

COMPUTING edge

- Security and Privacy (Development)
- Edge Computing
- Digital Trust
- Ethics

SEPTEMBER 2025

www.computer.org



IEEE Computer Society

Grants for EMERGING TECHNOLOGY ACTIVITIES

MAKE AN IMPACT | CREATE SOLUTIONS

Are you connecting the computing community with emerging technologies? Help advance emerging tech to create solutions for the betterment of humanity.

Every year, we give up to **\$50,000** in funding per project for these efforts.

Learn more at

computer.org/communities/emerging-technology-fund



STAFF

Editor

Lucy Holden

Periodicals Portfolio Senior Managers

Carrie Clark and Kimberly Sperka

Director, Periodicals and Special Projects

Robin Baldwin

Production & Design Artist

Carmen Flores-Garvey

Periodicals Operations Project Specialists

Priscilla An and Christine Shaughnessy

Senior Advertising Coordinator

Debbie Sims

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

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

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

Reuse Rights and Reprint Permissions: Educational or personal use of this material is permitted without fee, provided such use: 1) is not made for profit; 2) includes this notice and a full citation to the original work on the first page of the copy; and 3) does not imply IEEE endorsement of any third-party products or services. Authors and their companies are permitted to post the accepted version of IEEE-copyrighted material on their own Web servers without permission, provided that the IEEE copyright notice and a full citation to the original work appear on the first screen of the posted copy. An accepted manuscript is a version which has been revised by the author to incorporate review suggestions, but not the published version with copy-editing, proofreading, and formatting added by IEEE. For more information, please go to: http://www.ieee.org/publications_standards/publications/rights/paperversionpolicy.html. Permission to reprint/republish this material for commercial, advertising, or promotional purposes or for creating new collective works for resale or redistribution must be obtained from IEEE by writing to the IEEE Intellectual Property Rights Office, 445 Hoes Lane, Piscataway, NJ 08854-4141 or pubs-permissions@ieee.org. Copyright © 2025 IEEE. All rights reserved.

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

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

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

2025 IEEE Computer Society Magazine Editors in Chief

Computer

Jeff Voas, *NIST*

Computing in Science & Engineering

Jeffrey Carver,
University of Alabama

IEEE Annals of the History of Computing

Troy Astarte,
Swansea University

IEEE Computer Graphics and Applications

Pak Chung Wong, *Trovaes and Bill & Melinda Gates Foundation (Interim EIC)*

IEEE Intelligent Systems

Bo An, *Nanyang Technological University*

IEEE Internet Computing

Weisong Shi, *University of Delaware*

IEEE Micro

Hsien-Hsin Sean Lee,
Intel Corporation

IEEE MultiMedia

Balakrishnan Prabhakaran,
University of Texas at Dallas

IEEE Pervasive Computing

Fahim Kawsar, *Nokia Bell Labs and University of Glasgow*

IEEE Security & Privacy

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

IEEE Software

Sigrid Eldh, *Ericsson, Mälardalen University, Sweden; Carleton University, Canada*

IT Professional

Charalampos Z. Patrikakis, *University of West Attica*

SEPTEMBER 2025 • VOLUME 11 • NUMBER 9

COMPUTING
edge



8

On the
Measurability
and Testability of
IT Security

16

Advancing Data
Security and
Sustainability:
Establishing a
Circular Economy
for Storage

22

On Causality
in Distributed
Continuum
Systems

Security and Privacy (Development)

8 On the Measurability and Testability of IT Security

ANDREAS GRÜNERT, JAMES BRET MICHAEL, ROLF OPPLIGER,
AND RUEDI RYTZ

16 Advancing Data Security and Sustainability: Establishing a Circular Economy for Storage

JONMICHAEL HANDS AND TOM COUGHLIN

Edge Computing

22 On Causality in Distributed Continuum Systems

VÍCTOR CASAMAYOR PUJOL, BORIS SEDLAK, PRAVEEN KUMAR DONTA,
AND SCHAHRAM DUSTDAR

30 FlyNet: Drones on the Horizon

ALICIA ESQUIVEL MOREL, CHENGYI QU, PRASAD CALYAM, CONG WANG,
KOMAL THAREJA, ANIRBAN MANDAL, ERIC LYONS, MICHAEL ZINK,
GEORGE PAPADIMITRIOU, AND EWA DEELMAN

Digital Trust

40 Labeling “Things”

JOANNA F. DEFRANCO AND PHIL LAPLANTE

44 Can We Explain Privacy?

GÖNÜL AYCI, ARZUCAN ÖZGÜR, MURAT ŞENSOY, AND PINAR YOLUM

Ethics

50 How to “Sell” Ethics (Using AI): An Interview With Alexander Serebrenik

TIM MENZIES

54 Ethics: How Far Have We Come?

BRITTANY JOHNSON AND TIM MENZIES

Departments

4 Magazine Roundup

7 Editor’s Note: Developing Secure and Sustainable Systems

69 Conference Calendar

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



Magazine Roundup

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

Computer

Knowledge Monopoly Risks in Generative AI-Assisted Software Development Lifecycle

This article, featured in the July 2025 issue of *Computer*, examines the emerging knowledge monopolization risks in generative artificial intelligence (AI)-assisted software development, proposing a comprehensive assessment framework and practical intervention strategies to help organizations maintain technical sovereignty while leveraging AI capabilities for development efficiency.

Computing

Accelerating Innovative Energy Solutions Using Combustion Simulations

Combustion-based transportation, electricity generation, and industrial heating in manufacturing constitute the three largest sectors of energy demand. Several teams at the National Renewable Energy Laboratory have been actively advancing research in these areas by leveraging

computational modeling of combustion processes across the heavy-duty land-based transportation, aviation, and power generation sectors. This January–March 2025 *Computing in Science & Engineering* article summarizes some of these efforts, demonstrating the potential of advanced computational techniques to generate technological solutions that will transform the global energy system.

IEEE Annals

of the History of Computing

Neat Versus Scruffy:

How Early AI Researchers Classified Epistemic Cultures of Knowledge Representation

Accounts of the history of symbolic artificial intelligence (AI) often categorize researchers and their programs through binary oppositions: logicist versus procedural, or alternatively, neat versus scruffy. These terms were often referenced throughout the 1970s and 1980s by AI researchers at institutions in the United States and Europe to document epistemic and esthetic differences in approaches to knowledge representation. In this article, featured

in the April–June 2025 issue of *IEEE Annals of the History of Computing*, the author historicizes and situates these binary oppositions, showing how researchers leveraged them in discourse to demarcate the beliefs and commitments underpinning certain AI approaches.

IEEE Computer Graphics AND APPLICATIONS

Voting-Based Intervention Planning Using AI- Generated Images

The continuous evolution of artificial intelligence and advanced algorithms capable of generating information from simplified input creates new opportunities for several scientific fields. Currently, the applicability of such technologies is limited to art and medical domains, but it can be applied to engineering domains to help the architects and urban planners design environmentally friendly solutions by proposing several alternatives in a short time. This article, featured in the March/April 2025 issue of *IEEE Computer Graphics and Applications*, utilizes the image-inpainting algorithm for suggesting several



alternative solutions to four European cities.

IEEE Intelligent Systems

The Next Wave of AI for Social Impact: Challenges and Opportunities

The burgeoning field of artificial intelligence for social impact (AI4SI) represents a significant evolution in artificial intelligence, prioritizing measurable positive impact for vulnerable and under-resourced populations. This article, which was in the May/June 2025 issue of *IEEE Intelligent Systems*, examines the historical context and recent surge in AI4SI, driven by technological advancements and a growing awareness of societal challenges. It highlights the crucial role of interdisciplinary collaboration, ethical considerations, and the potential of emerging AI trends in addressing issues such as poverty, health, and environmental sustainability.

IEEE Internet Computing

Cognitive Digital Twins for the Microgrid: A Real-World Study for Intelligent Energy Management and Optimization

Digital twin (DT) technology is a promising solution for achieving

optimized microgrid control with enhanced efficiency, reliability, and sustainability. The authors of this article from the January/February 2025 issue of *IEEE Internet Computing* focus on a real-world microgrid in Singapore and develop a cognitive DT. They demonstrate the effectiveness of their DT in enabling real-time optimization and management of microgrid operations, paving the way for technology adoption in smart grids to achieve improved grid resilience and efficiency.

IEEE micro

Qualcomm Oryon CPU in Snapdragon X Elite: Micro-Architecture and Design

This article, featured in the May/June 2025 issue of *IEEE Micro*, presents the micro-architecture and design of the Qualcomm custom CPU, named Qualcomm Oryon CPU, that was introduced in 2024 in the Qualcomm Snapdragon X Elite system on a chip for the client computing market. It describes the micro-architecture of the CPU core and its cache and memory subsystem and is illustrative of a modern high-performance CPU with best-in-class energy efficiencies that is designed to be scalable across different product categories and price points.

IEEE MultiMedia

Terrain Segmentation Network in Wild Environments With Hybrid Plus Downsampling

Existing segmentation networks primarily use single downsampling to extract low-resolution semantic information, which may not adapt well to features of different scales, leading to information imbalance and distortion. The authors of this January–March 2025 *IEEE MultiMedia* article propose a hybrid plus downsampling method to address this issue. They also propose a terrain segmentation network (TSNet) for safe navigation of mobile robots in wild environments. Extensive experimental results on the wild datasets demonstrate that TSNet outperforms other state-of-the-art methods in recognizing wild unstructured terrain.

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

EtherealBreathing: A Holographic Biofeedback Game to Support Relaxation in Autistic Children

The authors of this article in the October–December 2024 issue of *IEEE Pervasive Computing* evaluate a novel biofeedback holographic

game, *EtherealBreathing*, designed to support autistic children. In *EtherealBreathing*, children practice box breathing to collect virtual elements to maintain the Earth's balance, using a wearable sensor to measure chest expansion for breath detection.

IEEE SECURITY & PRIVACY

Characterizing E-Commerce Harm by Investigating Online Communities: A Case Study With Abusive Dropshipping

E-commerce websites are targets of abusive individuals, though it is difficult to understand the methods and tools these individuals employ. The authors of this article, featured in the May/June 2025 issue of *IEEE Security & Privacy*, introduce a methodology and case study that leverage online communities as data sources to identify and analyze harmful activities.

IEEE Software

BPMN-LLM: Transforming BPMN Models Into Smart Contracts Using Large Language Models

"Law is code" is pivotal for advancing the intelligent judiciary. This article from the July/August 2025 issue of *IEEE Software* proposes a business process modeling notation-large language model (BPMN-LLM), which transforms BPMN models of legal contracts (LCs) into smart contracts (SCs) using LLMs in a user-friendly and cost-effective manner.

IT Professional

Securing the Industrial Internet of Things: A Comprehensive Digital Forensic Readiness Framework and Cybersecurity Approach

In the Industrial Internet of Things (IIoT), diversity exists in the operational and digital devices that work in tandem to deliver industry services. These systems, while ushering in efficiency, also introduce vulnerabilities that render them susceptible to cyberattacks. This article, featured in the May/June 2025 issue of *IT Professional*, focuses on digital forensic readiness (DFR) for the IIoT environment. It entails the identification

of challenges, development of a meticulously tailored DFR framework for IIoT networks and its alignment with an attack model, and the formulation of a comprehensive data artifact template. 🌐

**Join the IEEE
Computer
Society**

computer.org/join

IEEE DataPort™ STORE, SEARCH & MANAGE RESEARCH DATA

Individual subscriptions to IEEE DataPort are free for all IEEE society members and Young Professionals. Just log in and activate your subscription for unlimited access to datasets, data management tools, dataset storage for your own research, and more.





Editor's Note

Developing Secure and Sustainable Systems

While security and privacy systems have grown more advanced, so have the risks posed by cyber threats. This results in a lack of consumer trust in the security of IT products and services. There is also a concern that explosive data growth is outstripping storage and energy capabilities in security systems, creating unsustainable technology use. This issue of *ComputingEdge* explores testing techniques to make IT security more effective and dependable, creating sustainable and secure data systems through media sanitization, and addressing problems with digital trust related to Internet of Things (IoT) devices and privacy assistants. The articles further outline developments in edge computing and ethics in software engineering.

IT and data security must be tested, trustworthy, and sustainable. In *Computer* article "On the Measurability and Testability of IT Security," the authors consider if and how IT security can be measured and tested. *Computer* article

"Advancing Data Security and Sustainability: Establishing a Circular Economy for Storage" demonstrates how a well-defined media sanitization strategy can support both rigorous data protection and sustainable resource utilization.

Next on the horizon in the field of edge computing are expansions in distributed systems and enhanced data transport through a platform that also improves unpiloted aerial vehicle (UAVs) research. The authors of "On Causality in Distributed Continuum Systems," from *IEEE Internet Computing*, advocate for prioritizing causality and equity in DCS applications. In *IEEE Internet Computing* article "FlyNet: Drones on the Horizon," the authors present FlyNet, a workflow management system to support scientific research UAVs and other edge-to-cloud applications, such as data transport and service quality.

Privacy concerns associated with Internet of Things (IoT) devices as well as automated assistants need to be addressed to improve transparency, explainability, and

digital trust. In "Labeling 'Things,'" from *Computer*, the authors explain why labels are needed to increase awareness and the privacy risks involved in using devices such as electronic doorbells and home security cameras. The authors of "Can We Explain Privacy?" from *IEEE Internet Computing*, give an overview of privacy—what it is, what it means to users, and how automated privacy assistants can explain privacy decisions for greater transparency and understanding.

While there have been many achievements in software engineering ethics during the last decade, there's still a long way to go. *IEEE Software* article "How to 'Sell' Ethics (Using AI): An Interview With Alexander Serebrenik" discusses how to promote ethics in software engineering—by emphasizing what it does and how that can improve code and productivity. In *IEEE Software* article "Ethics: How Far Have We Come?" the authors reflect on advancements in ethics in recent years as well as areas that need improvement. 🌍

DEPARTMENT: CYBERTRUST

On the Measurability and Testability of IT Security

Andreas Grünert, *National Cyber Security Center*James Bret Michael , *Naval Postgraduate School*Rolf Oppliger  and Ruedi Rytz, *National Cyber Security Center*

In this article we consider the compound question of whether, and if so, how IT security can be measured and tested.

Even as best practices for assurance of systems continue to evolve, there are consumers of information technology-enabled products and services who are still dissatisfied with the level of dependability afforded by, and trust that can be placed in, those products and services. One aspect of assurance is testing, which involves (possibly multidimensional) measurements in which the test cases are deemed to have been passed if certain measured values are achieved. In this article, we consider the compound question of whether, and if so, how IT security can be measured and tested. As explained in the article, there are two approaches to assurance for security, verifying, and falsifying, both of which have their frailties and are therefore not convincing in their own right. At best, the two approaches can be used in combination to achieve the best possible, if not meaningful, measurement and test results.

SECURITY

To answer the previous question, we need to clarify what we mean by “security,” but this is difficult to do because there only exist context-dependent definitions for the term. Definitions of security, like those for safety, convey a notion of attaining some level of freedom from harm in the presence of threats and vulnerabilities over a particular interval of time, or that a stakeholder (for example, developer, user, owner, regulator) is sufficiently well prepared for all possible

security-related eventualities above a threshold of acceptable risk and can react effectively to minimize harm through risk-mitigation measures. However, what constitutes acceptable risk and being adequately prepared to react depends not only on the situation, but also on the observer. This means that the concept of security is subjective and depends on the observer’s perception and willingness to accept some level of risk, meaning that there is no precise definition or risk-taking threshold that is the same for everyone.

If you have a concept such as security, which is subjective and cannot be defined precisely and consistently for everyone, measurement or testing—according to the maxim “if you can’t define it, you can’t measure it”¹—is not possible or only possible with restrictions. The difficulty arises from the fact that although a metric or test criteria would be required that are clear and unambiguous, these cannot be found due to the imprecise definition. We are also familiar with this problem from other areas of everyday life: Because the term “intelligence” can be defined in different ways, there are also many different intelligence measurement and testing methods with their own metrics. In addition to the Turing test, in which an automaton is considered intelligent in the context of artificial intelligence if it cannot be distinguished from a human in a dialog, there are many other test procedures, many of which provide an intelligence quotient (IQ) as a result. However, an IQ also has advantages and disadvantages and is therefore the subject of controversial debate among experts. In any case, as a metric for a person’s intelligence, it is to a certain extent arbitrary and therefore controversial.



For a term that cannot be defined precisely and consistently for everybody, it seems difficult to establish metrics or test criteria that are neither arbitrary nor uncontroversial. Nevertheless, one can at least try to find properties and features that are characteristic of the term and, to a certain extent, measurable or testable. In the following, such properties and characteristics are referred to as “verifiable.” With regard to intelligence, for example, mathematical and logical skills and reasoning abilities can be measured and tested to a certain extent and are therefore also verifiable, while there are other properties and characteristics that also represent a form of intelligence, but which cannot be measured or tested and are therefore not verifiable.

Although there are many terms that can be defined with sufficient precision using verifiable properties and characteristics and are therefore themselves verifiable, there are other terms for which this is not so easily feasible, and which are therefore not verifiable either. Let’s take the example of a person’s health: If it were verifiable, it could somehow be measured by a doctor. We would then have something like a health quotient, although we could still argue about the sense and nonsense of such a value. In any case, such a value would represent a person’s health. However, we have neither a health quotient nor any other meaningful way of measuring a person’s health.

If we take the example further and ask how a doctor can proceed to provide information about a person’s health or state of health, two possibilities arise that are based on different definitions of the concept of health:

- If health is defined by properties and characteristics that correspond to standard values for certain measurements and tests, and if a person is said to be healthy as long as his or her measured values (for example, pulse, blood pressure, blood count, electrocardiogram, urine tests) do not deviate too much from the standard values,

the hypothesis that the person is healthy can be confirmed in the positive case and thus his or her health can be *verified*.

- If instead we define health as the absence of illness and say that a person is healthy if no (serious) illness is diagnosed, we can only refute and thus *falsify* the (assumed) health. To do this, the doctor must only diagnose a single illness.

In the first case, health is verifiable, but the definition of health is no longer comprehensive, that is verifiability is then at the expense of the quality and comprehensive validity of the definition (a person can be ill but still be considered healthy as long as the illness is not covered by the measurements and tests). In the second case, on the other hand, a comprehensive definition is assumed, but health itself is no longer verifiable. Strictly speaking, in this case the doctor would have to proceed by exclusion and try to rule out as many illnesses as possible. Specifically, the doctor would have to go through a list of possible illnesses and try to rule out each of these illnesses by means of suitable measurements or tests. If a disease cannot be ruled out and a corresponding diagnosis is confirmed, the person is considered ill or unhealthy. This refutes

DISCLAIMER

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Swiss National Cyber Security Center or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright annotations thereon.

or falsifies the (assumed) health. There are two problems to consider here: On the one hand, there are diseases that cannot be clearly diagnosed according to current knowledge, and on the other hand, the number of possible diseases is simply very large, and, above all, not exhaustive, that is new diseases can be added at any time. (Note that it is not possible to draw up a list of verifiable diseases even if all the diseases known in the medical literature are compiled, because there is always a possibility, albeit a small one, that new, previously unknown diseases will emerge—the Coronavirus disease (COVID-19) has brought this lesson home to us again.) Both problems make it difficult for a doctor to make conclusive statements about a person's health or state of health.

BECAUSE THE TERM "INTELLIGENCE" CAN BE DEFINED IN DIFFERENT WAYS, THERE ARE ALSO MANY DIFFERENT INTELLIGENCE MEASUREMENT AND TESTING METHODS WITH THEIR OWN METRICS.

There are basically two approaches to measuring or testing a person's health: A verifying approach, which could also be called constructive, because one tries to find a simple and accessible definition for the concept of health, and a falsifying approach, which could also be called destructive, because in the context of a measurement or test, one only tries to disprove the assumed health. It should be noted that the two approaches are not mutually exclusive. For example, a physician can use both approaches in combination to be able to make a meaningful statement within a reasonable amount of time. In any case, health is a concept that is difficult to handle from the point of view of measurability and testability, and unfortunately security also belongs to this category of concepts. Other examples not further elaborated here would be concepts like peace or justice.

The following sections discuss verifying (or constructive) and falsifying (or destructive) approaches that can be used to attempt to measure or test security in IT. With this, we turn away from health and toward IT security.

VERIFYING APPROACHES

In a verifying approach, a simpler definition must be constructed for the term under discussion, such as IT security, so that verification is feasible. One obvious possibility, for example, is to replace IT security with one or more other properties or characteristics that can be measured or tested and thus verified. Think of the classic CIA triad, in which IT security is replaced by the three properties of confidentiality, integrity, and availability. An IT system is described as secure if it protects the confidentiality, integrity, and availability of the stored, processed, and transmitted data. This seems to be more measurable or testable than the more comprehensive but more difficult-to-handle concept of IT security.

The second part of the Open Systems Interconnect reference model² provides another example. Here, a network is postulated to be secure if it implements and provides a set of security services, such as authentication, data confidentiality and integrity, and access control, with the help of suitable security mechanisms. Because the security services can be verified individually, it is also the security of a network. However, this has redefined the concept of security for networks: A network is no longer secure if it is not exposed to serious dangers, threats, and risks, but rather if it is able to provide certain services (it goes without saying that the services must continue to be offered even after verification, that is, the additional problem of continuity-of-service provision arises). In the analogy of a person's health, it would be like saying that a person is healthy as long as their breathing and heartbeat are functioning. In medicine, we know that this is oversimplified, and that health encompasses more than just the functioning of breathing and heartbeat. In the realm of network security, we also know that it is not enough to offer certain security services, but that it is important how exactly this is done. Reducing security to the provision of specific security services is too simplistic and hardly does justice to the problem. Accordingly, ISO 7498-2² has not yet established itself as a standard in the field.

A similar verifying approach is taken in the evaluation and certification of IT products: Based on the realization that the market for IT products is a "lemon market" in the sense of George A. Akerlof³ (a "lemon" is defined as a "Monday car" or a bad used car; the

market for used cars is characterized by an asymmetry of knowledge; while the seller likely knows the true condition of the car, this is difficult for the buyer to assess; and Akerlof has shown that prices erode in such markets, that is, comparatively low sales proceeds are achieved for both “lemons” and good quality used cars) and the knowledge that the quality of goods (in this case IT products) is eroding in such markets, many countries have been trying to intervene in a regulatory capacity since the early 1980s and enable and promote the evaluation and certification of IT products by independent and state-accredited bodies. Several criteria catalogues bear witness to this fact, such as follows:

- › Trusted Computer System Evaluation Criteria (TCSEC, also known as the “Orange Book”) in the USA⁴
- › Canadian Trusted Computer Product Evaluation Criteria (CTCPEC) in Canada
- › Information Technology Security Evaluation Criteria (ITSEC) in Europe and, in particular, France, Germany, the United Kingdom, and The Netherlands.⁵

Ultimately, these developments led to the Common Criteria for Information Technology Security Evaluation (referred to simply as “Common Criteria” and abbreviated as CC), which have been used since the turn of the millennium to evaluate the security of IT products and their components and to certify them for the international market on this basis.⁶ Here too, attempts are made to define security via verifiable properties and characteristics, such as functionality classes and trustworthiness (in terms of correctness of implementation and depth of testing).

However, neither the TCSEC, CTCPEC, or ITSEC, nor the CC have established themselves in the field of evaluation and certification, and the number of certified IT products is still comparatively small. In addition, certifiable products and components tend to be designed for individual functions and in this sense are not particularly complex, such as smartcards and smartcard readers or sensors. As soon as a product reaches a certain level of complexity, evaluation and certification is hardly possible and would, in principle, have to be repeated for each new release, at least in some

simplified form. (Continuous testing throughout the life cycle of the product would be desirable. This is not practicable, and the testing must be discrete regarding the testing times.) This is not only time-consuming and expensive, but also delays the publication of this release. Ultimately, there is no guarantee that a certified product has no vulnerabilities or weaknesses that could allow the product to be compromised under certain circumstances. In addition, security stands and falls with the configuration and operation of the product. Thus, any secure product can be configured and operated insecurely, and any insecure product can be configured and operated securely in a certain way (for example by operating it virtualized in a sandbox without interaction with other products). Accordingly, a CC or other certificate does not say much about how securely a product is used in an application context. To put it in the words of Robert H. Courtney: “You cannot say anything interesting (that is, significant) about the security of a system except in the context of a particular application and environment.”⁷ Accordingly, certificates for IT products without further information about the intended areas of application are illusory or even deceptive. Against this background, it will be interesting to see if and how successful the product certification program as required under the European Union’s Cyber Resilience Act (CRA)⁸ will be.

Finally, a verifying approach is also being attempted in the area of organizational security and IT-related business processes: If an organization wants to present itself as secure and acting in accordance with best practices in the area of IT security, the establishment, implementation, maintenance, and continuous improvement of a documented information security management system (ISMS) in accordance with ISO/IEC 27001⁹ is an option. To achieve greater visibility, such an ISMS can even be certified. Security is reduced to the existence of an ISMS with corresponding security control measures, that is, it is assumed that an organization is acting securely if it can rely on such an ISMS (without considering whether the security control measures in detail are actually suitable for fulfilling the organization’s security requirements). Whether this assumption is justified and whether security correlates with the existence of an ISMS in accordance with ISO/IEC 27001 is not clear a priori and requires further empirical investigation.

Between pure product audits in accordance with CC and audits focusing on an ISMS in accordance with ISO/IEC 27001, there are also audits in which it must be transparently demonstrated that certain security measures have been effectively implemented. Such audits typically refer to security standards and best practices such as the payment card industry, data security standard (PCI-DSS)¹⁰, the IT general controls (ITGC), or the requirements for a minimum viable secure product (MVSP) used in development. The success of this type of audit varies depending on the industry and level of commitment. However, it always increases transparency and thus improves assurance of security.

In summary, it can be said that although there are now several verifying approaches for measuring or testing IT security for both products (CC certificates) and services and organizations (ISO/IEC 27001 certificates), none of these approaches have gained widespread acceptance to date. Where such certificates are used, there are usually also corresponding regulatory requirements, all of which are aimed at influencing the trustworthy behavior of an organization and thus supporting customers. In addition to the difficulty of finding simplifying definitions with properties and features characteristic of IT security, the main problems with these approaches are that certifications are complex and expensive, and that not only the market for IT products but also that for certificates is a "lemon market," that is, there is a risk that the quality of the certificates themselves will erode over time. The only options available to influence and at best prevent such erosion are, on the one hand, state accreditation of the approved certification bodies and, on the other, regulation of the liability of these certification bodies. If a certificate can be issued without the simultaneous assumption of liability claims, then it appears difficult to maintain a high level of quality. In addition, the market for certificates is organized internationally, so that the control options of individual states are limited by default.

FALSIFYING APPROACHES

A falsification approach does not attempt to define IT security using verifiable properties and features. Instead, an attempt is made to show that an IT system under discussion can be compromised in a certain

way. This rather proves insecurity and thus disproves security. Such a (falsifying) approach is therefore basically designed for a test, whereby the test consists of the fact that the IT system cannot be compromised. At best, a metric could be derived from the difficulty of compromise, but such a metric would also be arbitrary to some extent.

In IT security, the system to be tested must therefore be compromised as part of a falsification approach. To do this, suitable attack vectors must be found and put together in such a way that a compromise or breach results. This is a heuristic and creative activity that largely depends on the expert knowledge and skills of the tester. Instead of a predefined list of properties and characteristics that can be checked individually, such a list is created in an ad hoc manner (that is, during the test). Although this has the advantage that a test is not predictable, the result depends on the tester and is to a certain extent arbitrary. This is not good, but is the best possible solution in many areas. A correspondingly large number of offers for such security tests are available on the market. Think of ethical hackers, penetration testers, bug bounty hunters, and red teams. They all offer audits and tests that can differ greatly in terms of methodology, effort, and intensity. Accordingly, even after a test, it is not clear whether other testers would have found the same or different attack vectors that could have compromised the object. Due to the strong dependence on the testers, objectification in the form of certification is not really feasible or is replaced by the reputation of the testers. If, for example, a testing company enjoys a great reputation in the security sector, then the confirmation that the company has not found any serious vulnerabilities and weaknesses will be just as meaningful and valuable in practice as a certificate, even if the significance cannot be directly validated.

In summary, it can be said that there are many falsifying approaches to measuring or testing IT security, especially for products and services and to a lesser extent for organizations. In addition to the high costs involved, the main problem with these approaches is the arbitrariness of the results and the associated dependencies on the testers' knowledge, skills, and abilities. Whether you can and want to live with this also depends on the industry in question. While security tests are widely used and accepted in

agile software development, for example, they have been rather unthinkable for suppliers to the aircraft industry. For instance, an aircraft manufacturer would hardly use a software module for which a randomly selected company commissioned with the testing has not found a vulnerability or weakness. In this industry, more assurance of IT security is usually required than can be achieved with such a security test. Instead, formal and verifying approaches will be required in this industry and falsifying approaches will at most be used in a complementary manner.

The explanations have shown that security is a concept that not only cannot be precisely defined, but strictly speaking can only be falsified. However, because a lack of falsification does not allow any statement to be made about security (an IT system that is not proven to be insecure is not necessarily secure), verifying and falsifying approaches are routinely used in combination in practice.

- › Verifying approaches deliver the desired results in terms of objectivity and traceability, and in this sense are also the declared aim of scientific endeavors (also in the field of IT security).
- › In contrast, falsifying approaches are based on security tests with unclear or difficult-to-interpret results. In this area, iterative improvements are often made and retested until further tests must be dispensed with for resource reasons. This is a “cat-and-mouse” or “cops-and-robbers” game that cannot be used to make reliable and scientific claims about effective IT security.

The bottom line is that verifying approaches (and therefore also CC certificates) are only possible for small and simple products or components for which a comprehensive and meaningful set of verifiable properties and characteristics can be defined. Above a certain size and complexity, this is no longer possible, so that falsifying security tests are required. To limit the arbitrariness of such tests, there is a need for security tests that are as objective or even standardized as possible. The European Standard EN 1143-1:2019¹¹ provides such a standard for secure storage units. Although this is difficult in IT security

because the attack vectors are as numerous as they are diverse, the question nevertheless arises as to whether a certain objectification or standardization of security tests would not also be possible and sensible. However, standardized security checks also have the disadvantage that attackers then know which attacks do not work and no longer need to mount them to compromise a product. This is certainly an argument against standardization in this area.

Due to their complexity, IT security audits of organizations (that is, their IT-dependent business processes) are generally more difficult. The ISO/IEC 27001 certificates issued in this area are based on the testing of technical and organizational characteristics, whereby the auditors must rely primarily on the organizational criteria to ensure that the information presented to them corresponds to the truth and that documented processes are also lived in this way. This means that purely technical audits are replaced or complemented by trust. The corresponding certificates are primarily driven by regulatory requirements and are subject to the danger that they will develop into a mass business and the quality will move downwards in favor of lower prices. We have already experienced this in the realm of ISO 9000 quality management standards. After the initial euphoria surrounding the topic, it has subsequently lost its significance. In this respect, certificates often develop as follows: As long as they are rare, they can serve as a distinguishing purpose and appear interesting (particularly for product differentiation). However, as soon as they have become widely accepted, they lose this advantage for companies and therefore also gradually lose their significance. How the state can meaningfully intervene and possibly even control this situation is an open question. This also applies to the liability issues affecting certification bodies. So far, the scope for states to exert influence beyond the accreditation of certification bodies has proven to be modest. Experience with the European Union’s Cyber Resilience Act⁸ will show whether the situation has improved in the meantime. Without effective and convincing options for measurement or testing, IT security will remain a matter of trust in many respects.

The situation can be summarized as follows: The complexity of a test object (IT product or organization) is inversely proportional to the possibility of defining

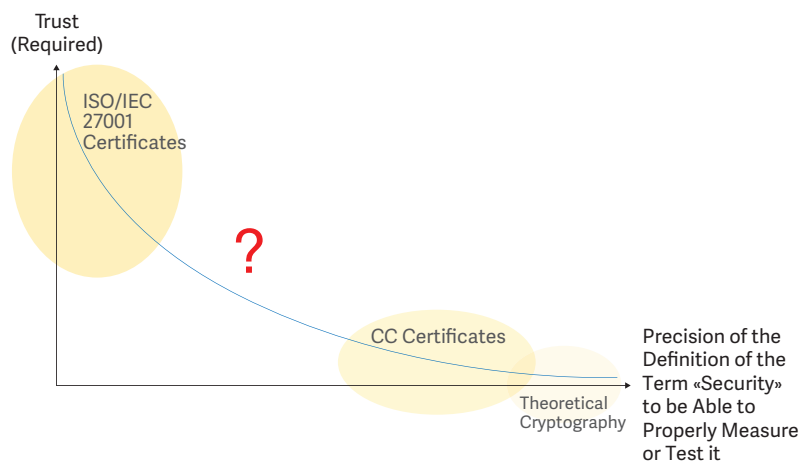


FIGURE 1. Testability curve.

its security, that is, the more complex the object is, the more difficult it will be to find a precise security definition for it and, as a result, a sufficiently large number of suitable and measurable or testable characteristics. In the field of theoretical cryptography, for example, we have very precise definitions for the security of certain cryptosystems and can therefore make meaningful measurements or tests. Nonetheless, the more these cryptosystems leave the “laboratories” and are used in the real world, the more difficult both the security definitions and the corresponding measurements or tests become. In addition to the purely mathematical properties of the cryptosystems, questions of implementation, key generation, and management, as well as user guidance, play a crucial role in security, and these questions are anything but easy to answer. Accordingly, we have conceptually simple test objects whose security can be defined reasonably precisely and measured or tested accordingly (for example, in the context of CC certificates), and more complex test objects whose security cannot be defined precisely and therefore cannot be measured or tested as meaningfully (for instance, in the context of ISO/IEC 27001 certificates). The less precise the definition and the less meaningful a measurement or test is, the greater the proportion of trust that must be placed in the security of the object (IT product or organization). This situation is illustrated schematically in the testability curve shown in Figure 1. It shows the required trust as a function of the precision of the definition of the term “security” as the basis for a measurement or test.

The big question concerns the measurement or testing of the security of larger IT products for which (due to their size and complexity) the CC are no longer applicable, but for which more is needed in terms of IT security than mere trust in the provider (that is, the testability curve should be pushed down as far as possible in this area). In this context, the questions of the applicative embedding of such products and their effects on the measurement or test results are largely unresolved. Whatever is proposed here will have to be measured

against the standard of the greatest possible transparency. In the meantime, cross-industry guidelines such as the MVSP offer welcome transitional solutions. A provider who commits to complying with these requirements is essentially just making a promise that the customer must trust again. The advantage, however, is that in this case it is at least transparent to the customer who and what exactly they are trusting. This is certainly better than blind trust. 🤖

REFERENCES

1. R. Oppliger and A. Grünert, “How to measure cybersecurity and why heuristics matter,” *Computer*, vol. 57, no. 2, pp. 111–115, Feb. 2024, doi: 10.1109/MC.2023.3334054.
2. ISO 7498-2, “Information processing systems - Open systems interconnection - Basic reference model, Part 2: Security architecture,” 1989. <https://www.iso.org/standard/14256.html>
3. G. A. Akerlof, “The market for lemons: Quality uncertainty and the market mechanism,” *Quart. J. Econ.*, vol. 84, no. 3, pp. 488–500, Aug. 1970, doi: 10.2307/1879431.
4. *DoD Trusted Computer System Evaluation Criteria*, Department of Defense (DoD) Standard 5200.28-STD, 1985.
5. *Information Technology Security Evaluation Criteria (ITSEC)*, version 1.2, Jun. 1991. https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Standards-und-Zertifizierung/Zertifizierung-und-Anerkennung/Zertifizierung-von-Produkten/Zertifizierung-nach-CC/IT-Sicherheitskriterien/ITSEC/itsec_node.html

6. ISO/IEC 15408-1, "Information security, cybersecurity and privacy protection - Evaluation criteria for IT security, Part 1: Introduction and general model," 2022. <https://www.iso.org/standard/72891.html>
7. R. W. Shirey, "Internet security glossary, version 2 (RFC 4949)," IETF, Aug. 2007. <https://datatracker.ietf.org/doc/html/rfc4949>
8. "Proposal for a regulation of the European Parliament and the Council on horizontal cybersecurity requirements for products with digital elements and amending Regulation (EU) 2019/1020," European Commission, Brussels, Belgium, Sep. 15, 2022. [Online]. Available <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52022PC0454>
9. ISO/IEC 27001, "Information security, cybersecurity and privacy protection - Information security management systems - Requirements," 2022. <https://www.iso.org/standard/27001>
10. *Payment Card Industry Data Security Standard*, version 4.0.1, PCI Security Standards Council, Wakefield, MA, USA, Jun. 2024. [Online]. Available https://www.pcisecuritystandards.org/document_library/?category=pcidss&document=pci_dss
11. *Secure Storage Units - Requirements, Classification*

and Methods of Test for Resistance to Burglary - Part 1: Safes, ATM Safes, Strongroom Doors and Strongrooms, European Standard EN 1143-1:2019, European Committee for Standardization, Brussels, Belgium, Apr. 2019.

ANDREAS GRÜNERT is an information security officer with the National Cyber Security Center, CH-3003 Bern, Switzerland. Contact him at andreas.gruenert@ncsc.admin.ch.

JAMES BRET MICHAEL is a professor in the Department of Computer Science and Department of Electrical and Computer Engineering, Naval Postgraduate School, Monterey, CA 93943 USA. Contact him at bmichael@nps.edu.

ROLF OPPLIGER is an information security officer with the National Cyber Security Center, CH-3003, Bern, Switzerland, and teaches at the University of Zurich, CH-8050 Zurich, Switzerland. Contact him at rolf.oppliger@ncsc.admin.ch.

RUEDI RYTZ is an information security officer with the National Cyber Security Center, CH-3003 Bern, Switzerland. Contact him at rudolf.rytz@ncsc.admin.ch.



IEEE COMPUTER SOCIETY
Call for Papers

Build your authority in the industry with exposure to a global network of 350K+ computing professionals.

 **GET PUBLISHED**
www.computer.org/cfp


 IEEE COMPUTER SOCIETY 

DEPARTMENT: MEMORY AND STORAGE

Advancing Data Security and Sustainability: Establishing a Circular Economy for Storage

Jonmichael Hands , Chia NetworksTom Coughlin , Coughlin Associates, Inc.

Recent extensions in the IEEE Standard 2883.1-2024 and initiatives by the OCP and the CHIPS Alliance enable a circular economy for digital storage devices. New storage disassembly equipment facilitates recovering useful materials from end-of-life storage devices.

rganizations must balance security imperatives with environmental stewardship in an era defined by explosive data growth and the rise of data-intensive applications, such as artificial intelligence. As we generate and store unprecedented volumes of information, the secure and sustainable handling of storage media has become a top priority. In our previous article,¹ we examined the foundational role of modern media sanitization techniques, particularly IEEE Standard 2883-2022,² in laying the groundwork for a circular economy in storage.

This article delves deeper into the evolving landscape of media sanitization, highlighting key industry trends, emerging technologies, and the implications of the newly published IEEE 2883.1-2024² recommended practice for organizational policies. Together, these developments underscore how a well-defined media sanitization strategy can support both rigorous data protection and responsible resource utilization.

THREE EMERGING TRENDS IN MEDIA SANITIZATION

Several key trends are shaping the future of media sanitization.

Transition from physical destruction to purge sanitization

Historically, organizations defaulted to physical destruction as a fail-safe method of safeguarding sensitive data. While effective from a security standpoint, this approach often undermines sustainability goals by rendering valuable hardware unusable. In contrast, “purge” sanitization methods, such as cryptographic erasure or block-level erase, make data recovery infeasible even under advanced laboratory scrutiny, all while preserving the underlying storage device for reuse and its components for recycling. This shift aligns with circular economy principles, extending device life spans and recapturing value at the end of operational life.

Hardware roots of trust

Infrastructural security technologies like Caliptra, supported by the Open Compute Project (OCP) L.O.C.K. initiative, are redefining standards of hardware integrity. By embedding robust security features directly into storage devices, these solutions help prevent unauthorized access to encryption keys and enable secure erasure at the hardware level. This hardware-anchored trust layer fortifies data protection, streamlines sanitization procedures, and mitigates sophisticated attacks against storage subsystems.

Advanced material recovery techniques

As reliable purge sanitization methods become more prevalent, organizations are looking beyond data protection toward the broader lifecycle implications



TABLE 1. Various sanitization methods defined in IEEE 2883 and evaluation of an adversary’s capability to recover information under different sanitization scenarios.

| Sanitization Method | Novice | Expert | Virtuoso |
|---------------------|-------------------|-------------------|-------------------|
| None | Almost certain | Almost certain | Almost certain |
| Clear | Unlikely | Likely | Almost certain |
| Purge | Almost impossible | Almost impossible | Unlikely |
| Destruct | Almost impossible | Almost impossible | Almost impossible |

of their storage assets. New, automated disassembly tools can reclaim high-value components, such as rare-earth magnets and other precious materials, thereby reducing environmental impacts and enhancing return on investment. By aligning security practices with end-to-end resource recovery, enterprises can contribute to a genuinely sustainable technology ecosystem.

IEEE 2883.1-2024: A NEW RECOMMENDED PRACTICE FOR DATA SANITIZATION

Building on the guidance established in IEEE 2883-2022, the newly released IEEE 2883.1-2024 recommended practice provides a more comprehensive framework for organizations implementing data sanitization policies. This guidance addresses risk assessment, technology selection, verification protocols, and the integration of sustainability considerations into sanitization strategies. Table 1 compares various sanitization methods and the relative difficulty adversaries face in recovering sanitized data.

KEY TAKEAWAYS FROM IEEE 2883.1-2024

Risk-based decision making

Organizations are advised to thoroughly assess breach implications and the likelihood of data recovery under various threat models, tailoring their sanitization measures accordingly.

Use of standardized methods

IEEE 2883 defines clear, purge, and destruct methods, each offering successively stronger assurances against data recovery. IEEE 2883.1 explains which sanitization method to use in the storage lifecycle and the policy and requirements for each.

Verification and assurance

The recommended practice emphasizes the necessity of verifying sanitization effectiveness. Demonstrable evidence builds stakeholder confidence and ensures that sensitive data are indeed irrecoverable.

Sustainability integration

By prioritizing purge sanitization over physical destruction, organizations retain the option to reuse storage media or salvage valuable materials. This supports both compliance and environmental objectives.

By adopting the recommendations of IEEE 2883.1-2024, organizations can create a reasonable policy around storage security, ensure regulatory compliance, streamline sanitization processes, and contribute to environmental sustainability.

CALIPTRA: ESTABLISHING A SILICON-BASED HARDWARE ROOT OF TRUST

Caliptra is an open source silicon-based root of trust (RoT) framework designed for modern data center systems-on-chip.³ Backed by hyperscalers such as Microsoft and Google and managed by OCP and the

CHIPS Alliance, Caliptra enforces firmware integrity, authenticates device identity, and ensures security at the silicon layer.

You can think of Caliptra as a foundational security “building block” baked directly into the silicon chip at the heart of a server or cloud computer. Every modern data center device, from CPUs to specialized accelerators, needs a trustworthy way to know it is running genuine, safe software. Without a secure starting point in hardware, it is tougher to guarantee the overall system is safe.

CORE CAPABILITIES

Hardware-embedded trust

Caliptra provides a secure launch environment, verifying code integrity before system boot and thwarting any tampering attempts at the foundational hardware level.

THIS HARDWARE-ANCHORED TRUST LAYER FORTIFIES DATA PROTECTION, STREAMLINES SANITIZATION PROCEDURES, AND MITIGATES SOPHISTICATED ATTACKS AGAINST STORAGE SUBSYSTEMS.

Robust measurement and attestation

By cryptographically hashing and signing boot code, Caliptra’s attestations enable third parties to confirm that a device is running legitimate firmware without having been stealthily replaced or compromised.

Unique device identity

Each chip can be cryptographically identified, streamlining supply chain integrity and lifecycle management.

Open standards and transparency

Caliptra aligns with widely recognized industry guidelines and is open source, inviting community scrutiny and continual refinement.

Extensibility for future needs

As technology evolves, Caliptra’s modular design simplifies integration and enhancement, allowing

the ecosystem to rapidly adapt to new security requirements.

OCP L.O.C.K.: STRENGTHENING CRYPTOGRAPHIC KEY MANAGEMENT

Building upon Caliptra’s foundation, OCP L.O.C.K.⁴ introduces a layered, open source cryptographic key management solution that prevents media encryption key leakage and enables secure cryptographic erasure through fuse-based purges.

KEY ADVANTAGES

Eliminating the need for physical destruction

OCP L.O.C.K. enhances cryptographic erase by providing audited, hardened key management that enables purge-level sanitization while maintaining the device for reuse.

Improved security posture

Tightly binding encryption keys to secure, externally supplied access keys ensures that even if the device firmware is compromised, internal secrets remain protected.

Comprehensive threat mitigation

Designed with advanced threat actors in mind, OCP L.O.C.K. counters physical and logical attacks, ensuring keys cannot be extracted from disassembled hardware or leaked through compromised firmware.

OCP L.O.C.K. has been designed explicitly for HDD and SSD controllers. A draft of the spec has already been released into OCP, and it is expected that vendors will start enabling drives that support it in the next few years. Enhancements in NVMe and TCG Opal will fully utilize the capabilities, finally getting cryptographic erase qualified in some of the most stringent environments like hyperscalers.

CLOSING THE LOOP: ROBOTIC HDD DISASSEMBLY AND RESOURCE RECOVERY

The circular economy for storage devices is further bolstered by new tools that disassemble HDDs to recover valuable subcomponents. These tools enable efficient

recovery of rare-earth metals and other valuable materials, reducing environmental impacts and maximizing lifecycle values.

Automation and robotics

According to the UN Global E-waste Monitor, only about 22.3% of e-waste was recycled in 2022, leaving significant value unrecovered.⁵ Robotics and automation offer a scalable, efficient solution by autonomously disassembling complex electronics to recover reusable components and materials. Companies like Molg⁶ use robotic microfactories to disassemble laptops, servers, and soon HDDs, preserving high-value parts for reuse or recycling.

Event Horizon microfactory system

Molg is developing the Event Horizon system, a high-precision, nondestructive HDD disassembly solution (Figure 1). By separating valuable rare-earth-containing subcomponents like magnet assemblies from the data-storing disk platters, a company can destroy only the disks and recycle and reuse the rest of the components. This reduces the environmental impacts of incineration and shredding while companies have time to change their policies to adopt media sanitization methods like purge and entire circular economy value chains.

Case in point: Microsoft's HDD disassembly robot

Microsoft's in-house robotic disassembly initiative⁷ aims to securely dismantle hard drives, preserve their most valuable components, and responsibly recycle obsolete parts. By diverting these drives from shredders and into targeted recovery workflows, Microsoft seeks to achieve a 90% reuse rate for servers and components, including HDDs, by 2025. This case illustrates how leading organizations can align data sanitization with broader sustainability goals.

Today's storage security landscape must align the right combination of policies, technologies, and material recovery strategies to meet evolving



FIGURE 1. The Molg HDD disassembly solution. (Source: Used with permission by Molg.)

demands. On one level, we see the introduction of recommended practices like IEEE 2883.1-2024, which equip organizations with clear guidelines to shift from default physical destruction toward more nuanced sanitization approaches, particularly purge methods that can effectively erase data without sacrificing the underlying storage medium. This policy layer sets a strategic framework for balancing security requirements with sustainability goals.

At the next level, hardware-based security solutions, embodied by Caliptra and OCP L.O.C.K., embed a foundational RoT deep within the data center silicon. These initiatives ensure that cryptographic keys remain protected, data are securely purged, and devices can be reliably verified as authentic. In contrast to the policy standards, these technologies operate in the technical trenches, securing data at their source and eliminating the guesswork and vulnerabilities associated with traditional key management methods.

Finally, the physical side of the equation comes into focus through emerging robotic and automated disassembly tools for hard drives. While standards guide practices and hardware RoTs safeguard encryption keys, these mechanical and material innovations reclaim valuable subcomponents and rare-earth metals from sanitized drives. By moving beyond shredding and toward selective disassembly, organizations ensure that the drive's embodied resources contribute to a circular economy, recycling and reusing materials that would otherwise become e-waste. 🌍

REFERENCES

1. J. Hands and T. Coughlin, "New IEEE media sanitization specification enables circular economy for storage," *Computer*, vol. 56, no. 1, pp. 111–116, Jan. 2023, doi: 10.1109/MC.2022.3218364.
2. *IEEE Standard for Sanitizing Storage*, IEEE 2883-2022 & IEEE 2883.1-2024, IEEE Standards Association, Aug. 17, 2022. [Online]. Available: <https://standards.ieee.org/ieee/2883/10277/>
3. "Caliptra." GitHub. Accessed: Dec. 31, 2024. [Online]. Available: <https://github.com/chipsalliance/Caliptra/tree/main>
4. Open Compute Project, "OCP layered open-source cryptographic key-management (L.O.C.K.): NVMe™ Key management specification," Oct. 2024. [Online]. Available: <https://www.opencompute.org/documents/ocp-nvme-key-management-specification-v05-docx-1-pdf>
5. "Global E-waste monitor 2024: Electronic waste rising five times faster than documented E-waste recycling," 2024. [Online]. Available: <https://unitar.org/about/news-stories/press/global-e-waste-monitor-2024-electronic-waste-rising-five-times-faster-documented-e-waste-recycling#:~:text=Meanwhile%2C%20less%20than%20one%20quarter,pollution%20risks%20to%20communities%20worldwide>
6. "Enabling circular manufacturing." MOLG. Accessed: Dec. 31, 2024. [Online]. Available: <https://www.molg.ai/>
7. C. Mellor, "Microsoft says it's opting for robots to recycle disk drives," *Blocks and Files*, Sep. 6, 2024. [Online]. Available: <https://blocksandfiles.com/2024/09/06/microsofts-robotic-disk-driven-disassembly-is-a-step-on-road-to-ending-its-e-waste/>

JONMICHAEL HANDS is with the Chia Network, San Francisco, CA 94104 USA. Contact them at jm@chia.net.

TOM COUGHLIN is the president of Coughlin Associates, Inc., San Jose, CA 95124 USA. Contact him at tom@tomcoughlin.com.

ADVERTISER INFORMATION

Advertising Coordinator

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

Advertising Sales Contacts

Mid-Atlantic US, Northeast, Europe, the Middle East and Africa:
Dawn Scoda
Email: dscoda@computer.org
Phone: +1 732-772-0160
Cell: +1 732-685-6068 | Fax: +1 732-772-0164

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

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

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

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

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

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

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

PURPOSE: Engaging professionals from all areas of computing, the IEEE Computer Society sets the standard for education and engagement that fuels global technological advancement. Through conferences, publications, and programs, IEEE CS empowers, guides, and shapes the future of its members, and the greater industry, enabling new opportunities to better serve our world.

OMBUDSMAN: Contact ombudsman@computer.org.

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

PUBLICATIONS AND ACTIVITIES

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

Periodicals: The IEEE CS publishes 12 magazines, 18 journals

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

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

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

Conferences/Education: The IEEE CS holds more than 215 conferences each year and sponsors many educational activities, including computing science accreditation.

Certifications: The IEEE CS offers three software developer credentials.

AVAILABLE INFORMATION

To check membership status, report an address change, or obtain information, contact help@computer.org.

IEEE 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

IEEE CS EXECUTIVE STAFF

Executive Director: Melissa Russell

Director, Governance & Associate Executive Director:
Anne Marie Kelly

Director, Conference Operations: Silvia Ceballos

Director, Information Technology & Services: Sumit Kacker

Director, Marketing & Sales: Michelle Tubb

Director, Membership Development: Eric Berkowitz

Director, Periodicals & Special Projects: Robin Baldwin

IEEE CS EXECUTIVE COMMITTEE

President: Hironori Washizaki

President-Elect: Grace A. Lewis

Past President: Jyotika Athavale

Vice President: Nils Aschenbruck

Secretary: Yoshiko Yasuda

Treasurer: Darren Galpin

VP, Member & Geographic Activities: Andrew Seely

VP, Professional & Educational Activities: Cyril Onwubiko

VP, Publications: Charles (Chuck) Hansen

VP, Standards Activities: Edward Au

VP, Technical & Conference Activities: Terry Benzel

2025–2026 IEEE Division VIII Director: Cecilia Metra

2024–2025 IEEE Division V Director: Christina M. Schober

2025 IEEE Division V Director-Elect: Leila De Floriani

IEEE CS BOARD OF GOVERNORS

Term Expiring 2025:

İlkay Altıntaş, Joaquim Jorge, Rick Kazman, Carolyn McGregor,
Andrew Seely

Term Expiring 2026:

Megha Ben, Terry Benzel, Mrinal Karvir, Andreas Reinhardt,
Deborah Silver, Yoshiko Yasuda

Term Expiring 2027:

Sven Dickinson, Alfredo Goldman, Daniel S. Katz, Yuhong Liu,
Ladan Tahvildari, Damla Turgut

IEEE EXECUTIVE STAFF

Executive Director and COO: Sophia Muirhead

General Counsel and Chief Compliance Officer:
Ahsaki Benion

Chief Human Resources Officer: Cheri N. Collins Wideman

Managing Director, IEEE-USA: Russell Harrison

Chief Marketing Officer: Jayne O'Brien

Chief Publication Officer and Managing Director:
Steven Heffner

Staff Executive, Corporate Activities: Donna Hourican

Managing Director, Member and Geographic Activities:
Cecelia Jankowski

Chief of Staff to the Executive Director: TBA

Managing Director, Educational Activities: Jamie Moesch

IEEE Standards Association Managing Director: Alpesh Shah

Chief Financial Officer: Kelly Armstrong

Chief Information Digital Officer: Jeff Strohschein

Managing Director, Conferences, Events, and Experiences:
Marie Hunter

Managing Director, Technical Activities: Mojdeh Bahar

IEEE OFFICERS

President & CEO: Kathleen A. Kramer

President-Elect: Mary Ellen Randall

Past President: Thomas M. Coughlin

Director & Secretary: Forrest D. Wright

Director & Treasurer: Gerardo Barbosa

Director & VP, Publication Services & Products: W. Clem Karl

Director & VP, Educational Activities: Timothy P. Kurzweg

Director & VP, Membership and Geographic Activities:
Antonio Luque



Director & President, Standards Association:
Gary R. Hoffman

Director & VP, Technical Activities: Dalma Novak

Director & President, IEEE-USA: Timothy T. Lee

DEPARTMENT: INTERNET OF THINGS, PEOPLE,
AND PROCESSESThis article originally
appeared in
IEEE Internet Computing
vol. 28, no. 2, 2024

On Causality in Distributed Continuum Systems

Victor Casamayor Pujol , Boris Sedlak , Praveen Kumar Donta , and Schahram Dustdar ,
Vienna University of Technology, Vienna, 1040, Austria

As distributed continuum systems (DCSs) are envisioned, they will have a massive impact on our future society. Hence, it is of utmost importance to ensure that their impact is socially responsible. Equipping these systems with causal models brings features such as explainability, accountability, and auditability, which are needed to provide the right level of trust. Furthermore, by combining causality with graph-based service-level objectives, we can cope with dynamic and complex system requirements while achieving sustainable development of DCSs' capacities and applications.

DCSs: NOVEL CHALLENGES

Distributed systems, specifically cloud computing, support many of our daily activities, spanning from the home to the workspace and including many facets of our social life and entertainment activities. These services have become ubiquitous due to their business model, performance, and security guarantees. Computing infrastructure centralization, i.e., the data center, has been a critical enabler for this technology, providing an efficient and scalable solution to manage worldwide-scale services.

Interestingly, the upcoming generations of Internet-distributed systems are moving away from these large and isolated data centers to be closer to the users at the *edge* of the network. This trend is justified by several reasons, such as reducing network congestion, reducing service latency, or improving privacy guarantees. The list is long, and many scientific publications emphasize this change, e.g., Shi et al.¹ and Satyanarayanan.² Furthermore, moving computational resources closer to the user brings new applications for our society, such as autonomous cars, individualized health tracking and assistance, optimized resource management, holographic communications, and many other opportunities.

The potential business opportunities are enormous. Hence, it is inevitable that a new generation of Internet-distributed systems, also known as *distributed continuum systems (DCSs)*,³ will eventually be a part of our

societal and technological ecosystems. Figure 1 shows a high-level view of this type of system. At the top, there are the resources from cloud computing; farther down we find fog and edge computing resources, which are more constrained, and, finally, at the bottom, we can find Internet of Things (IoT) devices and specific applications. This figure shows devices and their interconnections, but it is fundamental to understanding that there is no centralized logic behind it. The cloud does not govern all the devices in the figure, but each subset of devices is responsible for its tasks, and they are related in a loosely coupled manner. Similarly, as services would do in a service-oriented architecture, that is the reason behind using bidirectional arrows in Figure 1.

The technological challenges for these systems are manifold. DCSs are composed of a large diversity of devices and interconnections, which belong to different providers and are mutual to several tenants. Given the open and shared environment these systems inhabit, they need to adapt dynamically to changes while keeping strict constraints fulfilled in terms of performance, quality of service (QoS), or cost. There exist remarkable similarities between DCSs and natural ecosystems when their behaviors need to be explained, and both are self-adaptive systems built from self-adaptive components. We are studying the technological challenges presented by these new systems.⁴ We are confident that, sooner rather than later, all the challenges identified will be surmounted and these systems will become a reality.

DCS applications will have many interactions with people and a significant impact on society. Hence, it is

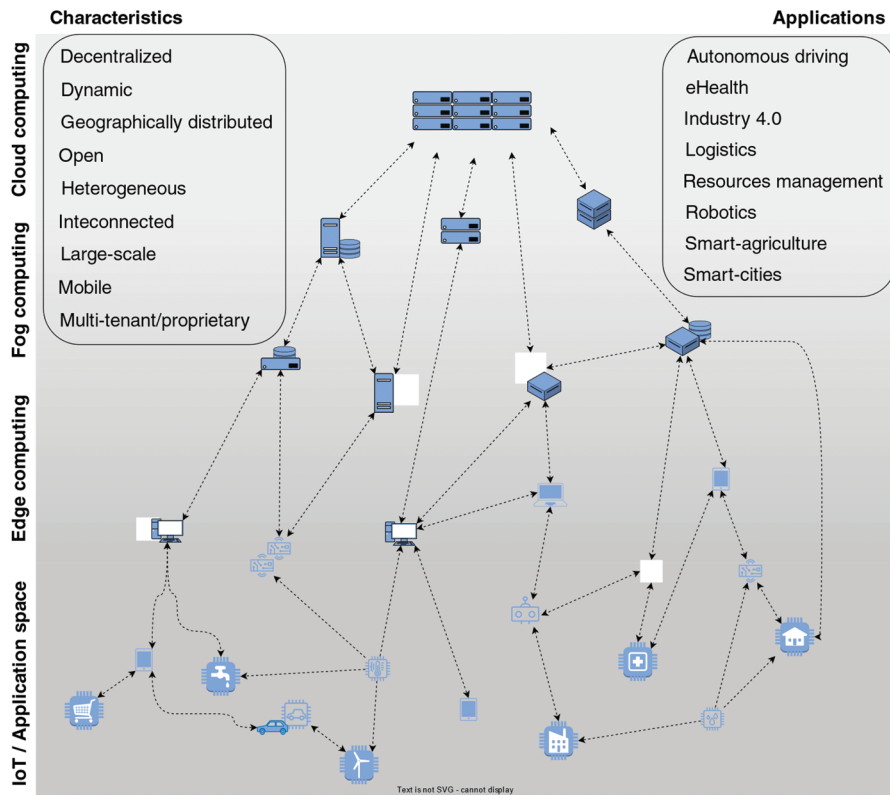


FIGURE 1. General overview of a distributed computing continuum system. IoT: the Internet of Things.

required that their complex behaviors be explainable, fair, and auditable. Causality is a technique that is able to offer this by composing models that explain each of the system's components. In this article, we start motivating causality from a bird's-eye view and take a closer look at how we envision its applicability in DCSs.

Top-Level View on Causality

At the highest conceptual level, causality aims at answering the why question. This question is rooted in human nature: our inherent way of comprehending how everything works. From this view, including causal models in new developments enables the capacity to understand the underlying reasons for their behaviors. This may be seen as an overhead, i.e., one just needs the system to do its job. However, this opens the door to conscious, open, and fair development of new systems. We claim that causality has to be the driving force for sustainable development of DCSs.

This is in contrast to having business growth as the main driving force. We have witnessed how cloud computing and big data analytics have allowed companies to provide new and exciting free services to users: e-mail, storage, social networks, and so on. However, these

services only looked free because their users had become valuable products for these companies.⁵ Similarly, machine learning (ML) and artificial intelligence (AI) have experienced exponential growth for the last one or two decades. Although all new capacities are discovered and enhanced through extensive training, many inconvenient aspects of this technology, such as a lack of trustworthiness⁶ or being resource devouring,⁷ have only emerged since ML/AI have become essential in our society. Fortunately, we are now aware of these perils, and legislation is slowly but steadily trying to solve these deficiencies. Our lesson learned is that business growth cannot be the main driving force for future DCSs.

Causality, as an overarching technique for DCS, can bring the necessary mechanisms to grasp their behavior and footprint fully. Sustainable development is not only about minimizing energy consumption, it is also about developing a holistic view of how these new systems interact with our society, and their impact. Causality is a cornerstone of DCSs' sustainable development as it accompanies concepts such as fairness, trustworthiness, resource efficiency, environmental responsibility, and so on.

CAUSALITY

In general, causality studies the cause and effect relationship between events and variables. A key difference with previous statistical analyses is that causality disconnects from spurious correlations. Hence, it looks for meaningful relationships between events and variables and provides methods for modeling them.

Causal models can be leveraged at three levels, or rungs, as explained in Pearl and Mackenzie.⁸ The first rung is observational, allowing inference on a system that cannot be interacted with, and the only available data are from observations. The second rung is interventional; in this case, the system can be interacted with, and the data obtained reflect this possible interaction. This offers the possibility to make inferences on the system's behavior after performing changes on it. Finally, the third rung is counterfactual, which aims to build models that can infer possibilities that cannot happen, such as, *What would have happened to the system if I had done action x instead of y?* Indeed, obtaining data and models for this type of query is challenging, but it is valuable reasoning to understand how a system can or cannot work. Further, it is a type of reasoning very familiar to us (people) as we commonly imagine situations and their possible outcomes (POs) before or after they have happened.

Embedding causal mechanisms within DCSs ensures human-interpretable backdoors to achieve explainability, accountability, and auditability, among the other benefits that this technology, such as parameter optimization, can bring to DCSs. As DCSs must be self-adaptive, they require knowledge about themselves, i.e., system knowledge, and the environment to which they need to adapt, i.e., context awareness. Crucially, by using causal-based techniques, it is possible to bring a third leg to this autonomy that considers its relationship with society, bringing the opportunity to become socially responsible and sustainable, i.e., societal responsibility (see Figure 2).

It is important to emphasize that causality is not merely a buzzword or trendy concept that we are introducing into distributed systems. It is a fundamental principle that is essential to understand to ensure the development of these new systems. In Figure 2, we can see all the features causality can provide to DCSs. Furthermore, it has been used in computer science for some time now. However, in most cases, it is used to solve specific problems for a specific domain and has never been used as the driving force for a new development.

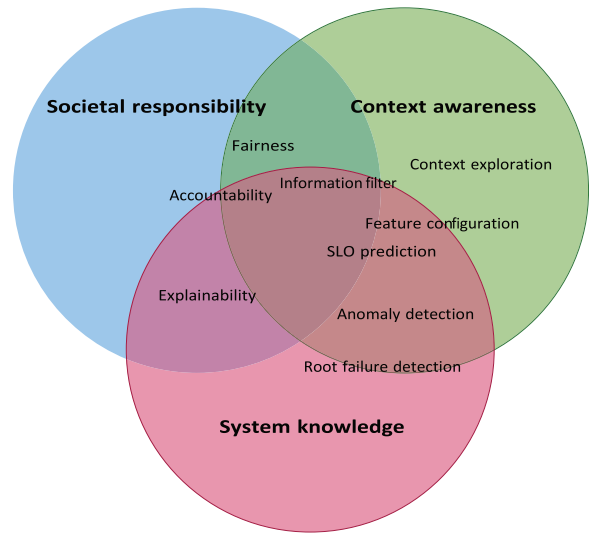


FIGURE 2. Causality features framed over the three main pillars to build DCSs. SLO: service-level objective.

Causality in Computer Science

In this section, we provide an overview of the use of causality in computer science to shed light on the possibilities of this technology for DCSs. Causality has been used in computer science for anomaly detection in distributed cloud-based systems. The complex relationships among services have required the development of fine-grained causal inference techniques to detect the root cause of failures or service-level objective (SLO) violations.⁹ Indeed, this complexity is increased in DCSs. However, the need to detect root-cause failures and SLO violations remains the same for these systems as an expansion of the computing units and services toward the network's edge.

Configurable software systems have also benefited from causality to better understand how system features relate among them and how the system performance is affected by proper configuration. It has been shown that causal relations help identify responsibilities and reasons for features and feature combinations, paving the way for future tailored optimized configurations. Moreover, this also offers the possibility to perform analysis on counter-factual configurations, enabling reasoning over different contexts, an utmost feature for self-adaptive systems.¹⁰ DCS are composed of a diversity of devices and connections. Hence, beyond horizontal or vertical scaling, properly configuring their usage is crucial to offering the expected QoSs.

Recommendation systems are another field where causality is starting to play a major role. These systems

can be understood as the primary information filters on the Internet, as they aim at providing only relevant information for its users. Hence, it is crucial for these systems to differentiate spurious correlations from real causal relations. In this regard, causality is essential to their inference models.¹¹ The volume of data that DCSs generate is huge; beyond the data related to the specific application, there is also a vast amount of data internal to the system, known as *big data*.¹² Hence, filtering these data for system management requires timely and effective processing and understanding.

The ubiquitousness of AI- and ML-inferring systems raises the need to find methods to explain these systems' outcomes; otherwise, it is unknown how much we can trust them. In that regard, causality is being explored to provide those systems with features such as interpretability, explainability, accountability, and fairness. In general, the main goal is to make these systems trustworthy and fair.¹³ Recently, several works have started to study the capacity of causal reasoning for large language models (LLMs).¹⁴ In brief, some causal queries are properly addressed, however, future research aims at incorporating causal models to provide this capacity to LLMs. This shows the relevance of causal reasoning in how we interact with anything. DCSs host sensitive applications, e.g., autonomous driving or e-health. Hence, explainability, fairness, and accountability are mandatory characteristics; otherwise, these applications will need more trustworthiness to be acceptable.

Many computer science fields use causality to boost or provide these missing functionalities to their state-of-the-art techniques. We opt for making causality a fundamental pillar for DCSs as it has the potential to revolutionize them by bringing capacities such as root failure and anomaly detection, SLO violation prediction, feature analysis and configuration, context exploration, information filtering, and decision explainability, while at the same time ensuring fairness and accountability.

Causally Enabled DCSs

Now our main focus is to identify the decisive elements from DCSs that are able to leverage causal models for seamless integration of all the benefits. In general, DCSs comprise cloud infrastructure, several fog nodes, more edge nodes, and many IoT devices, as shown in Figure 1. Their necessary functionalities are spread along the continuum. However, their requirements tend to be more specific to the computing tier. Hence, we need a mechanism that links functionalities and requirements, one that is aware of the specific

idiosyncrasies of the component and can embed the causal models. By the end of the "Control and Management Subsystems" section, we present the artifact able to do this.

Subsystem-Based View of DCSs

Providing a modular view of the system's necessary functionalities eases the understanding of where causal features are required. Further, we comment on how requirements affect functionalities according to the computer tier in which the functionality sits. Hence, the following is organized into subsystems,^a given that they represent system-specific functionalities. In short, any DCSs have the following subsystems: hardware, data, analytics, control, management, and network. Interestingly, in a DCS, there can be hardware components that are also a part of the payload, however, this is outside the scope of this article. Further, there can be other cross-cutting concerns, such as security, transversal to all others, which we understand as if each subsystem requires a *module* on security.

The hardware subsystem aims to manage the hardware components of the system. In that regard, it becomes more relevant in DCSs for two reasons. On the one hand, edge or IoT hardware has power constraints and mobility capacities and is geographically distributed. These aspects challenge the hardware requirements of any previous cloud-based system. On the other hand, future computing systems must be sustainable. Hence, most of the required considerations must be applied in this subsystem.

Hardware subsystem functionalities span from monitoring and tracking the health and performance of the hardware components to ensuring their desired behavior by performing adequate maintenance, including the required firmware updates and component substitutions. As a result, for this subsystem, it is crucial to incorporate anomaly detection and root-cause failure identification from causal methods. This would therefore enable fast and precise action at the hardware level with a high degree of accountability for the actions taken.

The data subsystem is in charge of handling data through its lifecycle, i.e., generation, (pre-)processing, storage, distribution, consumption, and deletion. Indeed, the specific needs of the application, combined with its location along the continuum, affects the requirements for each phase. Further, aspects such as data gravity and friction require special care in DCSs,¹⁵ where

^aThey are also known as planes or layers, but within a system's context, we found subsystems a more appropriate wording.

data may have to be moved through jurisdictional boundaries.

The two main causal features required for this subsystem are 1) context awareness, for defining the specific policies that have to affect the data, and related to that, 2) system accountability for the decisions made, given that they are sensitive.

The analytics subsystem collects system metrics and performs analytics to support the other subsystems. Indeed, the requirements for this subsystem are very different at the edge than at the cloud; however, the expected functionalities are very similar. Simply put, the edge requires lightweight and fast analytics, while cloud requirements can lead to cost-effective requirements. Indeed, performing analytics includes forecasting metric trends and SLO violations and providing optimization possibilities for system configurations.

In that regard, this is a crucial subsystem for applying causal methods. Causality can provide information filters to process only relevant features, i.e., causally dependent. Further, system configuration capabilities can be explored considering the context-awareness capability of causality, and moreover, it can use interventional and counter-factual reasoning to optimize the configuration options of the system. This subsystem must also have explainability, fairness, and accountability features derived from the use of causal models.

Control and management subsystems need conceptually similar mechanisms related to the causal features that can be used. However, in terms of functionalities, they have different perspectives on the system. The control subsystem is responsible for the lower-level needs of the system, typically in short timescales. Hence, it takes care of local tasks that require a fast response. As an example, the access control of resources can be managed by the control subsystem. On the contrary, the management subsystem takes a higher-level perspective on the system: it has larger timescales and works toward global tasks that have slower paces. Hence, provisioning or updating a computing edge cluster can be a part of the management subsystem's tasks. Similarly, as in previous subsystems, specific requirements for the functionalities are linked to their location in the continuum.

These subsystems require input from the system's analytics to act accordingly. In that regard, it is fundamental that they have causal models able to offer explanations for the decisions made. Further, in a more technical view, these systems can also benefit from feature configuration and counter-factual analysis to align the information obtained from the analytics to its specific domain of action. They therefore need to be

able to track previous decisions using the input from the analytics subsystem.

The network subsystem is in charge of managing the communication layer of DCSs. It takes control over specific functions of the communication network. For instance, it must ensure a certain level of throughput while having a dynamic number of requests. However, the techniques and resources vary depending on the computational tier. Generally, in cloud tiers, there are fast and reliable wired connections, while toward the edge, wireless and low-power communication is more typical.

The network subsystem requires to causally model its connectivity efficiency to account for the overall system performance. Furthermore, similarly to the previous subsystems, it also needs the feature configuration and counter-factual analysis features brought by causality.

CAUSALITY INTEGRATED IN DCS

A holistic integration of causality in DCSs is needed to harness their benefits; many of their features are required in all subsystems as key enablers for DCSs. From our experience in distributed systems, an enhanced version of SLOs can be the missing artifact to link causality with distributed systems. Functionalities need to be steered by requirements to be properly adapted to the computing continuum (CC) environment (see Figure 3). Further, DCSs are service oriented; therefore, SLOs can describe these functionalities while determining a specific requirement for them. It is worth noting that SLOs are commonly used in cloud computing systems, however, our enhanced version would reflect all the needs that services throughout the CC may have. This implies that SLOs reflect lower-level system needs, e.g., the maximum CPU usage of a

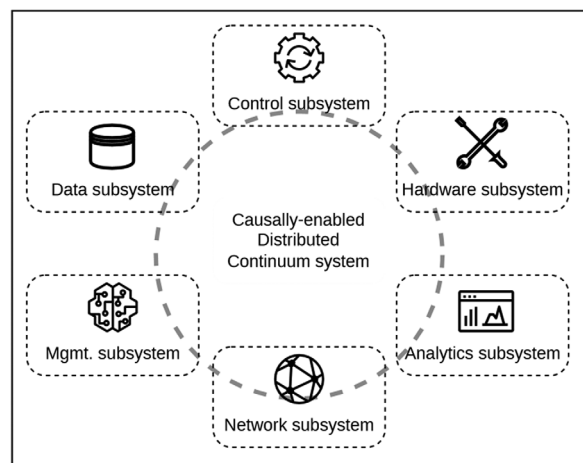


FIGURE 3. Subsystem view for distributed continuum systems.

device as well as higher-level application needs, e.g., the minimal accuracy for a medical-related inference, given that both relate to service requirements. Interestingly, DCSs are highly dependent on their underlying infrastructure; hence, accuracy may be as a result of the ML model as well as the sensor data's granularity.

Enhanced SLOs are built as causal graphs, where the leaf node is the SLO's compliance value, and its parents and grandparents are variables that causally influence the SLO's behavior. Moreover, these variables that influence the SLO's compliance can be parameterized. Looking at the accuracy example, the frames per second of a camera influences the SLO's compliance, but it can be adjusted to different rates. This implies that it is also possible to work at the interventional or counterfactual rungs with causally enhanced SLOs. As shown in Figure 4, an SLO is a causality graph where the values of its variables explain its behavior. These variables can also be deliberately modified, providing a framework for applying do-calculus and counterfactual reasoning in distributed systems. Further, these variables or parameters can relate to different SLOs, providing the capacity to explore the system's behavior beyond a single SLO.

Currently, there are two main frameworks for modeling causality: structural causal models (SCMs),¹⁶ based on causal graphs and structural equations, and the POs framework, which inherits from the development of randomized experiments (see Yao et al.¹⁷ for a

comprehensive survey). From our perspective, SCMs better suit the needs of DCSs, given that the graph view and its equations can easily model DCSs. Causal graphs can be expressed in terms of SCMs, where each parent of the SLO (Pa) together with a set of exogenous (nonobservable) variables (U) define the probability of SLO compliance [$P(\text{SLO})$].

$$P(\text{SLO}) \leftarrow f_i(Pa_i, U_i). \quad (1)$$

Consider that f is a function that relates a parent and the probability distribution of U to the probability distribution of SLO compliance. Also, the subindex on the right side of the equation accounts for different parents and exogenous variables, explaining that each can have a different relationship (i.e., function f) with the SLO. Hence, this formulation also offers the possibility to study the sensitivity of SLO compliance to its variables.¹⁸

Modeling these systems through SLO-based causal graphs brings the following crucial benefits:

- Defining SLOs through causal graphs embeds an information filter for each SLO. Given that every time an SLO is required to be analyzed, e.g., to understand why its compliance value has changed, only the information of the variables in its causal graph has to be assessed, which is the minimal set of information. This also relates to and emphasizes the use of Markov's blanket concept.³

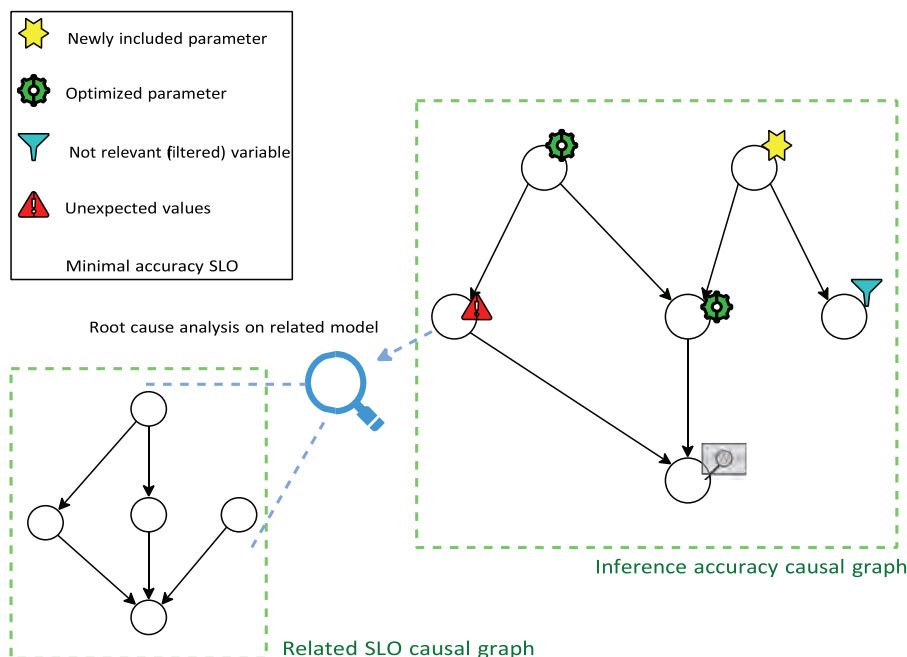


FIGURE 4. Causal graphs for an SLO-based designed system.

- › Causal models provide context awareness, given that they are specific to a service in a determined environment. However, they can also be updated, which means that changes in the service context can be integrated into the causal model, which adapts it to its new reality.
- › Exploring possible configurations of the system and their implications for the SLOs is possible when the configuration variables are included in the SLO-based causal graph. Hence, parameter explainability and optimization go together with the causal-enabled SLO.
- › The granularity of causal models is linked to the described SLO. Further, it is expected to find variables in a causal graph that are also related to other SLOs, which may have a different level of granularity. Hence, by following these relations, root cause and anomalies can be detected regardless of the distance between the observed effect and the root cause.
- › Causal graphical models are interpretable; hence, using them at the service-requirement level integrates the possibility of explaining and therefore accounting for the reasons for a service behavior.

Currently, we are making progress on defining SLOs for DCSs through a Bayesian network. Our next step is ensuring that the Bayesian network behaves as a causal graph so that we unfold all the benefits given by them.

CONCLUSION

The development of DCSs will have an enormous impact on our future society. Hence, they require embedding the capacity for being explainable, fair, accountable, and auditable. This is on top of all the other technical challenges that still need to be solved for these large-scale, heterogeneous, and dynamic systems. In this article, we motivated causality, in the form of causal graphs and SCMs, as the technique to be embraced by these systems to overcome all their technical challenges while also bringing these crucial capacities for being socially responsible. Further, we proposed their integration with SLOs to obtain this holistic framework for developing DCSs. 🌍

REFERENCES

1. W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637–646, Oct. 2016, doi: 10.1109/JIOT.2016.2579198.
2. M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 50, no. 1, pp. 30–39, Jan. 2017, doi: 10.1109/MC.2017.9.
3. S. Dustdar, V. C. Pujol, and P. K. Donta, "On distributed computing continuum systems," *IEEE Trans. Knowl. Data Eng.*, vol. 35, no. 4, pp. 4092–4105, Apr. 2023, doi: 10.1109/TKDE.2022.3142856.
4. V. Casamayor Pujol, P. K. Donta, A. Morichetta, I. Murturi, and S. Dustdar, "Edge intelligence—Research opportunities for distributed computing continuum systems," *IEEE Internet Comput.*, vol. 27, no. 4, pp. 53–74, Jul./Aug. 2023, doi: 10.1109/MIC.2023.3284693.
5. T. Morey, T. Forbath, and A. Schoop, "Customer data: Designing for transparency and trust," *Harvard Business Review*, May 2015. [Online]. Available: <https://hbr.org/2015/05/customer-data-designing-for-transparency-and-trust>
6. D. Kaur, S. Uslu, K. J. Rittichier, and A. Durrezi, "Trustworthy artificial intelligence: A review," *ACM Comput. Surv.*, vol. 55, no. 2, pp. 39:1–39:38, Jan. 2022, doi: 10.1145/3491209.
7. E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the dangers of stochastic parrots: Can language models be too big?" in *Proc. ACM Conf. Fairness, Accountability, Transparency*, 2021, pp. 610–623, doi: 10.1145/3442188.3445922.
8. J. Pearl and D. Mackenzie, *The Book of Why: The New Science of Cause and Effect*. New York, NY, USA: Basic Books, 2018.
9. P. Chen, Y. Qi, and D. Hou, "CausalInfer: Automated end-to-end performance diagnosis with hierarchical causality graph in cloud environment," *IEEE Trans. Services Comput.*, vol. 12, no. 2, pp. 214–230, Mar./Apr. 2019, doi: 10.1109/TSC.2016.2607739.
10. C. Dubslaff, K. Weis, C. Baier, and S. Apel, "Causality in configurable software systems," in *Proc. 44th Int. Conf. Softw. Eng. (ICSE)*, New York, NY, USA: ACM, Jul. 2022, pp. 325–337, doi: 10.1145/3510003.3510200.
11. C. Gao, Y. Zheng, W. Wang, F. Feng, X. He, and Y. Li, "Causal inference in recommender systems: A survey and future directions," Aug. 2022. [Online]. Available: <http://arxiv.org/abs/2208.12397>
12. P. K. Donta, B. Sedlak, V. C. Pujol, and S. Dustdar, "Governance and sustainability of distributed continuum systems: A big data approach," *J. Big Data*, vol. 10, no. 1, Apr. 2023, Art. no. 53, doi: 10.1186/s40537-023-00737-0.
13. N. Ganguly et al., "A review of the role of causality in developing trustworthy AI systems," Feb. 2023. [Online]. Available: <http://arxiv.org/abs/2302.06975>
14. C. Zhang et al., "Understanding causality with large language models: Feasibility and opportunities," Apr. 2023. [Online]. Available: <http://arxiv.org/abs/2304.05524>
15. B. Sedlak, V. C. Pujol, P. K. Donta, and S. Dustdar, "Controlling data gravity and data friction: From metrics to multidimensional elasticity strategies,"

in *Proc. IEEE Int. Conf. Softw. Services Eng. (SSE)*, Jul. 2023, pp. 43–49, doi: 10.1109/SSE60056.2023.00017. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/10234353>

16. J. Pearl, "Causal inference in statistics: An overview," *Statist. Surv.*, vol. 3, pp. 96–146, Jan. 2009, doi: 10.1214/09-SS057. [Online]. Available: <https://projecteuclid.org/journals/statistics-surveys/volume-3/issue-none/Causal-inference-in-statistics-An-overview/10.1214/09-SS057.full>
17. L. Yao, Z. Chu, S. Li, Y. Li, J. Gao, and A. Zhang, "A survey on causal inference," *ACM Trans. Knowl. Discovery Data*, vol. 15, no. 5, pp. 74:1–74:46, May 2021, doi: 10.1145/3444944.
18. C. Cinelli, D. Kumor, B. Chen, J. Pearl, and E. Bareinboim, "Sensitivity analysis of linear structural causal models," in *Proc. 36th Int. Conf. Mach. Learn.*, PMLR, May 2019, pp. 1252–1261. [Online]. Available: <https://proceedings.mlr.press/v97/cinelli19a.html>

VÍCTOR CASAMAYOR PUJOL is a project assistant (post-doctoral researcher) with the Distributed Systems Group, Vienna University of Technology, Vienna, 1040, Austria. Contact him at v.casamayor@dsg.tuwien.ac.at.

BORIS SEDLAK is a Ph.D. candidate with the Distributed Systems Group, Vienna University of Technology, Vienna, 1040, Austria. Contact him at b.sedlak@dsg.tuwien.ac.at.

PRAVEEN KUMAR DONTA is a postdoctoral researcher with the Distributed Systems Group, Vienna University of Technology, Vienna, 1040, Austria. Contact him at p.donta@dsg.tuwien.ac.at.

SCHAHRAM DUSTDAR is a full professor of computer science and heads the Research Division of Distributed Systems, Vienna University of Technology, Vienna, 1040, Austria. Contact him at dustdar@dsg.tuwien.ac.at.

Unleash Your Potential

ATTEND WORLD-CLASS CONFERENCES — Over 195 globally recognized conferences.

EXPLORE THE DIGITAL LIBRARY — Nearly 1 million articles covering world-class peer-reviewed content.

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

LEARN NEW SKILLS — Strengthen your resume with the IEEE Computer Society Course Catalog.

LEVEL UP YOUR CAREER — Search for new positions in the IEEE Computer Society Jobs Board.

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



Explore the benefits of membership today At the IEEE Computer Society

computer.org/membership



IEEE
COMPUTER
SOCIETY



FlyNet: Drones on the Horizon

Alicia Esquivel Morel , Chengyi Qu, and Prasad Calyam, *University of Missouri, Columbia, MO, 65201, USA*

Cong Wang, Komal Thareja, and Anirban Mandal, *Renaissance Computing Institute, University of North Carolina at Chapel Hill, Chapel Hill, NC, 27517, USA*

Eric Lyons and Michael Zink, *University of Massachusetts Amherst, Amherst, MA, 01003, USA*

George Papadimitriou and Ewa Deelman, *University of Southern California, Los Angeles, CA, 90089, USA*

Over the past few years, due to the boom of advances in image processing, edge computing, and wireless networking, unpiloted aerial vehicles, often referred to as drones, have become an important enabler to support a wide variety of scientific applications, ranging from environmental monitoring, disaster response, and wildfire monitoring to the survey of archaeological sites. In this article, we present the FlyNet platform, which extends an existing workflow management system to support and manage scientific workflows. FlyNet enables automated resource allocation, workflow instrumentation, and network service support to support researchers in their goal to analyze data for new scientific discoveries. In addition, FlyNet provides network services management to support quality of service for efficient data transport between edge devices, edge servers, and the cloud.

INTRODUCTION

Drones are literally on the horizon. Unpiloted aerial vehicles (UAVs) (often referred to as drones) are now supporting a wide range of scientific applications, ranging from environmental monitoring, disaster response, and wildfire monitoring to the survey of archeological sites. The success of these applications heavily depends on the ability to efficiently manage and analyze large volumes of data generated by drones. This is where scientific workflow support comes into play, providing researchers with the tools and techniques to better manage and analyze their data. In this context, scientific workflows can be characterized as a series of processes that are executed in a specific order to analyze the data generated by drones. Examples include the processing and analysis of video, imagery, and other sensor data. By using workflow management systems for scientific

UAV applications, researchers can create data management and analysis processes with the goal of efficiently and effectively extracting insights and new knowledge from the collected data.

In parallel, there has been an evolution of the cloud computing paradigm with the advent of edge computing, providing researchers with the opportunity to span their workflows across the edge-to-cloud spectrum based on the resource needs of their scientific applications. To streamline data management based on application requirements, resources across the spectrum need to be appropriately allocated. Unfortunately, selecting the appropriate set of resources for a specific scientific workflow is often a challenge for domain scientists who are not experts in distributed computer systems.

FlyNet introduces a platform to support scientific workflows from the edge to the core for UAV and other edge-to-cloud applications by automating the processes of resource allocation, workflow implementation, and network service support to support researchers in their goal to analyze data for new scientific discoveries.

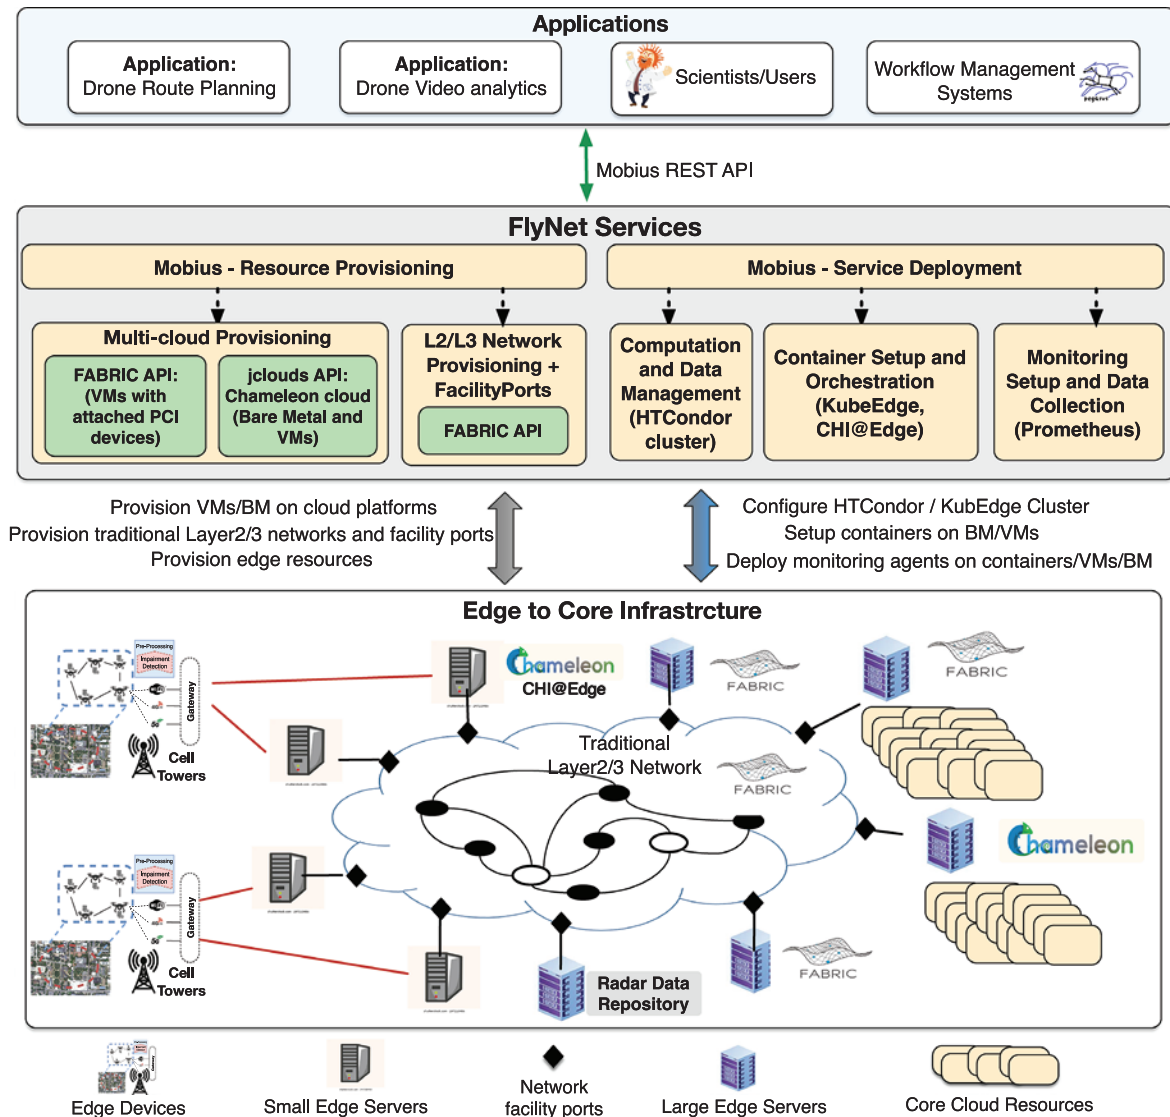


FIGURE 1. FlyNet System architecture showing how applications can leverage edge-to-core infrastructure via FlyNet services.

API: application programming interface; BM: bare metal; L: level; REST: Representational State Transfer; VM: virtual machine.

FLYNET SYSTEM ARCHITECTURE

The FlyNet architecture (shown in Figure 1) supports the composition of end-to-end (edge-to-core) workflows capable of supporting scientific UAV and other edge-to-cloud applications.

Edge-to-Core Infrastructure

The edge-to-core infrastructure depicted at the bottom of Figure 1 covers all points in the spectrum of response latency for application processing—the *latency spectrum*. While some processing needs to be performed on the devices and the network edge to support the

increasing scale of Internet of Things (IoT) applications, some computations need to be performed in network, and some can be offloaded to core computing resources “far” from the edge devices.

There are several categories in this latency spectrum—*edge devices*, *edge servers*, *in network*, and *core computing*. While edge devices provide minimum latency for response times, they have limited computational capabilities and/or power constraints. Thus, onboard resources are often not sufficient to support the UAV application processing needs. Edge servers or nodes that make up an edge computing infrastructure have more computational power and fast turnaround

times, but they support only limited scales of computation (e.g., they might be able to run very lightweight algorithms but not data- and compute-intensive workloads like deep learning models).

As the latency on the spectrum increases, processing packets and turning them around using in-network computing capabilities (either compute resources or specialized programmable hardware deployed in the network core) can be envisioned. This will reduce latency compared to cases where data have to be transmitted all the way to the computing core. For UAV data processing that needs substantially more computational resources (e.g., GPUs for training machine learning models for object detection), data need to travel all the way to core cloud resources. This incurs the maximum latency with the benefit that high processing power can be utilized.

FlyNet Services: Resource Provisioning

To implement this overall architecture, FlyNet uses a network-centric platform called Mobius¹ with support for provisioning programmable cyberinfrastructure comprised of FABRIC² and Chameleon Cloud³ testbeds. Mobius makes it easier for applications to provision and manage the appropriate infrastructure resources for their execution. It supports multiclouds and automated network provisioning to connect the clouds. It leverages the *jclouds* application programming interface (API), which supports OpenStack-based clouds, to provision bare metal (BM) nodes and virtual machines (VMs) from Chameleon. It uses the FABRIC FABlib API⁴ to 1) provision VMs from FABRIC with directly attached PCI devices—GPUs, network cards, non-volatile memory (NVMs) drives, and field-programmable gate arrays—and 2) to provision layer 2/3 networks and facility ports⁵ for connecting different FABRIC core and edge nodes with external infrastructure. Users, applications, and workflow management systems interact with Mobius using a Representational State Transfer (REST) API for provisioning resources and deploying services (see the next section).

FlyNet Services: Service Deployment Container Setup and Orchestration

Since we envision that edge servers will be shared by more than one application, the FlyNet architecture supports a container-based application deployment approach by using *KubeEdge*,⁶ which provides container orchestration at the edge. This containerized approach provides FlyNet with the required flexibility for workflows that support drone-based applications. The use of containers adds the benefit of simplified

deployments of applications on edge nodes and supports the migration of applications between edge nodes. The latter is an important requirement of drone-based applications, where the distance and, thus, the resulting latency between a drone and an edge node might become too large for effective and safe operations. In that case, migrating the application to a different edge node that is closer to the drone is critical. To support FlyNet, we extended Mobius to automatically deploy a container orchestration service using KubeEdge, which automatically instantiates KubeEdge clusters on the provisioned nodes. To support BM container orchestration on the edge resources, as on the Chameleon edge resources—*CHI@Edge*,⁷ Mobius takes advantage of the REST API⁸ to provision the containers.

Computation and Data Management Services

Mobius services also allow applications and workflow systems to deploy HTCondor⁹ clusters—HTCondor Master/scheduler and HTCondor workers—on the provisioned resources selected from (potentially) multiple cloud platforms (FABRIC and Chameleon), such that workflow/application tasks can be readily scheduled and executed. Mobius automates configurations for the networks, Internet Protocol addresses, and setup of the daemons and makes it easier for scientists and applications to use the provisioned infrastructure.

Monitoring Setup and Data Collection: Prometheus

Mobius also automatically deploys Prometheus¹⁰ monitoring agents on the provisioned resources—containers/VMs/BM. These agents monitor different resource metrics, e.g., CPU loads, continuously and stream the measurements to a central Prometheus server. The Prometheus server aggregates all of the monitoring time series data from the agents and exposes an API for applications. The applications can query on the observed performance attributes of the resources and make key decisions for resource management. Such monitoring data are critical for edge resource selection.

EDGE-TO-CLOUD WORKFLOW ORCHESTRATION

Challenges of Edge-to-Cloud Execution

Edge-to-cloud computing environments make it possible for applications and systems to capitalize on the desirable advantages offered by both computing paradigms: faster response times, data locality, cost savings at the edge, scalability, high availability, and reliability provided by the cloud. Effectively utilizing

both computing paradigms within such a complex execution environment for a given application presents a number of challenges. First, available resources and their states need to be visible to make scheduling decisions. Some environments with IoT devices may experience churn due to limited power and network connection. This is especially the case for UAVs that might come in and out of communication range when executing a mission. Second, scheduling decisions must be made. When running in the cloud, both compute and data movement costs may need to be considered. Incorporating the edge may involve taking into consideration energy consumption, limited compute capacity, and storage constraints. In addition to scheduling decisions, there may be resource provisioning decisions that can be made to better accommodate varying levels of expected load. Such provisioning can happen at the edge, for example, in a cloudlet or on idle edge devices. Third, software systems must be in place to execute computations at both ends and automatically handle failures when they occur. Finally, the ability to capture fine-grained performance metrics or provenance data is indispensable to optimizing executions on an edge-to-cloud continuum.

Edge-to-Cloud Workflow System Design

To orchestrate workflows that span edge and cloud resources, FlyNet uses the Pegasus Workflow Management System.¹¹ Pegasus has a number of key features that make it a particularly good candidate to provide the automation needed to span the edge-to-cloud continuum. Most importantly, it has the notion of an abstract workflow. This is a workflow description that is resource independent and captures the workflow at the science level: the codes used for the computations as well as the data needed for and generated by the workflow tasks. Pegasus takes this abstract workflow description and maps it to the available resources, generating the necessary resource-dependent scripts for job submission and adding the necessary data movement between jobs by invoking appropriate data transfer protocols. These resource-specific scripts produced by Pegasus form the executable workflow that is then passed to HTCondor's DAGMan¹² for execution.

Pegasus's architecture and the use of proven and versatile technologies, such as HTCondor, allowed us now to extend the workflows to the edge. HTCondor can run on any edge or cloud resource running Linux, macOS, or Windows, creating a hybrid edge-cloud infrastructure. To match jobs specifically with edge or cloud resources, we added an additional attribute,

which indicates whether or not that resource was an edge or cloud resource. During workflow generation, jobs can be annotated with the type of resources they should be matched with. During execution, HTCondor takes into account this requirement in addition to other job requirements and matches the job with the appropriate resources.

To support data movement operations, workflows are configured to use remote transfer protocols, such as HTTP and SCP, and local file system operations. These are managed by the *pegasus-transfer* utility. Pegasus-transfer is invoked for each job to handle staging in input data and staging out output data. For jobs that are scheduled on locations where input data already reside, *symlinks* are used by *pegasus-transfer* to avoid unnecessary data movements and reduce overall disk usage. One notable advantage of *pegasus-transfer* is that data movement operations are decoupled from the jobs themselves. For example, a change in the locations of initial input files would only require a workflow-specific configuration change with Pegasus.

Workflow Evaluation

For the evaluation, we used a drone application and two other edge-to-cloud workflows. We use these applications to demonstrate the feasibility of our approach and the benefits of using an infrastructure that provides resources across the edge-to-cloud continuum.

Typical UAV Workflow

This workflow¹³ was developed to represent data aggregation and analytics applications that run in edge-to-cloud environments. For such applications, initial input data are derived at the edge from multiple instruments, such as cameras and sensors, mounted on drones. Each input goes through preprocessing steps before being aggregated by a single job that outputs the final result.

Wind Workflow

The Wind workflow^{1,14} is designed to identify areas of maximum observed wind magnitudes from a network of overlapping Doppler weather radars. Single radar files in polarimetric format, from a total of seven radars, are regridded into a common coordinate system. At a centralized location, the workflow periodically takes any available scans collected over a given time interval and creates a new file in a latitude/longitude projection representing the highest winds that have been observed during the time period.

Orcasound Workflow

Orcasound¹⁵ is a community-driven project that leverages hydrophone sensors deployed in three locations

in the state of Washington (San Juan Island, Point Bush, and Port Townsend) to study orca whales in the Pacific Northwest region. The Orcasound Pegasus workflow¹⁶ processes the hydrophone data of one or more sensors in batches for each timestamp and converts them to a WAV format. Using the WAV output, the workflow creates spectrogram images that are stored in the final output location. Furthermore, using a pretrained Orcasound model developed by the community, the workflow scans the WAV files to identify potential sounds produced by the orcas.

Edge-to-Cloud Evaluation

To evaluate our approach, we executed each of the three workflows in edge-only, edge-to-cloud, and cloud-only scenarios. We emulated an edge-to-cloud scenario and provisioned nodes on both Chameleon sites at Texas Advanced Computing Center (TACC) and University of Chicago (UChicago). At TACC, we deployed our cloud site, where we assumed we could get unlimited resources, and, at UChicago, we used Docker to deploy our edge nodes and limit their processing capabilities.¹⁷

In Figure 2, we present the average makespan for 10 runs of each of the three workflows under the different scenarios as a percentage of the edge scenario. As can be seen, the wall clock time (makespan) for each of the three workflows is different for the three execution environments. While the typical UAV workflow performs best in an edge-only environment, the Wind and Orcasound workflows perform best in the cloud-only environment.

Additionally, in Figure 3, we present the average time the workflows spent transferring data over the wide area network (WAN) as a percentage of the edge scenario. This figure provides some insights as to why the cloud-only scenario does not perform the best in




FIGURE 2. Workflow makespans for 10 runs of each of the three workflows under different scenarios. UAV: unpiloted aerial vehicle.

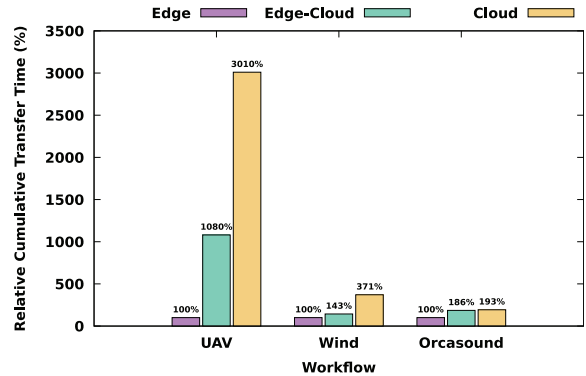


FIGURE 3. Cumulative time spent on transferring data over the wide area network.

all cases. The UAV workflow was designed to favor the edge-only scenario, and, without any computation at the edge, the workflow is forced to spend 30 times more on WAN transfers, negating any increase in compute power the cloud offers. On the other hand, the Wind and Orcasound workflows still have to spend about four times and two times more on WAN transfers, respectively, but the speed-up these workflows are getting from the cloud resources is enough to improve their overall makespans (Figure 2).

Overall, these results show the benefits and flexibility this approach provides. Without any additional development, Pegasus can map the workflows to edge and/or cloud resources, enabling optimizations under constraints utilizing different tradeoffs (e.g., a shorter makespan versus more network utilization).

NETWORK SERVICES FOR EDGE-TO-CLOUD WORKFLOWS

The edge-to-cloud orchestration presented in the “Edge-to-Cloud Workflow Orchestration” section shows the benefits of being able to explore the tradeoff between compute time, data transmission time, and queueing delays for different workflows. In addition to this workflow orchestration, we also investigate how network services that are based on programmable data planes can efficiently manage the transmission of data in the edge-to-cloud continuum. Such network services are an important component in the FlyNet architecture since they support efficient data transport between edge devices, edge servers, and the cloud. Figure 4 shows an example scenario for search-and-rescue operations, which requires the efficient transmission of video footage to adequate compute resources.

The advent of programmable data planes provides in-band telemetry (INT) capabilities that address network

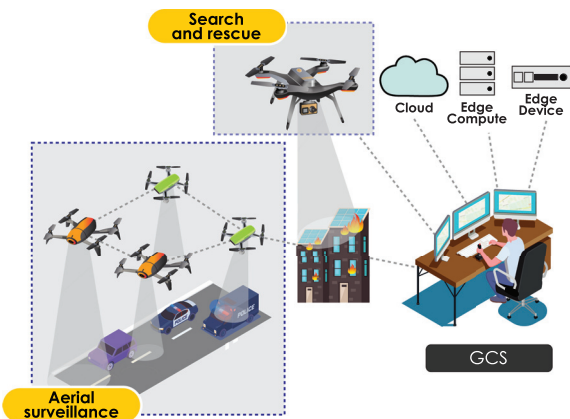


FIGURE 4. UAVs can be utilized for a wide variety of applications, such as, e.g., search and rescue as well as aerial surveillance. Challenges for network services management need to be overcome to guarantee the satisfactory performance of network-edge-based applications, such as video delivery. GCS: ground control station.

resource usage, identify resource contention, and provide detailed visibility into the network infrastructure. Based on these capabilities, INT can be used to enable network quality of service (QoS), ensuring that workflows receive the required network service.

INT

INT-based packet processors [e.g., programming protocol-independent packet processors (P4¹⁸)] enable the

generation of monitoring data. In contrast to existing approaches, INT based on P4 allows for the collection of network metrics (delay, jitter, BW, etc.) on a per-hop basis. Thus, QoS-related issues with a specific link can be pinpointed to a specific segment of the path, allowing network services to address these issues with the goal of maintaining the required QoS.

To further illustrate, Figure 5 depicts an INT implementation. At each of the programmable P4 switches, INT data in the form of the outgoing queue length are collected and added to the packets traversing the link. At the egress point, these metadata are removed from the packet (before it is forwarded to h_2) and analyzed. Queue sizes above a certain threshold might indicate that the required QoS can no longer be supported along this path. In this case, network services can be invoked to actively manage the network (rerouting or limiting of other traffic) to further guarantee the required QoS.

Network Services Control and Workflow Evaluation

As shown in Figure 1, the FlyNet architecture is designed to operate on advanced network infrastructures like FABRIC.² The availability of programmable network elements in FABRIC supports INT scenarios, as shown in Figure 5. The benefits of this approach can be demonstrated by a scenario in which a swarm of drones sends video footage from a search-and-rescue operation. Through the combination of INT and multihop route inspection (MRI), a control system can be created that is aware of the entire network topology between

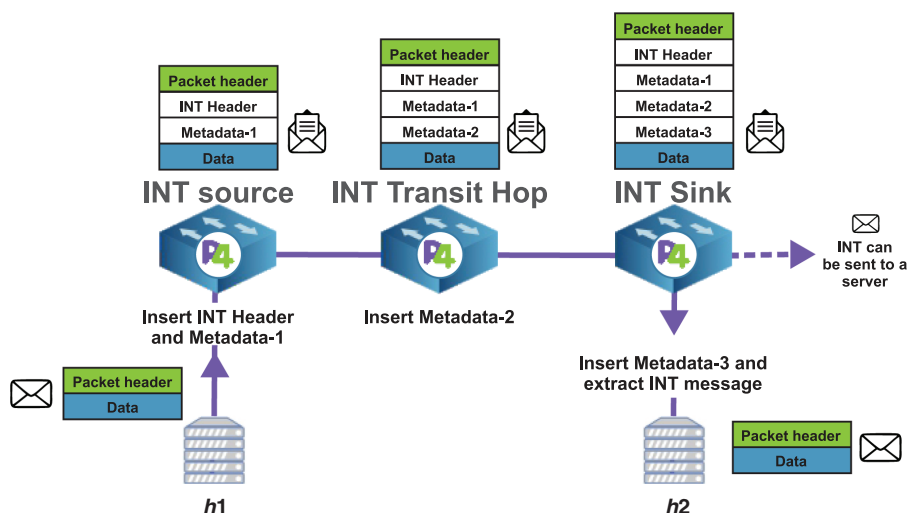


FIGURE 5. Illustration of the application of INT where data are transmitted between hosts h_1 and h_2 using three programmable network switches—the INT source, transit hop, and sink add headers—to report the time spent in the outgoing queues across the network path. INT: in-band telemetry.

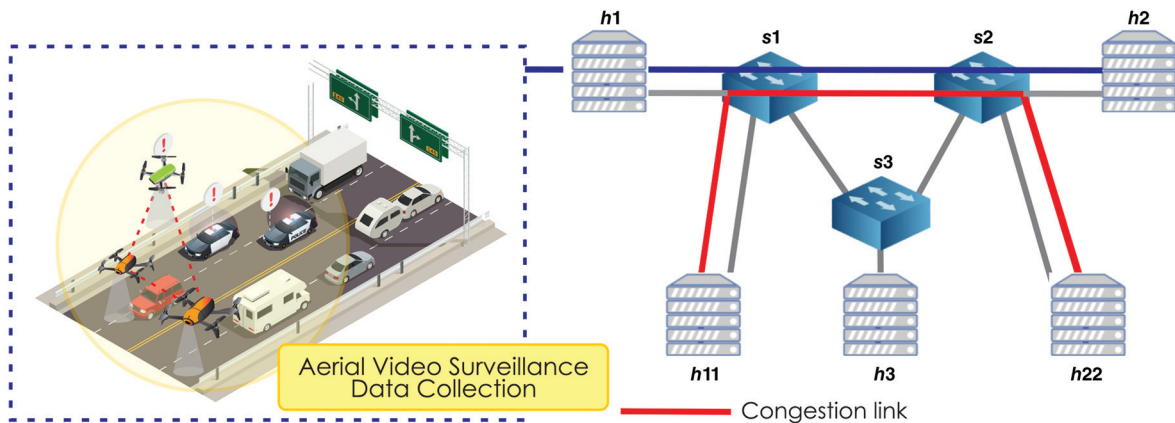


FIGURE 6. Aerial video surveillance data collection use case scenario of experiencing congestion bottlenecks without P4 programmable devices.

IoT devices (a swarm-of-drones scenario), edge servers, and the cloud. It allows the detection of congestion within that topology and can actively intervene to prevent it.¹⁹

Figure 6 depicts a scenario in which aggregated video streams from a swarm of drones are transmitted from edge server $h1$ to cloud server $h2$ via $s1$ and $s2$. Due to competing traffic between $h11$ and $h22$, packet loss and delay can occur for the video stream. With the aid of INT, the link on which this packet loss and delay occur can be identified, and MRI is invoked to reroute the competing traffic (from $h11$ to $h22$) via $s3$, mitigating the congestion on the $s1$ -to- $s2$ link.

As the results in Figure 7 show, this INT-based network service (implemented via P4 in FABRIC) is able to guarantee QoS for the video streams generated by the swarm of drones. While there is significant packet loss when no INT is applied (cases 1 and 2), there is no




FIGURE 7. Packet loss measurements to show the impact of increased congestion on the path between $s1$ and $s2$ with capacity of 200 Mb/s for the following cases: without P4 and congestion of 800 Mb/s (case 1), without P4 and congestion of 400 Mb/s (case 2), and with P4 (case 3).

packet loss when an INT-based network service is used (case 4).

CONCLUSION

UAVs, often referred to as drones, have become an important enabler for a wide variety of scientific and societally impactful applications. FlyNet supports these applications by providing automated resource allocation, workflow instrumentation, and network service management. It leverages the Pegasus workflow management system for supporting and managing scientific workflows spanning from the edge to the core cloud as well as Mobius, a resource-provisioning system that can build a virtual edge-to-cloud platform. In combination with network services that are based on programmable network elements, FlyNet is able to allocate network and compute resources to optimize the execution of these UAV workflows. As a result, researchers can collect and efficiently analyze data, make scientific discoveries, or react to information coming from remote locations.

While we have created a platform that supports drone-based research, there are many research issues that still need to be addressed in the future. For example, the interdependency between data collection and offloading under uncertain network connectivity conditions has not been sufficiently studied. Resource provisioning, task scheduling, and fault recovery that take into account a number of competing criteria, including performance, reliability, and power, are still challenging. We will address such research issues through the exploration of new algorithm design and experimentation with FlyNet on wireless testbeds like AERPAW.²⁰

In the future, we will utilize and extend the FlyNet platform to conduct new drone-based research—supporting new use cases like utilizing a network of drones for emergency management, using a network of edge computing systems to perform drone computations, and executing machine learning algorithms with varying computational requirements across the latency spectrum.

We also plan to harden, test, and expand its capabilities to make them available as part of the overall cyberinfrastructure ecosystem. This will allow scientists, engineers, and emergency managers to leverage FlyNet's capabilities for their work. 🤖

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Award CNS-1950873, Award CNS-1647182, and Award OAC-2018074. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

1. E. Lyons et al., "Toward a dynamic network-centric distributed cloud platform for scientific workflows: A case study for adaptive weather sensing," in *Proc. 15th Int. Conf. eSci.*, 2019, pp. 67–76, doi: 10.1109/eScience.2019.00015.
2. I. Baldin et al., "Fabric: A national-scale programmable experimental network infrastructure," *IEEE Internet Comput.*, vol. 23, no. 6, pp. 38–47, Nov./Dec. 2019, doi: 10.1109/MIC.2019.2958545.
3. K. Keahey et al., "Lessons learned from the chameleon testbed," in *Proc. USENIX Annu. Tech. Conf. (USENIX ATC)*, Berkeley, CA, USA: USENIX Association, Jul. 2020, pp. 219–233.
4. "FABLib API." FABRIC. Accessed: Jan. 12, 2023. [Online]. Available: <https://learn.fabric-testbed.net/knowledge-base/fablib-api/>
5. "Network services in FABRIC." FABRIC. Accessed: Jan. 12, 2023. [Online]. Available: <https://learn.fabric-testbed.net/knowledge-base/network-services-in-fabric/>
6. Y. Xiong, Y. Sun, L. Xing, and Y. Huang, "Extend cloud to edge with kubeedge," in *Proc. IEEE/ACM Symp. Edge Comput. (SEC)*, 2018, pp. 373–377, doi: 10.1109/SEC.2018.00048.
7. "Chameleon edge resources." Chameleon Cloud. Accessed: Jan. 12, 2023. [Online]. Available: <https://www.chameleoncloud.org/experiment/chiedge/>
8. "Containers service API." OpenStack. Accessed: Jan. 12, 2023. [Online]. Available: <https://docs.openstack.org/api-ref/application-container/>
9. D. Thain, T. Tannenbaum, and M. Livny, "Distributed computing in practice: The Condor experience," *Concurrency Comput., Pract. Experience*, vol. 17, nos. 2–4, pp. 323–356, Feb./Apr. 2005, doi: 10.1002/cpe.938.
10. "Overview." Prometheus. Accessed: Jan. 12, 2023. [Online]. Available: <https://prometheus.io/docs/introduction/overview/>
11. E. Deelman et al., "The evolution of the Pegasus workflow management software," *Comput. Sci. Eng.*, vol. 21, no. 4, pp. 22–36, Jul./Aug. 2019, doi: 10.1109/MCSE.2019.2919690.
12. "The directed acyclic graph manager," Univ. of Wisconsin–Madison, Madison, WI, USA, 2023. [Online]. Available: <https://research.cs.wisc.edu/htcondor/dagman/>
13. R. Tanaka and G. Papadimitriou, Jan. 2022, "Pegasus synthetic edge workflow," Zenodo, doi: 10.5281/zenodo.5889198.
14. G. Papadimitriou and S. C. Viswanath, Jan. 2022, "Pegasus casa wind workflow," Zenodo, doi: 10.5281/zenodo.5889207.
15. "The Orcasound project." Accessed: Jan. 12, 2023. [Online]. Available: <https://www.orcasound.net/>
16. G. Papadimitriou, Jan. 2022, "Pegasus orcasound workflow," Zenodo, doi: 10.5281/zenodo.5889225.
17. R. Tanaka et al., "Automating edge-to-cloud workflows for science: Traversing the edge-to-cloud continuum with pegasus," in *Proc. 22nd IEEE Int. Symp. Cluster, Cloud Internet Comput.*, Taormina, Italy, 2022, pp. 826–833, doi: 10.1109/CCGrid54584.2022.00098.
18. P. Bosshart et al., "P4: Programming protocol-independent packet processors," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 87–95, Jul. 2014, doi: 10.1145/2656877.2656890.
19. A. Esquivel Morel et al., "Network services management using programmable data planes for visual cloud computing," in *Proc. Int. Conf. Comput., Netw. Commun. (ICNC), Next Gener. Netw. Internet Appl. (NGNI)*, Honolulu, HI, USA, Feb. 2023, pp. 1–7.
20. AERPAW: Aerial Experimentation and Research Platform for Advanced Wireless. [Online]. Available: <https://aerpaw.org/experiments/>

ALICIA ESQUIVEL MOREL is with the University of Missouri, Columbia, MO, 65201, USA. Contact her at ace6qv@mail.missouri.edu.

CHENGYI QU is with the University of Missouri, Columbia, MO, 65201, USA. Contact him at cqy78@mail.missouri.edu.

PRASAD CALYAM is with the University of Missouri, Columbia, MO, 65201, USA. Contact him at calyamp@missouri.edu.

CONG WANG is with Renaissance Computing Institute (RENCI), University of North Carolina at Chapel Hill, Chapel Hill, NC, 27517, USA. Contact him at cwang@renci.org.

KOMAL THAREJA is with RENCi, University of North Carolina at Chapel Hill, Chapel Hill, NC, 27517, USA. Contact her at kthare10@renci.org.

ANIRBAN MANDAL is with RENCi, University of North Carolina at Chapel Hill, Chapel Hill, NC, 27517, USA. Contact him at anirban@renci.org.

ERIC LYONS is with the University of Massachusetts Amherst, Amherst, MA, 01003, USA. Contact him at elyons@umass.edu.

MICHAEL ZINK is with the University of Massachusetts Amherst, Amherst, MA, 01003, USA. Contact him at mzink@umass.edu.

GEORGE PAPADIMITRIOU is with the University of Southern California, Los Angeles, CA, 90089, USA. Contact him at georgpap@isi.edu.

EWA DEELMAN is with the University of Southern California, Los Angeles, CA, 90089, USA. Contact her at deelman@isi.edu.



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.



Get Published in the *IEEE Transactions on Privacy*

**This fully open access journal is
soliciting papers for review.**

IEEE Transactions on Privacy serves as a rapid publication forum for groundbreaking articles in the realm of privacy and data protection. Submit a paper and benefit from publishing with the IEEE Computer Society! With over 5 million unique monthly visitors to the IEEE Xplore® and Computer Society digital libraries, your research can benefit from broad distribution to readers in your field.

Submit a Paper Today!

Visit computer.org/tp to learn more.



Labeling “Things”

Joanna F. DeFranco  and Phil Laplante , The Pennsylvania State University

Internet of Things (IoT) devices are in our homes. Many, unbeknownst to us, intrude on our privacy. Labeling IoT products to inform consumers, offering understandable and relevant information about the “things,” is discussed.

Labels are everywhere. There are labels on supplement bottles to report percentages of the vitamins provided. Price stickers on cars show their features (for example, leather seats, Wi-Fi, high-end audio system, and so on). Most processed foods have labels for concerns such as sodium, calories, fat, country of origin, and so on. For the most part, for everyday items, there are labels, and they are usually easy to understand—in fact, government regulations standardize the content and look at these labels to make them so.

But “things,” which are the main ingredient of the Internet of Things (IoT), might not be so well understood. The IoT is not necessarily an everyday item for everyone. In the IoT, “things” could be a software system, sensor, Wi-Fi connection, device, laptop, and so on.

So therein lies the challenge and also the opportunity. Can we label IoT “things” in the same manner as we label everyday consumer products such that, for example, a system integrator of “things” knows a priori something about what the composite system will do (from a behavioral perspective)?

It is likely that the answer is yes, but only if we can determine the standard measures that offer

understandable and relevant information about “things.” So, let’s jump in and discuss this issue.

Consumer spending on IoT devices is on an upward trend and will reach US\$1 trillion if it hasn’t already. However, IoT market success should be celebrated cautiously as the security and privacy implications of bringing these trendy smart devices into your home are not insignificant.

Consumers should consider that market competition in this space sometimes causes more focus on product functionality and could shortcut the extremely important nonfunctional requirements: *security* and *privacy*. In addition, using some of these products in your home implies giving up your privacy—which, surprisingly, is not an enormous concern to many.

This article will focus on privacy specifically. Although security and privacy go hand in hand, we focused on security in a previous column and highlighted the agility of hackers. With new IoT devices, bad actors find vulnerabilities and quickly determine how to monetize them.¹ To easily differentiate these two concepts, a security vulnerability can be the situation some Ring camera owners experienced by using hacked passwords for their networks and devices (that is, hackers watching and talking to them through their cameras). Other times a privacy breach is an unintended feature of an IoT device because of video capture.

It is unfortunate that the average consumer isn’t thinking about privacy as much as product



functionality or isn't considering the privacy and security of the device he or she just purchased and placed in the home. Smart devices and the IoT introduce privacy vulnerabilities that did not exist in the past, and, more importantly, these vulnerabilities are exploitable by a much larger pool of threat actors. In the past, telephones might be wiretapped or, in rare cases, on-hook audio (obtaining audio from a landline phone even when it is not active) exploited, but these activities would have required a court order or significant skills and physical access. People who thought they might be spied on from another country through their vacuum cleaner or television would have been considered crazy then, but today they are just exercising appropriate caution.

Awareness needs to be heightened as this kind of violation is worse than a scheme to steal credit card or social security numbers. This is like a home invasion. If your location, pictures of your personal possessions, the layout of your home, and so on are all stolen, you can't change them as easily as a credit card number. Consumer awareness starts with education and sometimes legislation and laws. The awareness began with The Privacy Act of 1974, which was written to protect personally identifiable information (PII) collected by federal agencies. Privacy policies are now required to be included with products that collect PII. The policy specifies what information is collected and what will be done with it. However, how many privacy policies have you read? Is it understood that the policy doesn't mean you are protected? These policies are only a way for organizations to explain what they are doing with your information. Here are excerpts from a popular pet nanny camera (product name replaced with X):

"When you set up the X Camera, we collect any audio, video or pictures you create, upload, save or share through our Services (the

'Content'). We process Content data according to your configurations and settings. We may also collect video and audit information of individuals when they pass in front of the camera or speak when the X Camera is on."...
"We collect your geolocation data when you use our Services."

Those two things together are enough to get your house robbed. This statement could be refuted with the security measures to protect your data; however, your personal video isn't guaranteed security on that company's servers. Even if the company takes measures to keep it safe on their server, most of the time, third parties are the main security issue (that is, where your data are being sent for evaluation). In other words, the third party should be considered a weak link in the security chain.¹ Another argument might be that companies anonymize your data. Still, even if the anonymity is assured, predictive models have a high probability of revealing PII—therefore, anonymization is almost impossible.²

Another argument could be "the video in my home is of no value to a hacker." Consumers may not realize the value of these data. The value increases with companies wishing to improve their machine learning (ML) artificial intelligence (AI) algorithms. Much of the AI technology in our homes uses ML (algorithms that assess data to train the AI device). In other words, the device learns from consumers consenting to monitoring and data capturing by these these devices—inside their homes.

Because of the cost, many companies also use the AI-as-a-service model so organizations can test ML continuously on the cloud.³ What may not be realized is that part of the process involves humans to annotate (for example, categorize/label/contextualize) certain types of data. For US\$20 an hour, humans are sometimes paid to annotate pictures and videos

for ML purposes. These humans could be located anywhere in the world, and so can the ML algorithms. For example, the technology iRobot Roomba J7 images were sent to a third party that further sends the images to contracted workers to categorize the photos/video to train AI systems.⁴ Some of these images were “compromising,” of people and children inside homes. Some of these private photos ended up on social media. *Note:* The IRobot devices were “labeled,” the homeowners agreed to let the Roombas monitor them (for the purpose of AI ML), and the paid contractors also signed agreements to remove sensitive photos and video—or maybe you opted in from a privacy statement.

Here’s a scarier scenario. Any vacuuming robot or similar autonomous device, can, at your command, map each floor of your home and the placement of furniture. But it could just as easily identify other possessions, whether you have pets, estimate the number of inhabitants, create a schedule of comings and goings, and so on. In essence, the robot is conducting an ethnographic observation of your life in the home. At a minimum, this information can be used for marketing purposes, unwanted solicitations, and brushing scams (where sellers create fake accounts and order their own products to an address) or more nefariously to plan for home invasions, robberies, blackmail, and more.

LABELS

The United States has recognized the privacy issues concerning consumer-facing devices. An executive order was issued on 12 May 2021 on “Improving the Nation’s Cybersecurity,” addressing securing software development environments (<https://www.federalregister.gov/documents/2021/05/17/2021-10460/improving-the-nations-cybersecurity>). Part of the order outlines that the National Institute of Standards and Technology (NIST) will initiate pilot programs, informed by existing consumer product labeling programs, in an effort to educate consumers on the security capabilities of IoT devices and software development practices. The task includes incentivizing manufacturers to participate. In addition, together with the Federal Trade Commission, NIST is identifying IoT cybersecurity criteria for a consumer labeling program as well as

secure software development practices or criteria for a consumer software labeling program. Subsequently, there is also forthcoming legislation, called the “Informing Consumers About Smart Devices Act,” which will require manufacturers to let consumers know if there is a microphone or camera in an Internet-connected product.

Labeling isn’t a new concept. Voas (2000) proposed software warranties or certifications to address software quality due to the differing types of software and target environments.⁵ He suggested a framework for a certification to address the software assurance and integrity needs of the organization as well as a way to highlight the peculiarities of that software type. In 2021, Laplante recommended software labels to offer a consistent and coordinated way to assess the level of risk in software to decide if it needs to be labeled, like a food, drug, or hazardous material. It was further suggested that a label could expose important properties of the software to review its safety, security, privacy, and reliability.⁶ For example, information should be available to the developers when reusing software components, such as something similar to a food ingredient label: amount of reused (modified and unmodified) code, amount of new code (handwritten and auto-generated by tools), amount of open source code (and type of license), software complexity, testing methodology, and so on. Additionally, the software could be labeled for carbon (power) consumption—perhaps a simple green–yellow–orange–red system for excessive power consumption in typical or exceptional operation profiles.

NIST has provided a white paper in response to the executive order, called “Recommended Criteria for Cybersecurity Labeling for Consumer Internet of Things (IoT) Products” (<https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.02042022-2.pdf>). This work discussed a binary label, indicating if a product has met a baseline security standard. The label could be in the form of a uniform resource locator or a scannable (QR) code that would lead the consumer to additional information, such as (summarized)

- › intent and scope: to address potential misinterpretations
- › product criteria: cybersecurity properties

included in the baseline and how the criteria address, for example, security risks

- › a glossary of technical terms written in simple English
- › conformity assessment: evaluation of cybersecurity properties
- › declaration of conformity: referring to the baseline criteria, including the date of the last label
- › scope: the kinds of products eligible for the label and information to identify labeled products
- › changing applicability: the current state of this product's labeling as new cybersecurity threats and vulnerabilities emerge
- › security considerations and implications for end-of-life IoT products
- › expectations for consumers: consumer responsibility in securing software and how their actions (or inactions) can impact the software's cybersecurity
- › contact information for the labeling program.

The European Union's new Cyber Resilience Act will require manufacturers to provide consumer information on the security features of devices and how to securely configure them.⁷ Similarly, the Cyber Security Agency of Singapore (CSA) launched the Cybersecurity Labelling Scheme for consumer smart devices (<https://www.csa.gov.sg/Programmes/certification-and-labelling-schemes/cybersecurity-labelling-scheme/about-cls>). Finland and Germany have signed an agreement indicating they recognize the label issued by the CSA.

Labeling provides awareness and education not only to consumers, but to the developers creating these devices. This entire effort could start small with popular devices, such as electronic doorbells and home security cameras, to include multiple risk levels.

The bottom line is that IoT device buyers need to be aware of the risk involved in utilizing these devices in their homes, and developers need a reminder of what is important to include in these products. In addition to labeling, improving the process of data capture and analysis should be addressed: How can human involvement be made more efficient and safer? 🤖

ACKNOWLEDGMENT

The authors are appreciative of the valuable input for this article from Jeff Voas and Rick Kuhn at NIST.

REFERENCES

1. J. DeFranco and B. Maley, "Closing the security agility gap," *Computer*, vol. 55, no. 8, pp. 100–102, Aug. 2022, doi: 10.1109/MC.2022.3169400.
2. N. Kshetri and J. DeFranco, "Is privacy dead?" *IEEE IT Professional*, vol. 22, no. 5, pp. 4–12, Oct. 2020, doi: 10.1109/MITP.2020.2992148.
3. V. Dey, "AI-as-a-service making artificial intelligence and data analytics more accessible and cost effective," VentureBeat, San Francisco, CA, USA, Dec. 2022. Accessed: Jan. 2, 2023. [Online]. Available: <https://venturebeat.com/ai/ai-as-a-service-makes-artificial-intelligence-and-data-analytics-more-accessible-and-cost-effective/>
4. E. Guo, "A Roomba recorded a woman on the toilet. How did the screenshots end up on Facebook?" *MIT Technol. Rev.*, Dec. 2022. [Online]. Available: <https://www.technologyreview.com/2022/12/19/1065306/roomba-irobot-robot-vacuums-artificial-intelligence-training-data-privacy/>
5. J. M. Voas, "Limited software warranties," in *Proc. 7th IEEE Int. Conf. Workshop Eng. Comput.-Based Syst. (ECBS)*, Apr. 2000, pp. 56–61, doi: 10.1109/ECBS.2000.839861.
6. P. Laplante, "Software labels," *Computer*, vol. 54, no. 11, pp. 82–86, Nov. 2021, doi: 10.1109/MC.2021.3102360.
7. M. Nelson, "EU announces first ever move to legislate cybersecurity for IoT," *IoT Bus. News*, Oct. 2022. [Online]. Available: <https://iotbusinessnews.com/2022/10/12/63479-eu-announces-first-ever-move-to-legislate-cybersecurity-for-iot/>

JOANNA F. DEFRANCO is an associate professor of software engineering, associate director of the D.Eng. in Engineering program at The Pennsylvania State University, Malvern, PA 19355 USA, and an associate editor in chief of *Computer*. Contact her at jfd104@psu.edu.

PHIL LAPLANTE is a professor of software and systems engineering at The Pennsylvania State University, State College, PA 16801 USA, a Fellow of IEEE, and an associate editor in chief of *Computer*. Contact him at laplante@psu.edu.


DEPARTMENT: INTERNET ETHICS

Can We Explain Privacy?

Gönül Aycı , Bogazici University, Istanbul, 34342, Turkey

Arzucan Özgür , Bogazici University, Istanbul, 34342, Turkey

Murat Şensoy , Amazon Alexa AI, EC2A 2FA, London, U.K.

Pinar Yolum , Utrecht University, 3584CC, Utrecht, The Netherlands

Web users want to protect their privacy while sharing content online. This can be done through automated privacy assistants that are capable of taking actions by detecting privacy violations and recommending privacy settings for content that the user intends to share. While these approaches are promising in terms of the accuracy of their privacy decisions, they lack the ability to explain to the end user why certain decisions are being made. In this work, we study how privacy assistants can be enhanced through explanations generated in the context of privacy decisions for the user content. We outline a methodology to create explanations of privacy decisions, discuss core challenges, and show example explanations that are generated by our approach.

Millions of pictures and videos are being shared on social media platforms every day. Our personal data as well as the private content that we create circulate on the Web in ways we haven't imagined before. More and more, cloud services are the go-to locations for storing data, computing, and sharing content. While these services have a lot of benefits for end users, they may use many other third-party services to deliver value, and this may pose unprecedented privacy challenges.

The main method for handling privacy with these services is through consent, where the service provides information on how it will make use of the content, including the purpose and further processing that will be involved and the user of the service is asked to accept the conditions. The informed consent is aimed to detect whether personal data has been leaked or used against the person's will. Current models and implementations of consent prove cumbersome for users. The General Data Privacy Regulation (GDPR)¹ governs privacy and consent in Europe and is

influencing other jurisdictions. GDPR requires services to provide explanations, but those explanations are usually long texts. The users are not always clear if their declining to give consent will result in what parts of the service not being available. The services a user engages are numerous and include social media working with documents on the cloud. Thus, many users lack the capacity to even read the text to which they are giving consent.²

Privacy assistants can work with humans to help them with tasks related to managing their privacy.^{3,4} As users share content, it is necessary to think for whom the content is meant and how to configure its privacy settings.⁵ An important category of such content is images. Recent work helps users categorize whether a given image is private or not.⁶ This could be useful to help users avoid unintended sharing of private images on social networking sites. Privacy is personal and subjective: What one person identifies as private might be different from that of another. Thus, an assistant ought to provide personalized answers as to whether an image is private or not.^{7,8} For these privacy assistants to be adopted by end users, they need to be trusted. One important path to induce such trust is through explanations.⁹ We address a novel problem

concerning explanations and privacy: How can a privacy assistant explain why it identifies a certain piece of content as private or public?

EXPLANATIONS FOR PRIVACY

Explainable artificial intelligence offers methods that can help humans understand how algorithms work. Most of these methods are targeted to explain how a machine learning algorithm makes a classification. When the classification pertains to an image, visual explanations are useful. An example of a visual explanation would be highlighting the most relevant region in the image (e.g., highlighting a child in an image as to show that there is one). Saliency maps and attention maps are important tools for such explanations.¹⁰ However, these techniques are not immediately applicable to explaining privacy decisions. Specifically, an image is not categorized private or public because of a single segment. For example, an image of a child at home with her parents might be categorized as private, while an image of the same child participating in a school performance on stage might be categorized as public. Hence, identifying and highlighting the child in the image will not adequately explain why this image is private or public.

Another class of explanation techniques works on binary classification and considers what features of the input have been influential on the decision.¹¹ Thus, they provide a handle to interpret the decision. For example, these methods could say which features had an effect on classifying an image as private and what the strength of these features were. If the features are derived from images automatically, the features might not always translate to concepts that users would understand. Such knowledge may help algorithm developers but is not meaningful for end users. An alternative would be to use image tags for classification. While the tags themselves are understandable for the end user, the number of tags makes it difficult to generate succinct explanations.

METHODOLOGY FOR EXPLAINING PRIVACY

We propose to use the concept of *topics* as a way to capture explanations. We envision that each image belongs to multiple topics with different strengths, where the interplay between the topics leads to understanding why the image is private. For end users, the explanation needs to be either visual or short text and should touch on the most important aspects, rather than giving a comprehensive analysis of features. An explanation as to why an image is private or public will be described through a carefully selected subset of

topics that the image belongs to. In order to realize such explanations, we need to understand how we can associate images with topics and how we can decide on which topics to use for explanation.

Our proposed methodology has the following steps:

- 1) Start with a set of images already labeled as public or private. This could be the set of images that the user herself previously made a decision to share or not to share online.
- 2) Assign tags (i.e., keywords) for each image to describe its content. These tags could be provided by the users or can be generated automatically with a tool, such as Clarifai.¹²
- 3) Perform topic modeling. Each topic should pertain to images that could be described with similar tags. At the same time, each topic should be different from each other. Topic modeling is a technique that discovers latent topics within a collection of textual information, in this case, tags associated with images. As a result, each image is associated with one or more topics.
- 4) Create topic descriptions. Topics are intuitively meaningful and interpretable for humans. In order to improve the understanding for the user, it is useful to name the topics, for example, *Nature*, *Child*, and *Fashion*. This can be done manually as well as automatically. Different automated approaches can be applied such as identifying the most similar word to the tags as the topic name, where similarity can be computed by using the word embeddings of the tags or by using an ontology such as WordNet.¹³
- 5) Machine learning classification. Using the generated set of topics as features, it is necessary to train an interpretable machine learning model to perform binary classification on the images. Given a new image that is associated with topics, the classifier will assess whether it is private or public.
- 6) Evaluate the effect of each topic on the classification and determine which topics will be used for explanation. This can be done by different heuristics; for example, by identifying a largely influential single topic or multiple topics that make smaller contributions as well as identifying topics that have opposing classifications.
- 7) Create textual and visual explanation templates based on the interplay between topics. If only a single topic is influential, it is enough to mention that topic. If, on the other hand, opposing topics are present, the text should express their relation to each other.

Generating Topics

For Step 1, we use a balanced subset of the publicly available PicAlert dataset¹⁴, which is widely used for the privacy prediction for images.^{7,15} PicAlert contains Flickr images that are labeled as *private* or *public* by annotators. We consider an image as private if at least one annotator has annotated it as private and public only if all the annotators have annotated it as public. The balanced subset we work with contains 32,000 samples, comprising 27,000 training images and 5000 test images. For Step 2, we use Clarifai application programming interface¹² to automatically generate 20 different tags for each image. For Step 3, we explore 20 different latent topics using *Non-negative matrix factorization*¹⁶ as the topic modeling technique and name the topics based on the dominant keywords (Step 4).

To evaluate the representation of the images with the topics extracted using non-negative matrix factorization (NMF), we train a random forest classifier where the images are represented as term frequency-inverse document frequency (TF-IDF) vectors of these topics (Step 5). We study the precision, recall, and F1 score of the classifier for the private and public class separately. This is important because in many settings the cost of misclassifying a private image might be higher than that of misclassifying a public image.⁸ We observe that the scores for both classes are similar. For the private and public class, the classifier obtains a precision of 87% and 89%, recall of 90% and 87%, and F1 score of 89% and 88%, respectively. Overall, the accuracy of the classifier on the test set is 88%, indicating that the NMF-extracted topics are effective for privacy prediction. This performance matches that of state-of-the-art approaches for image privacy prediction^{6,8} and thus can be used to generate the explanations.

Generating Explanations From Topics

Even though now we have access to the topics associated with each image and that they are successful in classifying the images, generating explanations from this is still challenging. First, many topics are associated with each image; thus, listing all relevant topics is meaningless. Second, the topics that are mostly associated with the image do not have a clear prediction. Some topics such as *People* are associated more frequently with the private class, whereas others like *Nature* are associated more frequently with the public class. Note that although some topics are associated more frequently with one class, the topics do not have an explicit class to which they belong. Therefore, the topic itself does not directly signal a certain class, and as such it is not straightforward to generate an explanation

for the decision by simply looking at its class. For example, the *Performance* topic was associated with 35% of the private images as well as 30% of the public images and the *Competition* topic was associated with 37% of the private images as well as 25% of the public images. Thus, we need to consider to what extent a topic is related to the image as well as how the different topics come together in an image to explain privacy.

The TreeExplainer¹⁷ model provides the contributions of each feature in terms of Shapley values, which affect the model output of tree-based algorithms. A feature with a positive Shapley value denotes that the existence of that feature was influential in the prediction and vice versa for the negative value. For us, each feature corresponds to a topic. We obtained these topic-value pairs from the TreeExplainer. We remove the topics that are not relevant for the image based on its TF-IDF value as well as filter out the topics with values smaller than a threshold (Step 6). The explanations differ on how the remaining topics are related to each other as follows:

Single: It could be that a single topic defines the classification or

Multiple: that multiple topics make small contributions to classification or

Combination: combinations of various topics describe the classification.

For each type, we formulate a short text template and a visual that lists topics and tags that were important for the classification (Step 7).

EXAMPLE EXPLANATIONS

We consider two explanation templates that differ on how the topics relate to each other. The first template pertains to a case where the image has a number of topics, such that none of the topics by themselves would necessarily derive that the image would be of a given target class. However, the existence of multiple topics supports that the image should be of a particular target class. Figure 1 shows an example image that has been identified as private and the explanation generated by our algorithm. Three topics that are relevant to the image are provided with the tags that are important for the classification.

All three topics together strengthen the decision. An image could (and usually does) have many topics associated with it. However, to keep the explanations simple, we select only those topics that contribute the most to the decision. At the same time, it could be possible that the lack of certain topics would affect the underlying decision. For example, the fact that the image is not related to the topic *Outside* might have




FIGURE 1. Example image classified as private and the generated explanation.

been important in the classification. However, we do not include absent topics as part of the explanation.

The second template pertains to a case where the topics associated with an image do not always agree on whether the image should be private or public. In such situations, the explanation should indicate such opposing evidence as well as how and if it was resolved. One can naively expect that if the image is associated strongly with at least one private topic, then the privacy label of the image would also be private. However, interestingly private concepts when integrated into public space can yield images that are public.

Figure 2 shows an example image that has been identified as public and the generated explanation. Even though the topic *People* pushes this image to be classified as private, the fact that it is situated in the *Art* topic makes it public. By observing that the influence of the *Art* topic is larger than that of *People*, the algorithm can generate the explanation on the right side of Figure 2. Note that the explanation template

now is different than that of Figure 1 and reflects that there were opposing topics involved. These examples demonstrate how the interaction between various topics affect the outcome of the classification and thus the explanation that needs to be generated.

DIRECTIONS

The methodology that we propose generates explanations for end users to understand why a given image would be classified as private or public. It is based on exploring hidden topics using topic modeling from descriptive tags of images. It captures explanation templates that are based on the relationship between images and their associated topics and generates explanations automatically. An important direction is to design other explanation templates based on the interplay between topics that push the classification to private or public. Currently, we do not differentiate between topic characteristics; however, some topics, such as *Nature* or *Room*, pertain to the location context, whereas some topics, such as *Competition* or



FIGURE 2. Example image classified as public and visual explanation that shows the topics and relevant tags.

Performance denote public spaces. Capturing the semantics of these topics could lead to more detailed and better structured explanations of privacy. Our previous work focused on uncertainty and confidence of predictions,⁸ and in this work, we explain model predictions. These two directions for enhancing privacy decisions are complementary and combining them may help the user assess the explanations better. An important direction for future work is to be able to get feedback from end users and update the explanations. With a user study, we plan to understand if the generated explanations make sense to people and if the participants find the explanations useful. This would bring us closer to understand what other aspects need to go into explanations to make them viable for end users. 🌈

ACKNOWLEDGMENTS

The first author is supported by the Scientific and Technological Research Council of Turkey (TÜBİTAK) and Turkish Directorate of Strategy and Budget under the TAM Project number 2007K12 – 873. This research was partially funded by the Hybrid Intelligence Center, a 10-year program funded by the Dutch Ministry of Education, Culture, and Science through the Netherlands Organization for Scientific Research. This work does not relate to Şensoy's position at Amazon.

REFERENCES

1. C. J. Hoofnagle, B. Van Der Sloot, and F. Z. Borgesius, "The European Union general data protection regulation: What it is and what it means," *Inf. Commun. Technol. Law*, vol. 28, no. 1, pp. 65–98, Jan. 2019, doi: 10.1080/13600834.2019.1573501.
2. T. Vila, R. Greenstadt, and D. Molnar, "Why we can't be bothered to read privacy policies," in *Economics of Information Security*, L. J. Camp and S. Lewis, Eds. Boston, MA, USA: Springer-Verlag, 2004, pp. 143–153.
3. R. L. Fogues, P. K. Murukannaiah, J. M. Such, and M. P. Singh, "SoSharP: Recommending sharing policies in multiuser privacy scenarios," *IEEE Internet Comput.*, vol. 21, no. 6, pp. 28–36, Nov./Dec. 2017, doi: 10.1109/MIC.2017.4180836.
4. J. Colnago et al., "Informing the design of a personalized privacy assistant for the Internet of Things," in *Proc. CHI Conf. Human Factors Comput. Syst.*, 2020, pp. 1–13, doi: 10.1145/3313831.3376389.
5. O. Ulusoy and P. Yolum, "Panola: A personal assistant for supporting users in preserving privacy," *ACM Trans. Internet Technol.*, vol. 22, no. 1, pp. 1–32, Sep. 2021, doi: 10.1145/3471187.
6. A. Tonge and C. Caragea, "Image privacy prediction using deep neural networks," *ACM Trans. Web*, vol. 14, no. 2, pp. 1–32, Apr. 2020, doi: 10.1145/3386082.
7. A. Can Kurtan and P. Yolum, "Assisting humans in privacy management: An agent-based approach," *Auton. Agents Multi-Agent Syst.*, vol. 35, no. 1, pp. 1–33, Apr. 2021.
8. G. Ayçi, M. Şensoy, A. Özgür, and P. Yolum, "Uncertainty-aware personal assistant for making personalized privacy decisions," *ACM Trans. Internet Technol.*, vol. 23, no. 1, pp. 1–24, Mar. 2023, doi: 10.1145/3561820.
9. F. Mosca and J. Such, "An explainable assistant for multiuser privacy," *Auton. Agents Multi-Agent Syst.*, vol. 36, no. 1, pp. 1–45, Apr. 2022, doi: 10.1007/s10458-021-09543-5.
10. R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vision*, 2017, pp. 618–626.
11. M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?: Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.
12. Clarifai. [Online]. Available: <https://clarifai.com/clarifai/main/models/general-image-recognition>
13. G. A. Miller, "Wordnet: A lexical database for English," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, Nov. 1995, doi: 10.1145/219717.219748.
14. S. Zerr, S. Siersdorfer, J. Hare, and E. Demidova, "Privacy-aware image classification and search," in *Proc. 35th Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2012, pp. 35–44, doi: 10.1145/2348283.2348292.
15. A. Squicciarini, C. Caragea, and R. Balakavi, "Toward automated online photo privacy," *ACM Trans. Web*, vol. 11, no. 1, pp. 1–29, Apr. 2017, doi: 10.1145/2983644.
16. D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, Oct. 1999, doi: 10.1038/44565.
17. S. M. Lundberg et al., "From local explanations to global understanding with explainable AI for trees," *Nature Mach. Intell.*, vol. 2, no. 1, pp. 56–67, Jan. 2020, doi: 10.1038/s42256-019-0138-9.

GÖNÜL AYCI is a Ph.D. student at the Computer Engineering Department, Bogazici University, Istanbul, 34342, Turkey, and a visiting Ph.D. Researcher at Utrecht University. Her

research interests include designing and developing privacy assistants. Ayçi received her M.Sc. degree in computer science from Ozyegin University, Istanbul, Turkey. Contact her at gonul.ayci@boun.edu.tr.

ARZUCAN ÖZGÜR is a faculty member at the Computer Engineering Department, Bogazici University, Istanbul, 34342, Turkey. Her research interests include natural language processing and bioinformatics. Özgür received her Ph.D. degree in computer science and engineering from the University of Michigan. Contact her at arzucan.ozgur@boun.edu.tr.

MURAT ŞENSOY is an applied research scientist at Amazon Alexa AI, EC2A 2FA, London, U.K. His research interests include reliable machine learning systems. Şensoy received his Ph.D. degree in computer engineering from Bogazici University. Contact him at drmuratsensoy@gmail.com.

PINAR YOLUM is a professor in information and computing sciences at Utrecht University, 3584CC, Utrecht, The Netherlands. Her current research interests include trustworthy AI with an emphasis on privacy. Yolum received her Ph.D. degree in computer science from North Carolina State University. Contact her at p.yolum@uu.nl.



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



How to “Sell” Ethics (Using AI): An Interview With Alexander Serebrenik

Tim Menzies 

FROM THE EDITOR

In this issue, our column is a little more “AI for SE” than “SE for AI.” Here, I interview Prof. Alexander Serebrenik to discuss how software analytics (including machine learning algorithms) lets us better understand issues of diversity and exclusion in software engineering (SE) development teams.

But for future issues, what do you want to see in this “SE for AI” column? Do you have a surprising result or industrial experience? Something that challenges decades of conventional thinking in software engineering? If so, e-mail a one-paragraph synopsis to tim@menzies.us (subject line: “SE for AI: Idea: [Your Idea]”). If that looks interesting, I’ll ask you to submit a 1,000–3,000-word article (where each graph, table, or figure is worth 250 words) for review for *IEEE Software*. Note: Heresies are more than welcome (if supported by well-reasoned industrial experiences, case studies, or other empirical results).—Tim Menzies

“Most organizations are tone deaf when it comes to ethics,” says Prof. Alexander Serebrenik (Figure 1) of the Eindhoven University of Technology (<https://tue.academia.edu/AlexanderSerebrenik>). “I’ve been trying to talk discrimination, diversity, and inclusion with them for years, and frankly, I’ve given up”

Instead, he suggests, it is better to talk about productivity and how “community smells” can lead to bad “code smells” (a “code smell” is a characteristic in the source code that possibly indicates a deeper problem¹). “At least where I work,” says Dr. Serebrenik, “we live in a neoliberal society where everything is about money. So, to make them behave morally, I have to make an economic argument.”

“It’s software engineering [SE] 101,” he says. “Jim Herbsleb has been telling us for decades² (and he is right when he says it) that when development teams are not communicating effectively, then the software quality suffers. Organizations shoot themselves in the foot when they discriminate against certain developers because of their age³ or race or gender identity⁴ or other diversity issues.”

He argues that watching for discrimination and exclusion should be everyone’s job. “Developers can observe more detailed behavior than managers. Luckily, many developers have a professional pride in their work and truly care about the quality of their code. They know that developers can screw up code if they do not talk to each other.”

And managers have a special role to play. “It should be management 101—but it often isn’t—managers should be aware of the interplay between their developers and how those interactions can damage the

code.” To help in that task, Dr. Serebrenik is part of a large international team that has shown how “community smells” can cause “code smells.” That team argues convincingly that managers should watch out for community smells such as

- › *Organizational silo*: siloed areas of the developer community that do not communicate, except through one or two of their respective members
- › *Lone wolf*: unsanctioned or defiant contributors who carry out their work irrespective or regardless of their peers, their decisions, and their communications
- › *Radio silence*: an instance of the “unique boundary spanner” problem from social networks analysis: one member interposes themselves into every formal interaction across two or more subcommunities with little or no flexibility to introduce other parallel channels
- › *Black cloud*: information overload due to the lack of structured communication or cooperation governance.

SHOULD WE DRESS UP ETHICAL ISSUES AS “MERE” ECONOMIC ISSUES RELATED TO CODE QUALITY?

These community smells are clearly detrimental to code quality. Based on extensive research and some large-scale empirical studies, Dr. Serebrenik and colleagues have shown that these issues are related to many code bad smells (see Figure 2).^{5,6}

Dr. Serebrenik passionately believes that socio-technical perspectives are central to the continued success of SE. “Consider self-admitted technical debt,” he says. “How do we make it right for someone to admit something is wrong? What makes people admit that something is a shortcut solution that should be fixed in the future? I find this a fascinating question since, I hope, it will let us understand more about how we all contribute to productivity and well-being.”

As to the ongoing research implications of this work, Dr. Serebrenik says that research into developer diversity and inclusion can show how to resolve community and code smells. But he says we must move



FIGURE 1. Prof. Alexander Serebrenik of the Eindhoven University of Technology.

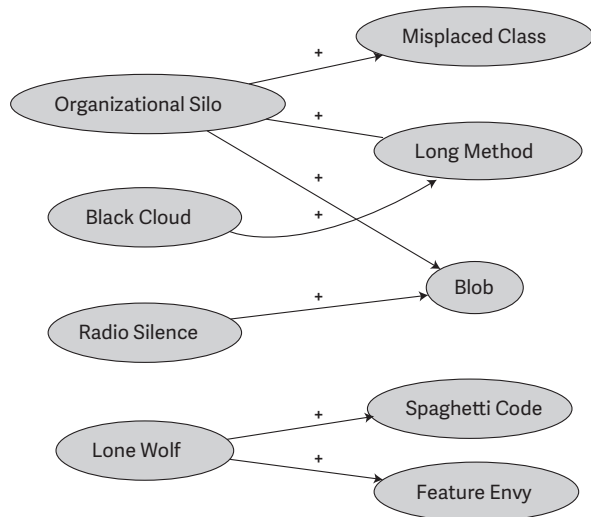


FIGURE 1. Community bad smells lead to code smells. (Source: Palomba et al.⁶)

forward cautiously. In a recent *IEEE Software* editorial,⁷ he and his colleagues wrote, “The techniques we use run the risk of oversimplification. Race and ethnicity are not globally applied uniform concepts, so we face a question: how can we collect data about race and ethnicity within global software development in our attempt to increase the percentage of people who are not white and who thus do not have

a privileged position?” To address that issue, Dr. Serebrenik discusses “inclusive data collection” methods that allow people to describe themselves in a more open-ended multifaceted manner. (The specific case of gender-inclusive data collection is discussed by Scheuerman et al.⁸)

Somewhat nervously, I asked Dr. Serebrenik if it was wrong, perhaps even pandering to industry, for researchers to dial back their rhetoric on diversity and inclusion. Should we dress up ethical issues as “mere” economic issues related to code quality?

In reply, he offered a very pragmatic answer, tuned by years of work in this field. “It is really an issue about what message different audiences are ready to hear,” he says. “The cost of poor communication in software development has been guesstimated to be US\$37 billion per year. Why not use that fact as a way to make organizations improve on how developers interact with each other?” And as to other audiences, even a cursory glance at his recent publications shows that Dr. Serebrenik presents his diversity message, loud and clear, for all to hear.

Dr. Serebrenik continues to work zealously and rigorously on issues of diversity and inclusion in SE. For more on this line of work, including methodological interventions that can incentivize changes to development teams and SE education, watch for the upcoming book *Equity, Diversity, and Inclusion in Software Engineering* (to appear in late 2023, from Apress Books).

As for Dr. Serebrenik himself, he is very hopeful for the future. “I might be overly optimistic on this point, but my feeling is that there is much recent discussion on this issue (diversity and inclusion). I’m disappointed, sure, that it is not going as fast as I would like it to go. But looking back at where we were just a decade ago, we can say that we are doing much better than before.” 🤞

REFERENCES

1. “CodeSmell.” Martin Fowler. [Online]. Available: <https://martinfowler.com/bliki/CodeSmell.html>
2. J. Herbsleb. *ICSE 2014 Inaugural Session Keynote: Socio-Technical Communication*. (2014). [Online Video]. Available: <https://www.youtube.com/watch?v=v0CSnYvd0C4>
3. S. Baltes, G. Park, and A. Serebrenik, “Is 40 the new 60?

How popular media portrays the employability of older software developers,” *IEEE Softw.*, vol. 37, no. 6, pp. 26–31, Nov./Dec. 2020, doi: 10.1109/MS.2020.3014178.

4. J. C. Carver, H. Muccini, B. Penzenstadler, R. Prikladnicki, A. Serebrenik, and T. Zimmermann, “Behavioral science and diversity in software engineering,” *IEEE Softw.*, vol. 38, no. 2, pp. 107–112, Mar./Apr. 2021, doi: 10.1109/MS.2020.3042683.
5. G. Catolino, F. Palomba, D. A. Tamburri, A. Serebrenik, and F. Ferrucci, “Gender diversity and women in software teams: How do they affect community smells?” in *Proc. IEEE/ACM 41st Int. Conf. Softw. Eng., Softw. Eng. Soc. (ICSE-SEIS)*, Montreal, QC, Canada, 2019, pp. 11–20. [Online]. Available: <https://www.win.tue.nl/~aserebre/ICSE2019SEIS-Gemma.pdf>, doi: 10.1109/ICSE-SEIS.2019.00010.
6. F. Palomba, D. Tamburri, F. Arcelli Fontana, R. Oliveto, A. Zaidman, and A. Serebrenik, “Beyond technical aspects: How do community smells influence the intensity of code smells?” *IEEE Trans. Softw. Eng.*, vol. 47, no. 1, pp. 108–129, Jan. 2021, doi: 10.1109/TSE.2018.2883603.
7. K. Albusays et al., “Storey: The diversity crisis in software development,” *IEEE Softw.*, vol. 38, no. 2, pp. 19–25, Mar./Apr. 2021, doi: 10.1109/MS.2020.3045817.
8. M. K. Scheuerman, K. Spiel, O. L. Haimson, F. Hamidi, and S. M. Branham. “HCI guidelines for gender equity and inclusivity.” Morgan-Klaus. [Online]. Available: <https://www.morgan-klaus.com/gender-guidelines.html>



TIM MENZIES is a full professor at North Carolina State University, Raleigh, NC 27606 USA. Contact him at tim@ieee.org.





CALL FOR SPECIAL ISSUE PROPOSALS

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

Please send us your high-quality proposals for the 2025–2026 editorial calendar. Of particular interest are proposals centered on:

- 3D printing
- Robotics
- LLMs
- AI safety
- Dis/Misinformation
- Legacy software
- Microelectronics

Proposal guidelines are available at:

www.computer.org/csdl/magazine/co/write-for-us/15911



DEPARTMENT: SE AND ETHICS

Ethics: How Far Have We Come?

Brittany Johnson  and Tim Menzies 

FROM THE EDITORS

Ethics in software engineering is a vast spectrum, spanning concerns from the foundational to the futuristic. It touches on how we design systems, interact with societal norms, and anticipate the consequences of our innovations. —*Brittany Johnson and Tim Menzies*

As we navigate into 2025, this column reflects on where we stand in addressing these challenges and what the big picture reveals about our collective responsibilities. What have we achieved, and what remains uncharted territory?

ETHICS AND LAW

Ethics in software engineering cannot exist in isolation from the broader societal constraints of legal and economic systems. These frameworks shape what is possible, permissible, and prioritized in software engineering (SE) projects. For example, laws governing data privacy influence how engineers design databases, while economic incentives can drive choices that may conflict with ethical practices. As Marc Canellas points out, “Engineers must actively participate in advocating for systemic change to ensure that legal and economic systems align with ethical technological progress.”¹

The transition from legal constraints to environmental impacts highlights the interconnected nature of ethical responsibilities across SE domains.

POWER ENGINEERING

The ethical implications of energy-intensive systems, such as data centers and blockchain technologies, demand urgent attention from software engineers. Federica Sarro highlights that “AI technologies are significant contributors to global energy consumption,” and she emphasizes that “software engineers must prioritize the design of energy-efficient systems to mitigate environmental impacts.”² SE teams must consider power efficiency as a core nonfunctional requirement, balancing performance with environmental sustainability.

From resource consumption, it is important to reflect on how ethical frameworks shape the decision-making processes within SE projects.

CRUTCH OR SUPERPOWER?

Ethical frameworks, if applied without deep understanding, risk becoming rote exercises devoid of meaningful impact in SE practice. Engineers must internalize these principles through practice and reflection. Otherwise, we risk a “ChatGPT effect,” where widespread usage generates little understanding and pervasive errors. As Chakraborty et al. note, “Superficial application of ethical guidelines often results in overconfidence in the outputs of AI systems, masking underlying biases and errors.”³ For SE practitioners,

this underscores the need to build ethical awareness into development workflows, such as during code reviews or model validation.

From individual understanding to systemic biases, fairness in AI offers another lens for ethical examination in SE.

BIAS AND FAIRNESS IN AI

Machine learning systems often inherit and amplify biases present in training data or design decisions, which software engineers must address. For example, Galhotra et al. describe methods for fairness testing in software systems, offering tools to detect discrimination and ensure ethical decision-making.⁴ Similarly, Themis, a tool introduced by Angell and colleagues, “automatically tests software for discriminatory behavior, providing developers with actionable insights to mitigate bias.”⁵ These tools highlight how SE practices can integrate fairness checks into the development pipeline to create more equitable systems.

Bias mitigation efforts naturally intersect with concerns around privacy, forming another critical axis of ethical responsibility in SE.

PRIVACY AND DATA PROTECTION

Ethical dilemmas surrounding user data persist, from issues of consent to challenges in anonymization and risks of data breaches. For SE practitioners, this means designing systems that prioritize user trust and comply with data protection regulations, such as General Data Protection Regulation, while maintaining scalability and efficiency.

From safeguarding privacy, we now turn to ensuring accessibility for all users.

ACCESSIBILITY AND INCLUSION

Despite decades of progress, accessibility remains underprioritized in many software projects. Ethical software engineering must ensure that systems are usable by all individuals, including those with disabilities. Johnson and Smith emphasize that “inclusive design is not just a technical challenge but an ethical mandate to ensure equity in digital spaces.”⁶ By embedding accessibility into development frameworks—such as through automated testing tools for Web Content Accessibility Guidelines compliance—SE teams

can make inclusivity a default feature rather than an afterthought.

Next, we consider the ethical implications of collaboration and ownership in SE projects.

INTELLECTUAL PROPERTY AND OPEN SOURCE ETHICS

The intersection of collaboration and commercialization in the open source community raises ethical questions about ownership, attribution, and fair use. Menzies and Johnson highlight that “open source ethics must evolve to balance the interests of individual

THE TRANSITION FROM LEGAL CONSTRAINTS TO ENVIRONMENTAL IMPACTS HIGHLIGHTS THE INTERCONNECTED NATURE OF ETHICAL RESPONSIBILITIES ACROSS SE DOMAINS.

contributors and commercial entities,” ensuring fair attribution while fostering innovation.⁷ For SE teams, these issues influence decisions about licensing, contribution policies, and the integration of third-party libraries into proprietary software.

Turning from ownership to accountability, we explore how SE handles ethical responsibility in autonomous systems.

RESPONSIBILITY FOR ALGORITHMIC DECISIONS

When software systems make decisions with harmful consequences, who is accountable? This question becomes critical as autonomous systems like self-driving cars and health care tools grow in prevalence. As Angell et al. suggest, “Algorithmic accountability frameworks are essential to delineate responsibility and ensure trust in autonomous systems.”⁵ For SE practitioners, this involves implementing mechanisms for auditability, traceability, and explainability within software, ensuring that stakeholders understand and can address unintended consequences.

Finally, we turn to the broader societal implications of software misuse.

SURVEILLANCE AND MISUSE OF SOFTWARE

Software that enables mass surveillance or invasive data collection poses significant ethical risks. Engineers must grapple with the implications of their work, questioning whether the systems they design align with societal values. Menzies and Hazard caution that “unchecked development of surveillance technologies can erode public trust and infringe on fundamental rights.”⁸ For SE teams, this calls for adopting ethical guidelines that explicitly prohibit the misuse of their software, supported by rigorous internal reviews and public accountability measures.

REFLECTIONS AND FUTURE DIRECTIONS

Ethics in software engineering is not a static concept but a continually evolving challenge. As we reflect on the spectrum of concerns discussed—from legal constraints to algorithmic accountability—we must also ask: what are we missing? Emerging technologies, such as quantum computing or brain-computer interfaces, bring new ethical questions to the table. How do we prepare for the unforeseen, and what structures do we need to ensure ethics remains central to innovation? As SE practitioners, our goal must be to embed ethical considerations into the fabric of our discipline, driving progress that is not only technologically impressive but also morally sound. 🧠

REFERENCES

1. B. Johnson and T. Menzies, “Fighting for what’s right: An interview with Marc Canellas,” *IEEE Softw.*, vol. 41, no. 2, pp. 104–107, Mar./Apr. 2024, doi: 10.1109/MS.2023.3340928.
2. T. Menzies and B. Johnson, “Powering down: An interview with Federica Sarro on tackling energy consumption in AI-powered software systems,” *IEEE Softw.*, vol. 41, no. 5, pp. 89–92, Sep./Oct. 2024, doi: 10.1109/MS.2024.3410011.
3. S. Galhotra, Y. Brun, and A. Meliou, “Fairness testing: Testing software for discrimination,” *Proc. 11th Joint Meeting Found. Softw. Eng.*, 2017, pp. 498–510, doi: 10.1145/3106237.3106277.
4. R. Angell et al., “Themis: Automatically testing software for discrimination,” *Proc. ACM Joint Meeting Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, pp. 871–875, 2018, doi: 10.1145/3236024.3264590.
5. J. Chakraborty, S. Majumder, and T. Menzies, “Bias in machine learning software: Why? How? What to do?” in *Proc. 29th ACM Joint Meeting Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, pp. 429–440, 2021, doi: 10.1145/3468264.3468537.
6. B. Johnson and J. Smith, “Towards ethical data-driven software: Filling the gaps in ethics research & practice,” in *Proc. IEEE/ACM Int. Workshop Ethics Softw. Eng. Res. Pract. (SEthics)*, pp. 18–25, 2021, doi: 10.1109/SEthics52569.2021.00011.
7. X. Ling, T. Menzies, C. Hazard, J. Shu, and J. Beel “Trading off scalability, privacy, and performance in data synthesis,” *IEEE Access*, vol. 12, pp. 26,642–26,654, 2024, doi: 10.1109/ACCESS.2024.3366556.
8. T. Menzies and C. Hazard, “The best data are fake data?: An interview with Chris Hazard,” *IEEE Softw.*, vol. 40, no. 5, pp. 121–124, Sep./Oct. 2023, doi: 10.1109/MS.2023.3286480.
9. T. Menzies and B. Johnson, “The ethical engineer’s dilemma: An interview with Jim Herbsleb,” *IEEE Softw.*, vol. 42, no. 2, pp. 103–106, Mar./Apr. 2025, doi: 10.1109/MS.2024.3520330.



BRITTANY JOHNSON is an assistant professor of computer science at George Mason University, Fairfax, VA 22030 USA. Contact her at johnsonb@gmu.edu



TIM MENZIES is a full professor at North Carolina State University, Raleigh, NC 27606 USA. Contact him at timm@ieee.org.

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

Get more citations by publishing with the *IEEE Open Journal of the Computer Society*

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



Submit your paper today!

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



Career Accelerating Opportunities

Explore new options—upload your resume today

careers.computer.org



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

Take advantage of these special resources for job seekers:



JOB ALERTS



TEMPLATES



WEBINARS



CAREER
ADVICE



RESUMES VIEWED
BY TOP EMPLOYERS

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





stay connected.

Join our online community! Follow us to stay connected wherever you are:



| @ComputerSociety



| facebook.com/IEEEComputerSociety



| IEEE Computer Society



| youtube.com/IEEEComputerSociety



| instagram.com/ieee_computer_society

IEEE SECURITY & PRIVACY

IEEE Security & Privacy is a bimonthly magazine communicating advances in security, privacy, and dependability from the top thinkers in the field.

computer.org/security

Find the latest research and practical articles alongside case studies, surveys, tutorials, columns, and in-depth interviews. Topics include:

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



**Join the IEEE Computer Society for
subscription discounts today!**
computer.org/product/csdl-full-access



**SUBMIT
TODAY**

IEEE TRANSACTIONS ON

SUSTAINABLE COMPUTING

► SCOPE

The *IEEE Transactions on Sustainable Computing (T-SUSC)* is a peer-reviewed journal devoted to publishing high-quality papers that explore the different aspects of sustainable computing. The notion of sustainability is one of the core areas in computing today and can cover a wide range of problem domains and technologies ranging from software to hardware designs to application domains. Sustainability (e.g., energy efficiency, natural resources preservation, using multiple energy sources) is needed in computing devices and infrastructure and has grown to be a major limitation to usability and performance.

Contributions to *T-SUSC* must address sustainability problems in different computing and information processing environments and technologies, and at different levels of the computational process. These problems can be related to information processing, integration, utilization, aggregation, and generation. Solutions for these problems can call upon a wide range of algorithmic and computational frameworks, such as optimization, machine learning, dynamical systems, prediction and control, decision support systems, meta-heuristics, and game-theory to name a few.

T-SUSC covers pure research and applications within novel scope related to sustainable computing, such as computational devices, storage organization, data transfer, software and information processing, and efficient algorithmic information distribution/processing. Articles dealing with hardware/software implementations, new architectures, modeling and simulation, mathematical models and designs that target sustainable computing problems are encouraged.

SUBSCRIBE AND SUBMIT

For more information on paper submission, featured articles, calls for papers, and subscription links visit:

www.computer.org/tsusc



IEEE

COMPUTER ARCHITECTURE

LETTERS

IEEE Computer Architecture Letters is a forum for fast publication of new, high-quality ideas in the form of short, critically refereed technical papers. Submissions are accepted on a continuing basis and letters will be published shortly after acceptance in IEEE Xplore and in the Computer Society Digital Library.

Submissions are welcomed on any topic in computer architecture, especially:

- Microprocessor and multiprocessor systems
- Microarchitecture and ILP processors
- Workload characterization
- Performance evaluation and simulation techniques
- Interactions with compilers and operating systems
- Interconnection network architectures
- Memory and cache systems
- Power and thermal issues at the architectural level
- I/O architectures and techniques
- Independent validation of previously published results
- Analysis of unsuccessful techniques
- Domain-specific processor architecture (embedded, graphics, network)
- High-availability architectures
- Reconfigurable computer architectures

www.computer.org/cal



Join the IEEE Computer Society
for subscription discounts today!

www.computer.org/product/journals/cal



IEEE
COMPUTER
SOCIETY



IEEE Internet Computing

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

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

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

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

www.computer.org/internet



Join the IEEE Computer Society
for subscription discounts today!

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



IEEE
COMPUTER
SOCIETY



IEEE TRANSACTIONS ON

COMPUTERS

Call for Papers

Publish your work in the IEEE Computer Society's flagship journal, *IEEE Transactions on Computers (TC)*. *TC* is a monthly publication with a wide distribution to researchers, industry professionals, and educators in the computing field.

TC seeks original research contributions on areas of current computing interest, including the following topics:

- Computer architecture
- Software systems
- Mobile and embedded systems
- Security and reliability
- Machine learning
- Quantum computing

All accepted manuscripts are automatically considered for the monthly featured paper and annual Best Paper Award.



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



IEEE
COMPUTER
SOCIETY



IEEE TRANSACTIONS ON BIG DATA

IEEE Transactions on Big Data is a quarterly journal that publishes peer-reviewed articles with big data as the main focus.

The articles provide cross-disciplinary, innovative research ideas and applications results for big data including novel theory, algorithms, and applications. Research areas include:

- Big data
 - Analytics
 - Curation and management
 - Infrastructure
 - Performance analyses
 - Semantics
 - Standards
 - Visualization
- Intelligence and scientific discovery from big data
- Security, privacy, and legal issues specific to big data

Applications of big data in the fields of endeavor where massive data is generated are of particular interest.

www.computer.org/tbd



Join the IEEE Computer Society
for subscription discounts today!

www.computer.org/product/journals/tbd



computer.org/pervasive

IEEE pervasive COMPUTING

Call for Articles

IEEE Pervasive Computing publishes accessible, useful peer-reviewed papers on the latest 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
bit.ly/3UoGDT7

Computing in Science & Engineering

The computational and data-centric problems faced by scientists and engineers transcend disciplines. There is a need to share knowledge of algorithms, software, and architectures, and to transmit lessons-learned to a broad scientific audience. *Computing in Science & Engineering (CiSE)* is a cross-disciplinary, international publication that meets this need by presenting contributions of high interest and educational value from a variety of fields, including physics, biology, chemistry, and astronomy. *CiSE* emphasizes innovative applications in cutting-edge techniques. *CiSE* publishes peer-reviewed research articles, as well as departments spanning news and analyses, topical reviews, tutorials, case studies, and more.

Read *CiSE* today! www.computer.org/cise




IEEE
COMPUTER
SOCIETY



IEEE Annals

of the History of Computing



IEEE Annals of the History of Computing publishes work covering the broad history of computer technology, including technical, economic, political, social, cultural, institutional, and material aspects of computing. Featuring scholarly articles by historians, computer scientists, and interdisciplinary scholars in fields such as media studies and science and technology studies, as well as firsthand accounts, *Annals* is the primary scholarly publication for recording, analyzing, and debating the history of computing.



www.computer.org/annals



IT Professional

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
- Health Information Technology

We welcome articles accompanied by web-based demos.

For more information, see our author guidelines at
bit.ly/4faGdch

computer.org/itpro



Draw Ahead of the Crowd—Publish in *IEEE Computer Graphics & Applications*



IEEE Computer Graphics and Applications (CG&A) Seeks Original Submissions

- Modeling
- HCI/User Interfaces
- Rendering
- VR, AR, XR, and MR systems
- Animation
- And More
- Haptics
- Data Visualization

IEEE
Computer Graphics
AND APPLICATIONS

In addition to peer-reviewed research papers, *IEEE CG&A* also publishes department articles, which can be shorter research papers, tutorials, surveys, position statements, or viewpoint articles.

For more information, check our author information page.

Calls for Papers



bit.ly/CGA_cfp





Conference Calendar

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

OCTOBER

6 October

- MASS (IEEE Int'l Conf. on Mobile Ad-Hoc and Smart Systems), Chicago, USA
- VL/HCC (IEEE Symposium on Visual Languages and Human-Centric Computing), Raleigh, North Carolina, USA

7 October

- SecDev (IEEE Secure Development Conference), Indianapolis, USA

8 October

- ISMAR (IEEE Int'l Symposium on Mixed and Augmented Reality), Daejeon, South Korea

10 October

- BDCloud (IEEE Int'l Conf. on Big Data and Cloud Computing), Shenyang, China
- ISPA (IEEE Int'l Symposium on Parallel and Distributed Processing with Applications), Shenyang, China
- SocialCom (IEEE Int'l Conf. on Social Computing and Networking), Shenyang, China
- SpaCCS (IEEE Int'l Conf. on Security, Privacy, Anonymity in Computation and Communication and Storage), Shenyang, China
- SustainCom (IEEE Int'l Conf. on Sustainable Computing and

Communications), Shenyang, China

12 October

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

14 October

- LCN (IEEE Conf. on Local Computer Networks), Sydney, Australia

21 October

- CBDDCom (IEEE Conf. on Cloud and Big Data Computing), Hakodate, Japan
- CyberSciTech (IEEE Cyber Science and Technology Congress), Hakodate, Japan
- DASC (IEEE Conf. on Dependable, Autonomic and Secure Computing), Hakodate, Japan
- DFT (IEEE Int'l Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems), Barcelona, Spain
- ISSRE (IEEE Int'l Symposium on Software Reliability Eng.), São Paulo, Brazil
- PICom (IEEE Conf. on Pervasive and Intelligent Computing), Hakodate, Japan

28 October

- SBAC-PAD (IEEE Int'l Symposium on Computer Architecture and High Performance

Computing), Bonito, Mato Grosso do Sul, Brazil

30 October

- Blockchain (IEEE Int'l Conf. on Blockchain), Zhengzhou, China
- CPSCoM (IEEE Int'l Conf. on Cyber, Physical and Social Computing), Zhengzhou, China
- Cybermatics (IEEE Cybermatics Congress), Zhengzhou, China
- GreenCom (IEEE Int'l Conf. on Green Computing and Communications), Zhengzhou, China
- iThings (IEEE Int'l Conf. on Internet of Things), Zhengzhou, China
- SmartData (IEEE Int'l Conf. on Smart Data), Zhengzhou, China

NOVEMBER

1 November

- VIS (IEEE Visualization and Visual Analytics), Vienna, Austria

2 November

- FIE (IEEE Frontiers in Education Conf.), Nashville, USA
- LDAV (IEEE Symposium on Large Data Analysis and Visualization), Vienna, Austria
- QAI (IEEE Int'l Conf. on Quantum Artificial Intelligence), Naples, Italy



3 November

- ICTAI (IEEE Int'l Conf. on Tools with Artificial Intelligence), Athens, Greece
- PRDC (IEEE Pacific Rim Int'l Symposium on Dependable Computing), Seoul, Korea

6 November

- BIBI (IEEE Int'l Conf. on Bioinformatics and Bioengineering), Athens, Greece

7 November

- CSCloud (IEEE Int'l Conf. on Cyber Security and Cloud Computing), NYC, USA
- EdgeCom (IEEE Int'l Conf. on Edge Computing and Scalable Cloud), NYC, USA

10 November

- CASCON (IEEE Int'l Conf. on Collaborative Advances in Software and COmputiNg), Toronto, Canada
- ICCD (IEEE Int'l Conf. on Computer Design), Richardson, Texas, USA
- ICEBE (IEEE Int'l Conf. on E-Business Eng.), Buraydah, Saudi Arabia

11 November

- CIC (IEEE Int'l Conf. on Collaboration and Internet Computing), Pittsburgh, USA
- CogMI (IEEE Int'l Conf. on Cognitive Machine Intelligence), Pittsburgh, USA
- TPS-ISA (IEEE Int'l Conf. on Trust, Privacy and Security in Intelligent Systems, and Applications), Pittsburgh, USA

12 November

- ICDM (IEEE Int'l Conf. on Data Mining), Washington DC, USA

13 November

- ICKG (IEEE Int'l Conf. on Knowledge Graph), Limassol, Cyprus

14 November

- AI + Congress (IEEE AI + Congress), Guiyang, China

17 November

- SmartIoT (IEEE Int'l Conference on Smart Internet of Things), Sydney, Australia

21 November

- IPCCC (IEEE Int'l Performance, Computing, and Communications Conf.), Austin, USA

DECEMBER

2 December

- RTSS (IEEE Real-Time Systems Symposium), Boston, USA

5 December

- ICA (IEEE Int'l Conf. on Agentic AI), Wuhan, China

8 December

- BigData (IEEE Int'l Conf. on Big Data), Macau, China

14 December

- FOCS (IEEE Annual Symposium on Foundations of Computer Science), Sydney, Australia
- ICPADS (IEEE Int'l Conf. on Parallel and Distributed Systems), Hefei, China

15 December

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

- iSES (IEEE Int'l Symposium on Smart Electronic Systems), Jaipur, India

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

18 December

- ESAI (Int'l Conf. on Embedded Systems and Artificial Intelligence), Fez, Morocco

2026

JANUARY

14 January

- ICOIN (Int'l Conf. on Information Networking), Hanoi, Vietnam



Learn more
about IEEE
Computer
Society
conferences

computer.org/conferences



IPDPS
2026 • New Orleans, USA

New Orleans, USA • May 25-29, 2026

40th IEEE International Parallel and Distributed Processing Symposium

ipdps.org



IPDPS 2026 CALL FOR PAPERS

IPDPS is an international forum for engineers and scientists from around the world to present their latest research findings in all aspects of parallel computation and distributed processing. In 2026, it will be held in New Orleans at the Marriott on Canal Street.

- **October 2, 2025 – Abstracts due**
- **October 9, 2025 – Papers due**

Authors are invited to submit manuscripts that present novel and impactful research in high performance computing (HPC) in parallel and distributed processing. In an effort to produce a standardized, long-lasting impact, **IPDPS 2026 is introducing a computational result reproducibility appendix**, to be appended to accepted papers, and aimed at describing the processes used to obtain results. Works focusing on emerging technologies, interdisciplinary work spanning multiple IPDPS focus areas, and novel open-source artifacts are welcome. Topics of interest include but are not limited to the following areas:

- **Algorithms**
- **Architecture**
- **Applications**
- **Machine Learning and Artificial Intelligence (ML/AI)**
- **Measurements, Modeling, and Experiments**
- **Programming Models, Compilers, and Runtime Systems**
- **System Software**

To see the full 2026 Call for Papers, visit ipdps.org.

For **IPDPS 2026**, we are holding all workshops during the first two days, and the main conference program will follow on the last 3 days. In addition to technical sessions of submitted paper presentations and invited speakers, IPDPS offers a student research forum, tutorials, and commercial presentations. Make plans to join us in New Orleans!

GENERAL CO-CHAIRS

Anu Bourgeois, Georgia State University, USA
Guillaume Pallez, INRIA, France

IPDPS 2026 PROGRAM CO-CHAIRS

Maxim Naumov, Meta, USA
Cristina Silvano, Politecnico di Milano, Italy

2026 PROGRAM TRACK CO-CHAIRS

- **Algorithms**
Sivan Toledo, University of Tel-Aviv, Israel
Oded Green, Nvidia, USA
- **Applications**
Lin Gan, Tsinghua University, China
Axel Huebl, Lawrence Berkeley National Lab, USA
- **Architectures**
Catherine Schuman, University of Tennessee, USA
Ioannis Sourdis, Chalmers University of Tech, Sweden
- **Machine Learning and Artificial Intelligence**
Jongsoo Park, OpenAI, USA
Alexey Tumanov, Georgia Tech, USA
- **Measurement, Performance and Experiments**
Marc Casas, Barcelona Supercomputing Center, Spain
Shirley Moore, University of Texas - El Paso, USA
- **Programming Models, Compilers, & Runtime Systems**
Dimitros Nikolopoulos, Virginia Tech, USA
Valeria Cardellini, University of Roma Tor Vergata, Italy
- **System Software**
Bogdan Nicolae, Argonne National Lab (ANL), USA
Ivy Peng, KTH Royal Institute Tech, Sweden

WORKSHOPS CHAIR

Suren Byna, The Ohio State University, USA

WORKSHOPS VICE CO-CHAIRS

Francieli Boito, Université de Bordeaux, France
Giulia Guidi, Cornell University, USA

To see the 2026 Call for Workshops, visit ipdps.org.

SPONSORED BY



IN COOPERATION WITH

