2019, *Vice* reported that ASUS “was used to unwittingly install a malicious backdoor on thousands of its customers’ computers last year after attackers compromised a server for the company’s live software update

tool.” The update process used digital signatures, and the updates themselves had valid signatures; the attackers got around the protections of this secure channel by compromising the other end of it. January 2020 brought news of a discovery of vulnerability in a new version Microsoft’s CryptoAPI used in Windows 10. This bug would enable adversaries to get around the protections of HTTPS communication channels based on elliptic-curve cryptography (ECC); the attractiveness of ECC for embedded systems suggests this could be a particular risk for the IoT, although I have not seen any documentation of IoT attacks so far.

Even more recently, CVE-2020-6961 flags medical telemetry devices where adversaries can access private keys. Whether these keys belong to the device or to some remote “trusted” party is not clear.

THE INTERNET OF BETRAYING DEVICES

By providing a digital connection to previously unconnected aspects of human existence, the IoT continues to lead to surprising privacy leaks. In January 2019, the German magazine *c’t* reported that a consumer used the EU General Data Protection Regulation to request his personal data from Amazon. The consumer received a 100-megabyte zip file that included much audio and transcribed audio from Alexa—but that audio material was all of someone else, as the consumer didn’t own an Alexa device. In February 2019, *Circuit Breaker* reported on the fallout from the revelation that Google’s “Nest Secure home security system included an on-device microphone.” In October 2019, it then reported that researchers had developed proof-of-concept apps for Amazon and Google smart speakers that made it through the vendors’ “security-vetting processes” and “surreptitiously eavesdropped on users and phished for their passwords.”

In July 2019, *ZDNet* reported that researchers discovered how Bluetooth devices can be tracked due to some

subtle flaws in address randomization techniques intended to defend against that. In November 2019, *The Washington Post* reported that “soldiers in an intelligence unit with top-secret clearances were ordered by their commander to download an information app” whose “terms of service said it could collect substantial amounts of personal data and that the developer has a presence overseas.” The same month, *NBC News* reported that researchers discovered smart TVs and other home devices were reporting data to remote third-party companies, such as Netflix, even if the consumer had no connection to them. In December 2019, *TechCrunch* reported that even the U.S. Federal Bureau of Investigation (in Portland, Oregon) was warning consumers about the privacy risks of smart TVs.

On a brighter note, the Associated Press reported in July 2019 that four Maryland teenagers vandalized their school but were caught because of surprising connectivity. “Though all four boys covered their faces during the hate crime, they didn’t realize their cellphones automatically connected to Glenelg High School’s Wi-Fi under their individual student IDs.”

IoT AND THE LAW

The interaction of the IoT and the (U.S.) legal process continues. In November 2018, *The New Hampshire Union Leader* published, “Judge Orders Amazon to Produce Recordings From Echo Device Seized From Crime Scene of Farmington Double-Murder.” In June 2019, *The Seattle Times* reported on a lawsuit alleging that “Amazon is recording children who use its Alexa devices without their consent, in violation of laws governing recordings in at least eight states, including Washington.” January 2020 brought more news from *The New Hampshire Union Leader*. In an ongoing criminal case, the prosecution wanted to use audio from an Amazon Ring, but the defense objected, as New Hampshire’s wiretapping law requires all parties (including, in this case, the alleged

criminal) to consent to being recorded. Recently, we have seen governmental entities, such as California, the United Kingdom, and the National Institute of Standards and Technology, all promulgate regulations and standards for IoT security.

Due to the speed of development and the potential advantages, humanity tends to deploy IT applications before thoroughly thinking through the issues. As we have seen over the last two years, the IoT has been no exception. We, the members of the IT community, have our work cut out for us: continuing research on

technological approaches that might systematically avoid repeating the traditional blunders of the IoC, continuing effort in the public policy space to mitigate the risks, and continuing public discussion and awareness of the issues so that society can make better-informed choices about what it really wants and how to get there. Two years after the book was published, I can say “I told you so”—but in 10 years, I’d rather be celebrating how we got it right. **■**

REFERENCES

1. S. W. Smith and J. S. Erickson, “Never mind pearl harbor: What about a cyber love canal?” *IEEE Security Privacy*, vol. 13, no. 2, pp. 94–98, Mar./Apr. 2015. doi:10.1109/MSP.2015.37.
2. S. W. Smith, *The Internet of Risky Things: Trusting the Devices That Surround Us*. Sebastopol, CA: O’Reilly Media, 2017.

SEAN W. SMITH is a professor in the Department of Computer Science at Dartmouth College. Contact him at sws@cs.dartmouth.edu.

Call for Articles

IEEE Pervasive Computing

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

Author guidelines:

www.computer.org/mc/pervasive/author.htm

Further details:

pervasive@computer.org

www.computer.org/pervasive

Digital Object Identifier 10.1109/MC.2020.2993514

IEEE
pervasive
COMPUTING
MOBILE AND UBIQUITOUS SYSTEMS



The Rise of the Quantum Internet

Marcello Caleffi, Daryus Chandra, Daniele Cuomo, Shima Hassanpour, and Angela Sara Cacciapuoti, University of Naples Federico II

The Internet just turned 50: five decades that shaped the world we live in. But what comes next, the so-called Quantum Internet, will be even more revolutionary, likely in ways we can't imagine yet.

On 29 October 1969, the first successful message was exchanged over the Arpanet, the predecessor to what we now know as the Internet. In the five decades since, the Internet has revolutionized communications to the extent that its impact on our lives is not only technological but rather has affected almost every facet of business and lifestyle, throughout the structure of society.

The Internet itself evolved amazingly during these decades, from a network comprising a few static nodes in the early days to a leviathan interconnecting half of the world's population through billions of devices. Yet the fundamental underlying assumption—the Internet's primary purpose of transmitting messages that can be successfully encoded in a sequence of classical bits—has been unchanged since the beginning.

The advent of the engineering phase of quantum technologies is challenging the Internet's fundamental assumption because quantum devices require—as communication primitives—the ability to transmit quantum information. Hence, research groups throughout the world, and ours as well, are investing their efforts to design and engineer the Quantum Internet.¹⁻⁶ But there's still a long way to go and no guarantee of getting there very soon.

THE QUANTUM REVOLUTION

Quantum technology advances have successfully enticed tech giants, such as IBM, Google, and Intel, to participate in the so-called quantum race. Several start-up companies also have been founded to join in this monumental endeavor. A very significant milestone was achieved at the end of 2019 by a group of researchers at Google, which announced quantum supremacy by solving a classically intractable problem with its quantum processor^{7,8} (see “The Quantum Supremacy”).

Immense interest in the future of quantum technologies is not only displayed by industry but also by governments around the world. To mention some initiatives, in April 2017, the European Commission launched a 10-year, €1 billion flagship project to accelerate European

quantum technologies research.⁹ Meanwhile, across the Atlantic, in September 2018, the U.S. House of Representatives unanimously approved the establishment of a National Quantum Initiative funded with US\$1.25 billion over 10 years.¹⁰

Within this context of a real quantum revolution, the ultimate vision

necessary steps as well as the novel challenges we will face on our journey toward the Quantum Internet design and deployment.

THE QUANTUM INTERNET

The Quantum Internet is a network enabling quantum communications among remote quantum devices. What

decoherence rapidly corrupts quantum information, making it challenging to rely on quantum memories.

Another constraint that makes things harder is the no-cloning theorem. Indeed, the classical Internet operates by extensively duplicating information among the different components of a network node and among different nodes. In the Quantum Internet, the no-cloning theorem forbids copying an unknown qubit. Hence, the commonly used methods for keeping the integrity of information, for example, retransmission of the same information, are now forbidden. Finally, quantum states cannot be read without affecting their states. Any attempt to measure a qubit makes its state collapse into a classical bit value—0 or 1. For this particular reason, and for the no-cloning theorem as well, the direct transmission of qubits so far appears limited to relatively short distances in the context of specific applications that can tolerate low-transmission success rates.

It becomes evident that a paradigm shift is required. Indeed, the very concept of information transmission has to be rethought and reformulated for Quantum Internet design. Thankfully, quantum mechanics provides us an amazing tool for transmitting quantum information, the quantum teleportation process, astonishingly, without the physical transfer of the qubit.

BEYOND DIRECT QUBIT TRANSMISSION

By using a unique feature of quantum mechanics, known as *entanglement* (see “Introducing Entanglement”), in 1993 Bennett et al.¹³ showed that it is possible to instantaneously transfer the quantum state encoded in a qubit at a certain sender to a qubit stored at a certain receiver without, surprisingly, the physical transfer of the qubit at the sender.³ This quantum communication protocol, already experimentally verified, is known as *quantum teleportation*.

In a nutshell, the teleportation process, portrayed in Figure 1 for a

The direct transmission of qubits so far appears limited to relatively short distances in the context of specific applications that can tolerate low-transmission success rates.

is to build a quantum network infrastructure—also known as the *Quantum Internet*—to interconnect remote quantum devices so that quantum communications among them are enabled.^{2,3} The reason behind this vision is that the Quantum Internet is capable of supporting functionalities with no direct counterpart in the classical Internet—ranging from secure communication⁵ to blind computing¹¹ through distributed quantum computing^{1,2}—as recently overviewed by the Internet Engineering Task Force.¹²

Although it is too early to tell when and how this quantum network will be deployed, our goal here is to describe how the Quantum Internet differs from the current Internet. For this, we introduce the very basic idea of the Quantum Internet and its underlying foundation, and we highlight the

sets it apart from the classical Internet is the ability to transmit quantum bits (qubits), which differ fundamentally from classical bits, and create distributed, entangled quantum states with no classical equivalent.³

Specifically, the Quantum Internet is governed by the laws of quantum mechanics. Hence, phenomena with no counterpart in classical networks, such as entanglement, the impossibility to safely read and copy the quantum information impose terrific constraints for the network design. That means most techniques adopted within the classical Internet cannot be reused here.²

Just consider how important storing information for long periods at network nodes is to classical Internet functionalities. This cannot be taken for granted in the Quantum Internet because the phenomenon known as

THE QUANTUM SUPREMACY

The term *quantum supremacy* was coined by J. Preskill in 2011^{S1} to describe the moment when a programmable quantum device would solve a problem that cannot be solved by classical computers, regardless of the usefulness of the problem.

Reference

S1. J. Preskill, “Quantum computing and the entanglement frontier,” in *Proc. 25th Solvay Conf. Physics*, Oct. 2011.

single qubit, requires 1) the generation and distribution of a maximally entangled pair of qubits (referred to as an EPR pair) between the source and destination, and 2) a classical transmission to send two classical bits. Consequently, a classical link for sending classical information and a quantum link for entanglement generation and distribution need to be established in advance.

Moreover, each teleportation process destroys the entanglement-pair member at the source. A successive teleporting requires the generation and distribution of a new entangled pair between source and destination. This, in turn, implies radically new challenges from a network design perspective, completely changing the classical concepts of network connectivity and throughput. Indeed, the connectivity between two quantum nodes is strictly determined by the availability of a shared entangled pair, and it inherently varies in time as a consequence of the depletion of the entanglement-pair member at the source.

The challenges are not limited to the above-mentioned ones. In fact, long-distance entanglement distribution still constitutes a key issue due to the decay of the entanglement distribution rate as a function of the distance.^{1,3} And because qubits cannot be copied due to the no-cloning theorem, classical signal amplification techniques cannot be employed. In this context, quantum teleportation relies on intermediate nodes, known as *quantum repeaters*, that are capable of entangling distant nodes—without physically sending an entangled qubit through the entire distance—by swapping the entanglement generated through shorter links,¹⁴ as illustrated in Figure 2.

It is evident that the design of the Quantum Internet constitutes a breakthrough from an engineering perspective. Each network functionality must be redesigned and reengineered with a solid integration of classical and quantum communications resources.¹⁵ In this regard, the classical resources for

transmitting classical bits will likely be provided by integrating such classical networks as the current Internet with the Quantum Internet.²

Paving a journey toward the Quantum Internet is indeed not a straightforward task. Historically, predictions about technological developments prove themselves true

hardly or in ways the predictor didn't expect at all. Hence, there will definitely be twists and turns in the design of the Quantum Internet, with uncertainty on when and how this goal will be accomplished (see "Realizing the Qubit").

However, we may envision roughly three subsequent necessary steps, whose complexity scales as a function of the time and the level of platform

INTRODUCING ENTANGLEMENT

Entanglement is one of the most distinguishing quantum phenomena with no counterpart in the classical world, in which the quantum states of two or more particles become inextricably linked even if they are separated by a great distance. The entanglement of quantum particles demonstrates a relationship between their fundamental properties that cannot happen arbitrarily. When a measurement is performed on one of the particles, the other particle will be instantly influenced.




FIGURE 1. A general schematic of quantum teleportation protocol, where the standard bra-ket notation $|\psi\rangle$ is adopted for describing quantum states. Notice in the figure that after quantum teleportation, the original qubit and the entanglement are destroyed. As weird as it seems, quantum teleportation fully obeys the fundamental principles of quantum mechanics. Therefore, the cost of transmitting quantum information can be exchanged with entanglement and classical communications. Because the entanglement is always destroyed after every single teleportation, it constitutes the primary consumable resource in the Quantum Internet, which means it needs to be generated continuously.

REALIZING THE QUBIT

Currently, there exist multiple technologies for realizing a qubit (quantum dots, transmons, ion traps, photons, and so forth), with each technology characterized by different pros and cons. This hardware heterogeneity will impose its own additional challenges to create an integrated Quantum Internet ecosystem.

heterogeneity, as portrayed in Figure 3. The very first step involves interconnecting multiple quantum processors within a single quantum computer. The qubits are likely to be homogeneous among the different processors, although heterogeneity may arise within due to different hardware technologies underlying memory and computational units. The link for connecting the qubits is very short, and the network topology is fixed so that only a simple addressing and routing protocol is required. Timing and synchronization need to be carefully designed.

Network functionalities that are unavailable in classical networks must be designed and implemented. For instance, quantum decoherence must be carefully accounted for within the network design so that it can be used to represent a key metric for the network functionalities. Local operations among qubits within a single processor must be complemented by remote operations—operations among qubits placed at different processors. The tradeoff between entangled qubits devoted to computation and entangled qubits devoted to communication represents a key issue with

no counterpart in the classical network design. The very concept of distributed quantum algorithm design must explicitly take such a tradeoff and the delay induced by remote operations into consideration.

The second step involves interconnecting multiple quantum computers within the same farm. At this stage, the hardware heterogeneity among the different quantum computers may arise. Such heterogeneity must be considered in network functionalities. The entanglement distribution benefits from the controlled farm environment and relatively short distances. Delay imposed by classical communication times is slightly longer compared to interprocessor wiring. Hence, this requires more sophisticated timing and synchronization. The network topology is more complex, and it may vary in time as the number of nodes in the network changes. This, in turn, induces dynamics at the network bootstrap/functioning, which

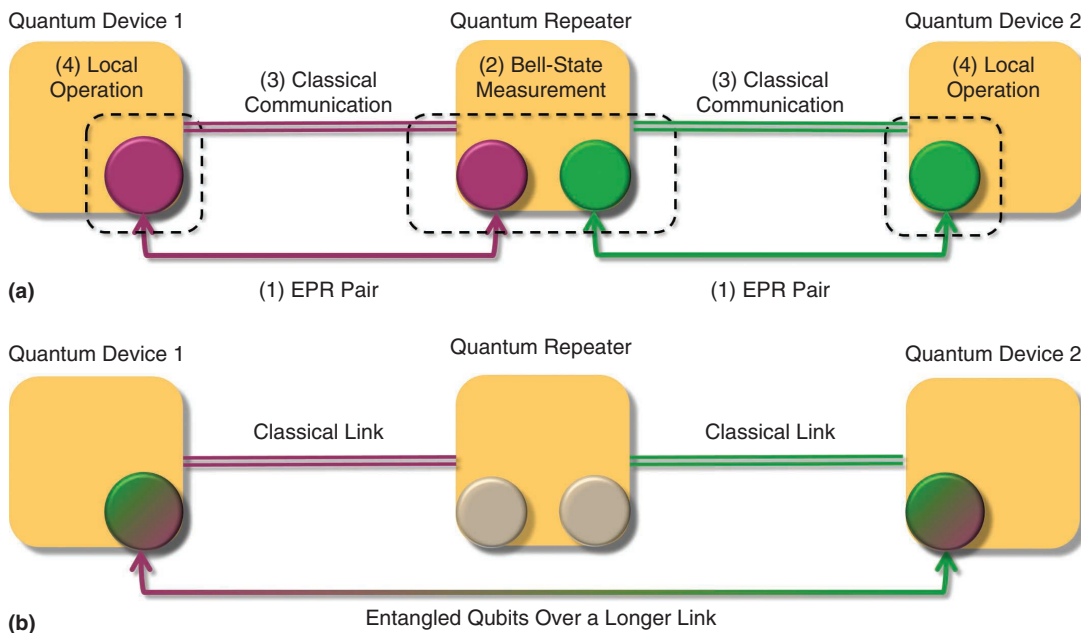


FIGURE 2. The entanglement swapping portrait. (a) Each quantum device shares an EPR pair with an intermediate node, the quantum repeater. The repeater performs Bell-state measurement on the two qubits in its possession, which results in the collapse of their quantum states into classical bits. The repeater sends the classical bits obtained from the measurement operation to the quantum devices. Finally, based on the received bits, the quantum devices perform local operations to complete the swapping process. (b) The result is that the entanglement between the quantum devices is created over a longer distance.




FIGURE 3. The necessary steps toward the envisioned Quantum Internet. We hypothesize that the complexity scales proportional to the distance of connectivity and level of platform heterogeneity among quantum farms.

requires more sophisticated strategies for routing and access as well as for mitigating quantum errors. Finally, the balance between local and remote operations—between computational and communication qubits—becomes even more intricate.

The final long-term step involves interconnecting multiple geographically distributed quantum farms. One of the key challenges is the heterogeneity among different quantum farms, which may be operated by different companies. This requires significant efforts in terms of network standardization. Furthermore, the heterogeneity among quantum links, for example, optical, free space, or satellite, will arise. The delays induced by the distances will introduce severe challenges on the entanglement generation and distribution. The increasing number of quantum devices to be wired and the heterogeneity of the environments hosting the quantum computers must be taken into account.

One of the judicious questions raised from this discussion is when will we

see the Quantum Internet? There is no definite answer to this question. However, we firmly believe this is a goal that requires a collaborative effort and a multidisciplinary approach between academics and industry. The required competences and skills are many and diverse, and each is interconnected with and vital to the others. **□**

REFERENCES

1. M. Caleffi, A. S. Cacciapuoti, and G. Bianchi, "Quantum Internet: From communication to distributed computing!" in *Proc. 5th ACM Int. Conf. Nanoscale Computing and Communication*, 2018, pp. 1–4. doi: 10.1145/3233188.3233224.
2. A. S. Cacciapuoti, M. Caleffi, F. Tafuri, F. S. Cataliotti, S. Gherardini, and G. Bianchi, "Quantum Internet: Networking challenges in distributed quantum computing," *IEEE Netw.*, vol. 34, no. 1, pp. 137–143, Jan. 2020. doi: 10.1109/MNET.001.1900092.
3. A. S. Cacciapuoti, M. Caleffi, R. Van Meter, and L. Hanzo, "When

- entanglement meets classical communications: Quantum teleportation for the Quantum Internet," *IEEE Trans. Commun.*, to be published.
4. S. Pirandola and S. L. Braunstein, "Physics: Unite to build a quantum internet," *Nature*, vol. 532, no. 7598, pp. 169–171, 2016. doi: 10.1038/532169a.
 5. S. Wehner, D. Elkouss, and R. Hanson, "Quantum Internet: A vision for the road ahead," *Sci.*, vol. 362, no. 6412, p. eaam9288, 2018. doi: 10.1126/science.aam9288.
 6. D. Castelvecchi, "The quantum internet has arrived (and it hasn't)," *Nature*, vol. 554, no. 7692, Feb. 2018. doi: 10.1038/d41586-018-01835-3.
 7. F. Arute et al., "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019. doi: 10.1038/s41586-019-1666-5.
 8. E. P. DeBenedictis, "Beyond quantum supremacy," *Computer*, vol. 53, no. 2, pp. 91–94, Feb. 2020. doi: 10.1109/MC.2019.2958446.
 9. E. Gibney, "Europe's billion-euro quantum project takes shape," *Nature News*, vol. 545, no. 7652, p. 16, 2017. doi: 10.1038/545016a.
 10. E. Gibney, "Quantum gold rush: The private funding pouring into quantum start-ups," *Nature*, vol. 574, no. 7776, pp. 22–24, 2019. doi: 10.1038/d41586-019-02935-4.
 11. A. Broadbent, J. Fitzsimons, and E. Kashefi, "Universal blind quantum computation," in *Proc. IEEE 50th Annu. Symp. Foundations Computer Science*, 2009, pp. 517–526.
 12. C. Wang, A. Rahman, and R. Li, "Applications and use cases for the Quantum Internet," Internet-Draft Draft-Wang-Qir g-Quantum-Internet-Use-Cases-04, Internet Engineering Task Force, Fremont, CA, Mar. 2020. Work in Progress.
 13. C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters "Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels," *Phys. Rev. Lett.*, vol. 70, no. 13,

pp. 1895-1899, 1993. doi: 10.1103/PhysRevLett.70.1895.

14. R. Van Meter, *Quantum Networking*. Hoboken, NJ: Wiley-ISTE, Apr. 2014.

15. W. Kozłowski, S. Wehner, R. V. Meter, B. Rijsman, A. S. Cacciapuoti, and M. Caleffi, "Architectural principles for a Quantum Internet,"

Internet-Draft Draft-Irtf-Qirg-Principles-03, Internet Engineering Task Force, Fremont, CA, Mar. 2020. Work in Progress.

MARCELLO CALEFFI is the codirector of the FLY: Future Communications Laboratory, University of Naples Federico II, Italy. He is also with the Laboratorio Nazionale di Comunicazioni Multimediali, National Inter-University Consortium for Telecommunications, Naples, Italy. Contact him at marcello.caleffi@unina.it.

DARYUS CHANDRA is with the www.QuantumInternet.it Research Group, FLY: Future Communications Laboratory, University of Naples Federico II, Italy. Contact him at daryus.chandra@unina.it.

DANIELE CUOMO is with the www.QuantumInternet.it Research Group, FLY: Future Communications Laboratory, University of Naples Federico II, Italy. Contact him at daniele.cuomo@unina.it.

SHIMA HASSANPOUR is with the www.QuantumInternet.it Research Group, FLY: Future Communications Laboratory, University of Naples Federico II, Italy. Contact her at shima.hassanpour@unina.it.

ANGELA SARA CACCIAPUOTI is the codirector of the FLY: Future Communications Laboratory, University of Naples Federico II, Italy. Contact her at angelasara.cacciapuoti@unina.it.



IEEE Security & Privacy magazine provides articles with both a practical and research bent by the top thinkers in the field.

- stay current on the latest security tools and theories and gain invaluable practical and research knowledge,
- learn more about the latest techniques and cutting-edge technology, and
- discover case studies, tutorials, columns, and in-depth interviews and podcasts for the information security industry.



computer.org/security



Inference Acceleration: Adding Brawn to the Brains

Mark Campbell, Trace3

Much attention in artificial intelligence (AI) is focused on creating and training smarter models. As these models grow in accuracy, they increase size and inefficiency. Inference acceleration enables AI to solve ever-larger problems on ever-smaller devices.

Since the artificial intelligence (AI) spring, a burgeoning ecosystem of model frameworks, reference algorithms, vast inexpensive training data repositories, and tools to streamline model development, training, and deployment has emerged. How to execute these models after creation is left as an exercise for the reader.

Digital Object Identifier 10.1109/MC.2020.2984870
Date of current version: 4 June 2020

Until recently, discussions about model optimization, execution platforms, and model tuning were not pressing issues for most AI shops. The conventional thinking has been that model creation is the tricky part. If you obtain enough of the right data, get a hefty training platform, and train the model successfully, then the heavy lifting is over. Just throw the trained model onto whatever hardware you have, stand back, and let it do its “thang.” For the most part, this approach works. After all, the precious hours of a data scientist’s day are not needed to execute the model but to make it deeper and smarter, collect more and better training data, and create new purpose-built

versions for edge cases. This all makes very good sense, except this approach is breaking down.

THE AI BLOAT PROBLEM

AI models do not grow fast—they grow very, very, very fast. One example that illustrates this model bloat is a group of natural language generators called *transformers*. In late 2018, Google AI Language published a landmark paper on Bidirectional Encoder Representations from Transformers (BERT), a natural language transformer with 345 million

parameters, which, at the time, was the largest of its kind in existence. By February 2019, OpenAI dropped jaws with Generative Pretrained Transformer 2 (GPT-2), which touted 744 million parameters, and rumors flew about a monster chained up in the organization's basement that had 1.4 billion parameters.¹ Yet, only six months later, NVIDIA announced an 8.3 billion-parameter behemoth called the Megatron Language Model,² which was soon doubled by Microsoft's DeepSpeed at 17 billion parameters in February 2020.³

information outside data centers than inside.⁷ As the center of data gravity shifts from the data center to the edge, it drags applications with it, and a great proportion of today's applications have embedded AI models. But this presents a problem.

As noted earlier, AI models are growing rapidly, making them very ill-suited to the space and power constraints of edge devices. The cumulative size of embedded AI models will soon swamp even the beefiest edge device. Edge AI applications need to be orders of magnitude smaller

to permeate the limited capacity of edge platforms.

MODEL OPTIMIZATION TRENDS

There are several model optimization techniques that reduce the size of the trained model and speed up its inference performance.

Knowledge distillation

One such technique, knowledge distillation, takes a large trained neural network (the teacher) and trains a smaller model (the student) to replicate the behavior of the teacher at every level. This results in a much smaller final student model that can then be deployed.¹⁰ A dramatic example is TinyBERT, which, as the name implies, is a distilled version of BERT that is 7.5 times smaller and 9.4 times faster while only 3% less accurate.¹¹

Pruning

Another method, pruning, takes advantage of the phenomenon that not all neurons and connections are equally valuable to the output of the overall network, and many contribute nothing. The elimination of low- or no-value connections and neurons results in a sparse network that is smaller, more compressible, and faster.¹² On a broader scale, pruning can also be applied to the filter or layer levels, yielding even more marked improvements. Pruning recently enabled several computer vision projects to reduce their models by nine to 13 times, with only a marginal degradation in accuracy.⁸ A variant of pruning, called the *lottery ticket hypothesis*, takes a dense, randomly initialized network and identifies subnetworks that, when trained in isolation, yield shockingly accurate results with only a subset of the original network. These "winning tickets" are typically 10–20% of the size of the original network and can even yield a higher accuracy than the entire network.¹³

Kernel tweaking

Another technique applied both during and after training is quantization.

AI training compute demand has grown 300,000 times since 2012, doubling approximately every 100 days.

This is a 50-fold increase in the transformer model size in 15 months, or a doubling roughly every 80 days.

THE AI COMPUTE PROBLEM

As the model size grows so does the compute power to train and run it. An oft-quoted study by OpenAI indicates that AI training compute demand has grown 300,000 times since 2012, doubling approximately every 100 days,⁴ eerily close to the growth curve seen in the preceding transformer example. Yet, as we enter the tail phase of Moore's law, microprocessor performance has increased only marginally during the same period.⁵ This increasing imbalance between the model demand and compute supply predicts that, in the not-so-distant future, the model size could be limited as a function of computational capacity.

THE EDGE AI PROBLEM

While AI training has found a home inside the chilled walls of the data center, AI inference is moving increasingly toward the edge. The edge device market expands at a 32.8% compound annual growth rate, with no signs of slowing down.⁶ By 2025, this massive surge to the edge will generate more

and faster yet still retain the accuracy of their data center-bound versions.⁸ Not the easiest of engineering feats.

But all is not lost. There are three promising solution trends changing all this.

MODEL SPECIALIZATION TRENDS

The number and variety of neural network models and model frameworks have ballooned in recent years, so much so that an open source community, the Open Neural Network Exchange (ONNX), has blossomed to enable developers to use and share neural network models on a wide spectrum of frameworks, tools, runtimes, and compilers. ONNX and other AI communities have facilitated a much higher degree of model specialization. One extreme example, SqueezeNet, is a model specifically designed to decrease size and increase inference performance. It is an image classifier with an AlexNet-level of accuracy, but it needs only 2% of the parameters that AlexNet requires and has a compressed model size of just 500 KB (50 times smaller than AlexNet).⁹ Other slimmed-down models now abound, such as MobileNet, Fritz AI, and ShuffleNet, that enable AI

Model weights and activations are typically represented as high-precision floating point numbers, which take more space and computational time than a lower-precision floating point or even integers. A precision decrease for weights and activation values saves memory bandwidth, creates improved cache locality, and reduces power consumption, with only nominal drops in accuracy.¹⁴

Unlike the model optimizations presented thus far, which are applied during and after training, kernel tweaking is applied at runtime. One such kernel tweeker is NVIDIA's TensorRT, which quantizes the model, selects a platform-specific kernel, and adjusts its memory footprint to yield the best performance for that specific model. NVIDIA conducted one test in which TensorRT yielded a performance on a ResNet50 that was 19 times faster than one run on native TensorFlow.¹⁵ Another popular kernel tweeker is the TensorFlow domain-specific accelerated linear algebra compiler (also known as XLA), which generates a sequence of custom computation kernels based on the model's TensorFlow graph to yield a modest, but welcome, 15% performance boost.¹⁶

Most model optimization efforts use a combination of distillation, pruning, and quantization to produce extremely streamlined models with only a slight loss of accuracy. However, more techniques are in the experimental stage and will undoubtedly be added to the arsenal soon.

INFERENCE PLATFORM TRENDS

For most of AI's history, CPUs have provided the platform for both training and inference. As model sizes and performance needs exploded during the past decade, it became apparent that general-purpose CPUs were not a very effective solution and that alternate computational platforms were required. In a fortuitous and now legendary breakfast discussion between Andrew Ng and Bill Dally in 2010,

graphics processing units (GPUs) entered the AI arena. After taking the idea to NVIDIA, Ng and Dally used the inherent parallelism in 48 GPUs to solve a problem that previously had taken 16,000 CPUs. So what was the problem they tackled for such a historic event? Recognizing cat images, of course. Since then, a thriving environment of programming languages, optimizers, and deployment tools helped make GPU-based platforms, such as NVIDIA's, the default choice for AI initiatives.¹⁷

The cumulative size of embedded AI models will soon swamp even the beefiest edge device.

In 2016, Google unveiled its own on-device machine learning chip, the tensor processing unit (TPU), built around a revolutionary matrix multiplier unit capable of hundreds of thousands of matrix operations (multiply and adds) per clock tick.¹⁸ In conjunction with the TensorFlow framework and associated tools, TPUs not only power Google's online services, such as Search, Street View, Photos, and Translate, but are offered to Google Cloud Platform customers as a service.

Other incumbents soon developed their own inference monsters. Intel announced its cloud-based Pohoiki Springs neuromorphic platform, which touts 100 million neurons in an array of its 128-core Loihi chip.¹⁹ In 2019, Amazon entered the fray with Inf1 cloud server instances based on its Inferentia inference acceleration chip.²⁰ Not to be left behind, Microsoft now offers a specialized AI cloud service based on the Graphcore intelligence processing unit,²¹ and Facebook is working on its own AI chip.²²

Exciting newcomers

Several recent start-up entrants in the custom AI chips market offer bespoke application-specific integrated circuit

(ASIC) and field-programmable gate array platforms optimized for blisteringly fast inference acceleration with surprisingly low power consumption. One emerging platform of note is Tartan AI, which began as a research project in 2014 at the University of Toronto and led to several techniques, including Laconic, TensorDash, and Shape-shifter. Tartan combines an inference optimization software stack run on a custom hardware platform and leverages an on-chip compression engine to squeeze out network sparsity, reduce

memory traffic, and yield a 40× inference execution acceleration and an order of magnitude power reduction across a wide selection of neural networks.^{23, 24}

The latest development for inference acceleration platforms comes as an offshoot of the optical quantum computing world. Silicon photonics uses silicon as an optical medium to guide photons, much like electronics uses conductors to guide electrical signals. However, unlike its electronics cousin, light serves as a nearly perfect medium for encoding and transmitting information at high speeds and with high fidelity, owing to its fundamental physical properties. One company applying silicon photonics to AI inference is Lightmatter, founded by scientists from the Massachusetts Institute of Technology's quantum computing labs. Its full AI stack includes a standard ASIC to control the optics and perform "housekeeping" activities, coupled with a silicon photonics processor that executes matrix operations at literally the speed of light. The platform supports common plug-ins and frameworks, such as PyTorch and TensorFlow. Still in development, early results show a tenfold increase in power

efficiency and similar speed growth in inference execution, and long-term expectations are even higher.²⁵

Although AI training techniques and smart models will continue to evolve for the foreseeable future, the focus is shifting from increasing accuracy to accelerating the speed of execution, adding brawn to the brains. During the coming years, specialized models, advanced optimization techniques, and new platforms will enable AI to bring solutions to ever-larger problems on ever-smaller devices. **■**

REFERENCES

1. K. Johnson, "OpenAI releases curtailed version of GPT-2 language model," *Venture Beat*, Aug. 20, 2019. [Online]. Available: <https://venturebeat.com/2020/03/20/best-cbd-oil/>
2. K. Johnson, "Nvidia trains world's largest Transformer-based language model," *Venture Beat*, Aug. 19, 2019. [Online]. Available: <https://venturebeat.com/2019/08/13/nvidia-trains-worlds-largest-transformer-based-language-model/>
3. K. Johnson, "Microsoft trains world's largest Transformer language model," *Venture Beat*, Feb. 10, 2020. [Online]. Available: <https://venturebeat.com/2020/02/10/microsoft-trains-worlds-largest-transformer-language-model/>
4. "AI and compute," OpenAI, San Francisco, 2018. [Online]. Available: <https://openai.com/blog/ai-and-compute/>
5. K. Rupp, "42 years of microprocessor trend data," Karl Rupp, 2020. [Online]. Available: <https://www.karlrupp.net/2018/02/42-years-of-microprocessor-trend-data/>
6. M. Sinnarkar, B. Jagtap, and S. Bau, "Edge computing market outlook: 2025," Allied Market Research, 2019. [Online]. Available: <https://www.alliedmarketresearch.com/edge-computing-market>
7. D. Reinsel, J. Gantz, and J. Rydning, "The digitization of the world from edge to core," IDC, Framington, MA, White Paper, 2018.
8. J. Toole, "Deep learning has a size problem," *Fritz AI*, 2019. [Online]. Available: <https://heartbeat.fritz.ai/deep-learning-has-a-size-problem-ea601304cd8>
9. F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5MB model size. 2016. [Online]. Available: [arXiv:1602.07360](https://arxiv.org/abs/1602.07360)
10. P. Ganesh, "Knowledge distillation: Simplified," *Towards Data Science*, Aug. 12, 2019. [Online]. Available: <https://towardsdatascience.com/knowledge-distillation-simplified-dd4973dbc764>
11. X. Jiao et al., TinyBERT: Distilling BERT for natural language understanding. 2019. [Online]. Available: [arXiv:1909.10351](https://arxiv.org/abs/1909.10351)
12. R. Singh, "Pruning deep neural networks," *Towards Data Science*, July 30, 2019. [Online]. Available: <https://towardsdatascience.com/pruning-deep-neural-network-56cae1ec5505>
13. J. Frankle and M. Carbin, The lottery ticket hypothesis: Finding sparse, trainable neural networks. 2018. [Online]. Available: [arXiv:1803.03635](https://arxiv.org/abs/1803.03635)
14. P. Zhao, X. Chen, Z. Qin, and J. Ye, "Model quantization for production-level neural network inference," *MXnet*, Apr. 16, 2019. [Online]. Available: <https://medium.com/apache-mxnet/model-quantization-for-production-level-neural-network-inference-f54462ebba05>
15. S. Prasanna, "Deep learning deployment with NVIDIA TensorRT," NVIDIA Deep Learning Institute, New York, 2019.
16. TensorFlow. *XLA: Optimizing Compiler for Machine Learning*. (2020).
17. S. Narasimhan, private communication, Mar. 5, 2020.
18. K. Sato, C. Young, and D. Patterson, "An in-depth look at Google's first Tensor Processing Unit (TPU)," Google, Mountain View, CA, 2017.
19. S. K. Moore, "Intel's neuromorphic nose learns scents in just one sniff," *IEEE Spectrum*, Mar. 16, 2020. [Online]. Available: <https://spectrum.ieee.org/tech-talk/artificial-intelligence/machine-learning/intels-neuromorphic-nose-learns-scents-in-just-one-sniff>
20. K. Quach, "Amazon sticks AI inference chip up for rent in the cloud for machine-learning geeks," *The Register*, Dec. 4, 2019. [Online]. Available: https://forums.theregister.co.uk/forum/all/2019/12/04/aws_reinvent_ai/
21. W. Knight, "Microsoft sends a new kind of AI processor into the cloud," *Wired*, Nov. 13, 2019. [Online]. Available: <https://www.wired.com/story/microsoft-sends-a-new-kind-of-ai-processor-into-the-cloud/>
22. S. Shead, "Facebook plans to develop its own AI chips," *Forbes*, Feb. 19, 2019. [Online]. Available: <https://www.forbes.com/sites/samshead/2019/02/19/facebook-plans-to-develop-its-own-ai-chips/#377017018379>
23. S. Sharify, private communication, Feb. 27, 2020.
24. S. Sharify et al., "Laconic deep learning inference acceleration," in *Proc. 46th Int. Symp. Computer Architecture*, 2019, pp. 304-317. doi: 10.1145/3307650.3322255.
25. N. Harris, private communication, Mar. 26, 2020.
26. R. Horev, "BERT explained: State of the art language model for NLP," *Towards Data Science*, Nov. 10, 2018. [Online]. Available: <https://towardsdatascience.com/bert-explained-state-of-the-art-language-model-for-nlp-f8b21a9b6270>

MARK CAMPBELL is the chief innovation officer for Trace3. Contact him at campbell@collabrial.com.



Managing Your Open Source Supply Chain—Why and How?

Nikolay Harutyunyan, Friedrich-Alexander University of Erlangen-Nürnberg

More than 90% of software products include open source components, most of which are not directly added by your own developers. Instead, they are an inseparable part of the software supply chains that virtually all companies depend on. This article covers the related risks of ungoverned open source use and provides industry best practices to practitioners.

players and much more. Companies use such components to address their product requirements for non-differentiating functionalities while focusing their internal development efforts on core differentiating features. By extension, software supply chains consist of multiple supplier tiers that all feed into each other and thus accumulate open source software that eventually gets into companies that sell complex products, such as original equipment manufacturers (OEMs).

A recent European Commission report estimated that using free/libre and open source software (FLOSS) saves the European economy roughly €114 billion per year directly and up to €399 billion per year overall.¹ FLOSS components are an essential part of software infrastructure ranging from operating systems and web servers to media

WHY FLOSS GOVERNANCE FOR SOFTWARE SUPPLY CHAINS?

While FLOSS is highly critical and relevant for industry, many companies are unaware that they use open source software, either disregarding it or delegating it to developers. This approach of ungoverned open source use carries a number of the following potential risks for companies:

- › legal risks caused by open source license noncompliance or incompatible licensing

Digital Object Identifier 10.1109/MC.2020.2983530
Date of current version: 4 June 2020

FROM THE EDITOR

Welcome back! This month's article on governing the open source software supply chain pulls together many different aspects from past articles into a comprehensive picture. A company needs to have an open source program office to coordinate all open source activities; it needs to understand the risks that stem from using open source, including open source hiding in components sourced from third-party suppliers; and it needs to actively manage those suppliers and their deliveries. Like all articles in this column, this one can stand alone, but I still recommend that you review past articles if you haven't done so yet to get the most out of this month's piece. Happy hacking, and be safe and healthy! — *Dirk Riehle*

- › financial risks resulting from preliminary injunctions (sales stop) or supplier replacement costs
- › technical risks stemming from the forced replacement of supplied open source software components.

As an example, let's look at the legal risk category. Companies have supplier contracts that provide some protection against the aforementioned legal risks. However, when it comes to open source license compliance,

such as additional costs for replacing the supplied software and technical resources for maintaining it in house (when possible).

FLOSS GOVERNANCE BEST PRACTICES FOR SOFTWARE SUPPLY CHAINS

Our analysis of open source governance expert interviews suggests an industry best practice of working with the supplier from the get-go to ensure open source governance at the supplier's site during software development, as opposed to a checkup upon software

FLOSS components are an essential part of software infrastructure ranging from operating systems and web servers to media players.

this protection is often overestimated. According to open source governance experts we interviewed, most companies at the end of the supply chain are much larger than their smaller suppliers. Once such a large company faces litigation over license noncompliance or copyright violation associated with FLOSS use² that stems from supplied code, it's possible but impractical to shift the legal responsibility to the supplier (or to a company further down the supply chain). If you adopt the latter strategy, you might end up running your smaller supplier out of business, which would create more problems,

delivery. Our data analysis suggests that some companies with an advanced understanding of supply chain FLOSS governance should focus their efforts on preventive steps, such as providing license checking and approval guidance during the development phase.

In a larger study of open source governance, we conducted and analyzed 21 expert interviews and FLOSS governance guidelines to learn about current industry best practices for using open source software in products.^{3,4} This article discusses a small subset of our findings, providing insights for the following best practice categories

related to supply chain open source governance:

- › the supply chain management (SCM) policy and process
- › preventive governance
- › corrective governance
- › bill-of-materials (BOM) management.

SCM POLICY AND PROCESS

SCM in terms of FLOSS governance is a complex and multifaceted task involving in-house software development teams, procurement offices, suppliers, and lawyers. Coordinating these stakeholders and ensuring a company-wide approach to open source governance is essential. Industry experts recommend that companies set up an SCM policy that strategically defines enterprises' FLOSS governance, informing all stakeholders.

Without a comprehensive policy, different parts of the company might apply dissimilar rules (or no rules at all) when dealing with open source software as part of the supplied code. For example, if a third-party software component is purchased to be used in a product, it is rarely checked for open source license compliance, especially in companies with little FLOSS governance awareness. Instead, firms rely on supplier contracts for any potential intellectual property issues, considering license noncompliance as one such concern. However, such clauses cannot guarantee that your products including open source components are license compliant. Instead, they can act only as a corrective measure if an issue is discovered by a customer, and even then, they are not a universal solution to the risks of ungoverned FLOSS use. Open source governance on the topic of SCM goes beyond supplier contracts and compliance checks, requiring a systematic approach and a company-wide policy.

In the course of our study, we found an industry best practice whereby an SCM policy should address governance aspects, such as

- › company goals for supplier management
- › metrics for efficient supplier management
- › recommendations for automating supplier management through tools
- › rules for suppliers that use open source components.

The policy should be defined by the open source program office to ensure a consistent approach to SCM within the company. It should be maintained, revised, and communicated through time. While the SCM policy defines the company's strategic take, it needs to be translated into the day-to-day processes of product development and software procurement. To achieve this, experts recommend operationalizing the policy through an SCM process.

The SCM process guides product managers, technical product managers, software developers, and others dealing with software supply chains. It also helps procurement managers and IT managers with external tasks related to SCM, such as dealing with supplier requests and contracts. The SCM process covers, among other things, assessing the open source governance maturity of a supplier, requesting supplier certification, and auditing suppliers.

The SCM process should be integrated into the daily workflows of the company. It should be easy to read, and it should be created in collaboration with the stakeholder engineers and managers. The process should be directly related to the tasks of software development and solve problems that engineers and managers face in their work when dealing with FLOSS governance.

PREVENTIVE GOVERNANCE

We found that industry experts recommend focusing on preventive open source governance. Companies should take steps to prevent supplier-related FLOSS governance issues. The initial preventive measure applies to choosing suppliers. We found a best practice

to choose the right supplier, taking into account the supplier's open source governance and compliance awareness and maturity. To do so, companies should design supplier contracts with open source governance aspects in mind and consider requesting supplier certification. Such certifications can be conducted internally (self-certification) or using existing standard certification frameworks for FLOSS governance in supply chains. A leading framework on the topic is being developed by the OpenChain Project.⁵

Another best practice for preventive FLOSS governance of software supply

chains focuses on supplier contracts. In certain cases, these contracts can include strict provisions, such as specific templates that suppliers must follow before any anticipated use of an open source component in the software development of the to-be-supplied code. A supplier would have to use the template to send open source component requests to the client for approval. The suppliers would also be encouraged to employ a similar practice with their own suppliers, which would, in turn, make the whole supply chain safer in terms of open source compliance.

CORRECTIVE GOVERNANCE

We found a number of industry best practices for addressing the issues of FLOSS governance and compliance caused by software supply chains. Going beyond the preventive measures, companies should also establish corrective open source governance in the context of SCM. Though preventive governance best practices mitigate the potential issues that result from lacking SCM, companies should be ready to address any cases of noncompliance as well as other issues caused by suppliers.

One expert recommendation is to conduct regular and surprise audits of software suppliers and their code to find potential issues, such as unintended open source licenses and missing copyright data. If risks are found, companies should proceed to mitigate them by assessing the threats' criticality and costs as well as by triggering supplier contract clauses and working with suppliers to take care of the issues, when possible. However, industry experts recommend against running suppliers out of business when conducting corrective governance, as the potential losses could

Industry experts recommend against running suppliers out of business when conducting corrective governance.

increase with the bankruptcy of a small supplier.

MANAGING BILLS OF MATERIALS

Most open source components end up in company products through software supply chains. Given the complex dependencies between open source components and libraries, as well as with companies' proprietary code,^{6,7} enterprises need to use systematic and consistent instruments to ensure the complete and transferable documentation of open source use that is introduced by their suppliers and their own developers. BOMs are such an instrument; however, they need to be extended beyond the traditional format of merely listing the software components of a product (supplied or own). To address the specifics of open source governance, BOMs must include additional metadata for open source components, such as accurate license information, versions, copyright details, and export-restriction tags.

Leading industry experts recommend using existing standards for BOM documentation and exchange within

software supply chains. The current leading standard is called the *Software Package Data Exchange (SPDX)*,⁸ which is an open standard for communicating software BOM information that enables the specialized documentation of open source component metadata. Using this format can be of high value to an OEM because doing so ensures full transparency when it comes to the open source use in products, including the awareness of FLOSS components originating in the supply chain.

It's an industry best practice to ask your suppliers for their software's BOM in the SPDX format, which can be checked and combined with the BOMs from other suppliers, eventually forming the BOM of an OEM's final product. As a consequence, an OEM would have an updated and ready BOM for its own products if a customer requested it. The experts we interviewed mentioned further benefits of the aforementioned approach, including the use of a machine-readable format compatible with most open source governance and compliance tools as well as

the method's industry-wide recognition as a leading standard.

During the course of our research on SCM in terms of FLOSS governance, we identified a number of industry best practices akin to the previously mentioned one. Putting some of these best practices together, we propose workflows or processes that practitioners can adjust and use in their companies. Figure 1 presents an example of such a workflow for BOM management. Starting with identifying the used FLOSS components and their metadata, companies should track and document this use, employing machine-readable exchange formats as well as ensuring the license compliance of and self-hosting backups for the used components.

This article presented a snapshot of our larger findings on the topic of open source governance,^{3,4} building upon our previous work on managing software supply chains in the context of FLOSS

governance.⁹ Going beyond the presented best practices from industry, it is crucial to use tools to automate various aspects of open source governance, such as license scanning and compliance checking, documenting open source components as part of product architecture, and managing BOMs. **□**

REFERENCES

1. European Commission, "The economic and social impact of software & services on competitiveness and innovation (SMART 2015/0015)," Luxembourg: Publications Office of the European Union, 2017. [Online]. Available: <https://op.europa.eu/en/publication-detail/-/publication/480eff53-0495-11e7-8a35-01aa75ed71a1>
2. H. Schoettle, "Open source license compliance: Why and how?" *Computer*, vol. 52, no. 8, pp. 63–67, 2019. doi: 10.1109/MC.2019.2915690.
3. N. Harutyunyan, "Corporate open source governance of software supply chains,"




FIGURE 1. A workflow for BOM management.

Ph.D. dissertation.
Friedrich-Alexander-Univ.
Erlangen-Nürnberg, 2019. [Online].
Available: <http://nbn-resolving.de/urn:nbn:de:bvb:29-opus4-122727>

4. N. Harutyunyan and D. Riehle, "Getting started with open source governance and compliance in companies," in *Proc. 15th Int. Symp. Open Collaboration*, 2019, pp. 1–10. doi: 10.1145/3306446.3340815.

5. OpenChain. Accessed on: Mar. 25, 2020. [Online]. Available: <https://www.openchainproject.org/>

6. A. Bauer, N. Harutyunyan, D. Riehle, and G.-D. Schwarz, "Challenges of tracking and documenting open source dependencies in products: A case study," in *Proc. 16th Int. Conf. Open Source Systems*, to be published.

7. T. Gustavsson, "Managing the open source dependency," *Computer*, vol. 53, no. 2, pp. 83–87, 2020. doi: 10.1109/MC.2019.2955869.

8. Software Package Data Exchange. Accessed on: Mar. 25, 2020. [Online]. Available: <https://spdx.org/>

9. D. Riehle and N. Harutyunyan, "Open-source license compliance in software supply chains," in

Towards Engineering Free/Libre Open Source Software (FLOSS) Ecosystems Impact Sustainability, B. Fitzgerald, A. Mockus, and M. Zhou, Eds. Singapore: Springer, 2019, pp. 83–95.

NIKOLAY HARUTYUNYAN is an open source researcher and postdoc at Friedrich-Alexander University of Erlangen-Nürnberg, Germany. Contact him at nikolay.harutyunyan@fau.de.



IEEE TRANSACTIONS ON BIG DATA

▶ SUBSCRIBE AND SUBMIT

For more information on paper submission, featured articles, calls for papers, and subscription links visit: www.computer.org/tbd

TBD is financially cosponsored by IEEE Computer Society, IEEE Communications Society, IEEE Computational Intelligence Society, IEEE Sensors Council, IEEE Consumer Electronics Society, IEEE Signal Processing Society, IEEE Systems, Man & Cybernetics Society, IEEE Systems Council, and IEEE Vehicular Technology Society

TBD is technically cosponsored by IEEE Control Systems Society, IEEE Photonics Society, IEEE Engineering in Medicine & Biology Society, IEEE Power & Energy Society, and IEEE Biometrics Council

Digital Object Identifier 10.1109/MC.2020.2993516





An Ambassador for Neural Networks

David Alan Grier, Djanghe, LLC

Neural networks were never a new part of computing, but at one point they needed a dedicated ambassador to move them into computing practice.

We should not be surprised to find a tutorial on artificial neural networks in the list of *Computer's* most influential articles. Neural networks are now a common subject in computer science, and they have been taught to two or three generations of undergraduates. They would be an obvious topic for a clear and well-organized article, such as the one written by Anil K. Jain, Jianchang Mao, and K.M. Mohiuddin. This article ranks number 30 on the list of the most influential articles that have been published by *Computer*. It has 5,128 downloads and 740 citations.

Yet, there are several surprises in the article that help us understand both its position in the body of knowledge

and the role of *Computer* in building that body of knowledge. Perhaps the prominent surprise is that the article has never been cited by another article in *Computer*, the publication that brought it to the field.

This fact, which seems anomalous in the technical literature, leads us to consider the role and position of artificial neural nets in the field of computing. They actually have a long and storied history that dates to the very first days of the field. You can find them, in a primitive form, in John von Neuman's 1945 technical report,

ARTICLE FACTS

- » Article: "Artificial Neural Networks: A Tutorial"
- » Authors: A.K. Jain, J. Mao, and K.M. Mohiuddin
- » Citation: *Computer*, vol. 29, no. 3, pp. 31–44, March 1996
- » Computer influence rank: #30 with 5,128 downloads and 740 citations



“First Draft Report on the EDVAC.”¹ We generally remember this report as describing the structure that we now call the von Neuman Architecture, yet von Neuman drew on the work of Warren McCulloch and Walter Pitts to show the connection between the structure of the brain and mechanical computation.² He found the connection so compelling that he returned to this idea in his final work, *The Computer and The Brain*, which he prepared as the Silliman Memorial lecture on science for Yale University.³

Even with its connection to von Neuman’s work, artificial neural networks would not be a major technology for the artificial intelligence research of the 1950s and 1960s. At that time, there was no good theory that described the properties of these networks or how they operated. Instead, researchers focused their attention on symbolic or search-based systems. Familiar examples include Newell, Saw, and Simon’s General Problem Solver;⁴ Simon and Feigenbaum’s Elementary Perceiver and Memorizer;⁵ Shortliffe’s Expert System Mycin;⁶ and Greenblatt’s chess program.⁷

Many writers claim that the work of Marvin Minsky and Seymour Papert pushed neural networks out of computer science temporarily, but this point of view misrepresents the work of Minsky and Papert and simplifies a complex period in the history of computing. In their 1969 book, *Perceptrons*, the two authors started to build a theoretical model for neural nets by looking at a perceptron, a linearized version of a general neuron. In their book, Minsky and Papert⁸ demonstrated the limitations of a single perceptron but, at the same time, laid the foundation for a general theory of neural networks and an important result. This result stated that as a tool for approximating functions, networks of perceptrons could approximate any mathematical function to any degree of accuracy.⁹

Building a theory of neural networks also included the work of re-conceptualizing these computing structures and placing them in a new context. By the mid-1970s, researchers had recognized that neural networks were classifiers and had properties that were similar to existing classification algorithms, such as those derived from statistical models. Paul Werbos was one of the great pioneers of this concept, and his great contribution was the back propagation algorithm.

times by a broad collection of journals. (Other citation sources give a number that is almost four times larger.) It has been cited by articles on medicine, tomography, communications, agriculture, chemistry, and encryption. It is repeatedly cited by articles that are presenting new classification algorithms.

Perhaps somewhat unusual for a *Computer* article, it is primarily cited by conference papers. Well over half of the cites in *IEEE Xplore* are to such papers,

The article, as Jain described it, was a simple, straightforward article “to inform readers with little or no knowledge of neural networks and help them understand detailed technical publications.”

Starting with his 1974 Ph.D. thesis, his work sparked a growing interest in neural nets and expanded the number of potential applications for them.¹⁰ In 1990, the IEEE created *IEEE Transactions on Neural Networks*, and six years later, the editors of *Computer* felt that it was time to do a special issue on the topic. They asked our authors to prepare a tutorial for that issue. “My research interest,” wrote lead author Jain, “was primarily in the field of statistical pattern recognition and data clustering.” As neural nets became more important, he “was fascinated by their generalization and feature extraction abilities.”

The article, as Jain described it, was a simple, straightforward article “to inform readers with little or no knowledge of neural networks and help them understand detailed technical publications.” As a result, it has become a classic in the field. It is regularly cited by papers that are first efforts to apply neural nets or machine learning to new applications. According to *IEEE Xplore*, it has been cited 740

which suggests its role as a tutorial. It has explained neural nets not only to researchers who work in computer science but also to those work in other fields. This role is suggested by the number of citations, which declined slightly around 2010 as deep networks and machine learning were becoming common in computer science but rose again as these technologies moved into other fields. Reinforcing this conclusion is the fact that the article is cited far more frequently by periodicals published by outside organizations than those published by the IEEE Computer Society. In fact, *IEEE Xplore* lists no citations of this article by any IEEE Computer Society magazine.

In the body of computing literature, this article by Jain, Mao, and Mohiuddin is an ambassador that takes the work of our members and delivers it to other technical communities around the world.


It has needed no other recognition in *Computer*.

Until now. 

REFERENCES

1. J. von Neumann, "First draft report on the EDVAC," *IEEE Ann. Hist. Comput.*, vol. 15, no. 4, pp. 27–75, 1993. doi: 10.1109/85.238389.
2. J. von Neumann, *The Computer and the Brain*. New Haven, CT: Yale Univ. Press, 1958.
3. W. S. McCulloch and W. Pitts, "A logical calculus of ideas immanent in nervous activity," *Bull. Math. Biophys.*, vol. 5, pp. 115–133, Dec. 1943. doi: 10.1007/BF02478259.
4. A. Newell, J. C. Shaw, and H. Simon, "Report on a general problem-solving program for a computer," in *Proc. Int. Conf. Information Processing*, 1960, pp. 256–264.
5. E. Feigenbaum, "The simulation of verbal learning behavior," in *Proc. IRE-AIEE-ACM '61 (Western)*, 1961, pp. 121–132. doi: 10.1145/1460690.1460704.
6. E. Shortliffe, "A rule-based computer program for advising physicians regarding antimicrobial therapy selection," in *Proc. Annual ACM Conf. (ACM '74)*, Jan. 1974, vol. 2, p. 739. doi: 10.1145/1408800.1408906.
7. R. Greenblatt, D. Eastlake, and S. Crocker, "The Greenblatt Chess Program," in *Proc. Fall Joint Computer Conf.*, 1967, pp. 801–810. doi: 10.1145/1465611.1465715.
8. M. Minsky and S. Papert, *Perceptrons: An Introduction to Computational Geometry*. Cambridge, MA: MIT Press, 1969.
9. G. Cybenko, "Approximations by superpositions of sigmoidal functions," *Math. Control, Signals, Syst.*, vol. 2, no. 4, pp. 303–314, 1989. doi: 10.1007/BF02551274.
10. P. Werbos, "Beyond regression: New tools for prediction and analysis in the behavioral sciences," Ph.D. dissertation, Dept. of Applied Mathematics, Harvard Univ., Cambridge, MA, 1974.

DAVID ALAN GRIER is a principal with Djaghe, LLC. He is a Fellow of the IEEE. Contact him at grier@gwu.edu.



IEEE TRANSACTIONS ON

COMPUTERS

Call for Papers: IEEE Transactions on Computers

Publish your work in the IEEE Computer Society's flagship journal, *IEEE Transactions on Computers*. The journal seeks papers on everything from computer architecture and software systems to machine learning and quantum computing.

Learn about calls for papers and submission details at www.computer.org/tc.

Digital Object Identifier 10.1109/MC.2020.2993517





CS CONNECTION

IEEE COMPUTER SOCIETY SPONSORS NEW JOURNAL ON ARTIFICIAL INTELLIGENCE

IEEE Transactions on Artificial Intelligence—financially cosponsored by the IEEE Computer Society (CS)—opened for submissions in April 2020 and will publish its inaugural issue in August. Led by founding editor-in-chief Hussein Abbass, the new journal will cover theories, methodologies, and applications in the field of artificial intelligence. For more information, including how to submit a paper,


Digital Object Identifier 10.1109/MC.2020.2986872
Date of current version: 4 June 2020

visit <https://cis.ieee.org/publications/ieee-transactions-on-artificial-intelligence>.

RICHARD E. MERWIN SCHOLARSHIP OPEN FOR APPLICATIONS

The CS is offering US\$40,000 in student scholarships, each US\$1,000 or more, to recognize and reward active student volunteer leaders in Student Branches or Chapters who show promise in their academic and professional efforts. This scholarship was created in honor of the late Richard E. Merwin, a past president of the CS, to recognize and reward student leadership.

The winners of this award will have the opportunity to serve as CS student ambassadors for the particular IEEE Region to which they belong. Duties as student ambassadors will include collecting and disseminating information to Student Branches or Chapters in their region and serving as a liaison between the CS Member and Geographic Activities Board and Student Members in their region. More than a dozen scholarships of US\$1,000 and up are available, lasting one academic year (approximately nine months).

Apply by 30 September 2020. For more information, visit www.computer.org/volunteering/awards/scholarships/merwin. 



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

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

EXECUTIVE COMMITTEE

President: Leila De Floriani; **President-Elect:** Forrest Shull; **Past President:** Cecilia Metra; **First VP:** Riccardo Mariani; **Second VP:** Sy-Yen Kuo; **Secretary:** Dimitrios Serpanos; **Treasurer:** David Lomet; **VP, Membership & Geographic Activities:** Yervant Zorian; **VP, Professional & Educational Activities:** Sy-Yen Kuo; **VP, Publications:** Fabrizio Lombardi; **VP, Standards Activities:** Riccardo Mariani; **VP, Technical & Conference Activities:** William D. Gropp; **2019-2020 IEEE Division VIII Director:** Elizabeth L. Burd; **2020-2021 IEEE Division V Director:** Thomas M. Conte; **2020 IEEE Division VIII Director-Elect:** Christina M. Schober

BOARD OF GOVERNORS

Term Expiring 2020: Andy T. Chen, John D. Johnson, Sy-Yen Kuo, David Lomet, Dimitrios Serpanos, Forrest Shull, Hayato Yamana
Term Expiring 2021: M. Brian Blake, Fred Douglass, Carlos E. Jimenez-Gomez, Ramalatha Marimuthu, Erik Jan Marinissen, Kunio Uchiyama
Term Expiring 2022: Nils Aschenbruck, Ernesto Cuadros-Vargas, David S. Ebert, William Gropp, Grace Lewis, Stefano Zanero

revised 1 May 2020

BOARD OF GOVERNORS MEETING

24 – 25 September 2020 in McLean, Virginia, USA

EXECUTIVE STAFF

Executive Director: Melissa A. Russell; **Director, Governance & Associate Executive Director:** Anne Marie Kelly; **Director, Finance & Accounting:** Sunny Hwang; **Director, Information Technology & Services:** Sumit Kacker; **Director, Marketing & Sales:** Michelle Tubbs; **Director, Membership Development:** Eric Berkowitz

COMPUTER SOCIETY OFFICES

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

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

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

IEEE BOARD OF DIRECTORS

President & CEO: Toshio Fukuda
President-Elect: Susan K. "Kathy" Land
Past President: José M.F. Moura
Secretary: Kathleen A. Kramer
Treasurer: Joseph V. Lillie

Director & President, IEEE-USA: Jim Conrad; **Director & President, Standards Association:** Robert S. Fish; **Director & VP, Educational Activities:** Stephen Phillips; **Director & VP, Membership and Geographic Activities:** Kukjin Chun; **Director & VP, Publication Services & Products:** Tapan Sarkar; **Director & VP, Technical Activities:** Kazuhiro Kosuge



Cybersecurity Analysis.

Protect Your Organization.



Advance Your Career with Cybersecurity

International Institute of Business Analysis™ (IIBA®) and IEEE Computer Society have partnered to provide a robust learning and certification program on what business analysis professionals need to know to be prepared for today's cybersecurity challenges.

Learn more about Cybersecurity Analysis at IIBA.org/cybersecurity.

