COMPUTING edge

> **Blockchain**
> **Internet of Things**
> **Internet Trends and Applications**
> **Virtual and Augmented Reality**

JANUARY 2019

www.computer.org

IEEE

IEEE COMPUTER SOCIETY

COMPSAC 2019

DATA DRIVEN INTELLIGENCE FOR A SMARTER WORLD

Hosted by Marquette University, Milwaukee, Wisconsin, USA
July 15-19 *www.compsac.org*

In the era of "big data" there is an unprecedented increase in the amount of data collected in data warehouses. Extracting meaning and knowledge from these data is crucial for governments and businesses to support their strategic and tactical decision making. Furthermore, artificial intelligence (AI) and machine learning (ML) makes it possible for machines, processing large amounts of such data, to learn and execute tasks never before accomplished. Advances in big data-related technologies are increasing rapidly. For example, virtual assistants, smart cars, and smart home devices in the emerging Internet of Things world, can, we think, make our lives easier. But despite perceived benefits of these technologies/methodologies, there are many challenges ahead. What will be the social, cultural, and economic challenges arising from these developments? What are the technical issue related, for example, to the privacy and security of data used by AI/ML systems? How might humans interact with, rely on, or even trust AI predictions or decisions emanating from these technologies? How can we prevent such data-driven intelligence from being used to make malicious decisions?

Authors are invited to submit original, unpublished research work, as well as industrial practice reports. Simultaneous submission to other publication venues is not permitted. All submissions must adhere to IEEE Publishing Policies, and all will be vetted through the IEEE CrossCheck portal. **For full CFP and conference information, please visit the conference website at WWW.COMPSAC.ORG**

IMPORTANT DATES

October 1, 2018: Workshop proposals due
October 30, 2018: Workshop proposal notifications
January 21, 2019: Abstracts and full papers due
April 7, 2019: Paper notifications
April 15, 2019: Workshop papers due
May 1, 2019: Workshop paper notifications
May 17, 2019 – Camera ready submissions and advance author registration due

IEEE
IEEE computer society
MARQUETTE UNIVERSITY
Be The Difference.

ORGANIZING COMMITTEE

General Chairs: Jean-Luc Gaudiot, University of California, Irvine, USA; Vladimir Getov, University of Westminster, UK

Program Chairs in Chief: Morris Chang, University of South Florida, USA; Stelvio Cimato, University of Milan, Italy; Nariyoshi Yamai, Tokyo University of Agriculture & Technology, Japan

Workshop Chairs: Hong Va Leong, Hong Kong Polytechnic University, Hong Kong; Yuuichi Teranishi, National Institute of Information and Communications Technology, Japan; Ji-Jiang Yang, Tsinghua University, China

Local Organizing Committee Chair: Praveen Madiraju, Marquette University, USA

Standing Committee Chair: Sorel Reisman, California State University, USA

Standing Committee Vice Chair: Sheikh Iqbal Ahamed, Marquette Universiity, USA

COMPUTING edge

IEEE COMPUTER SOCIETY computer.org • +1 714 821 8380

Printed with inks containing soy and/or vegetable oils

IEEE

IEEE COMPUTER SOCIETY

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.

JANUARY 2019 • VOLUME 5, NUMBER 1

COMPUTING edge

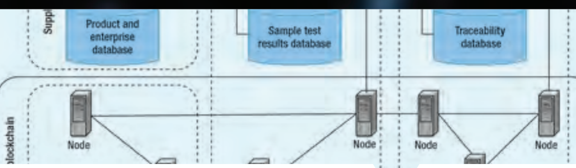

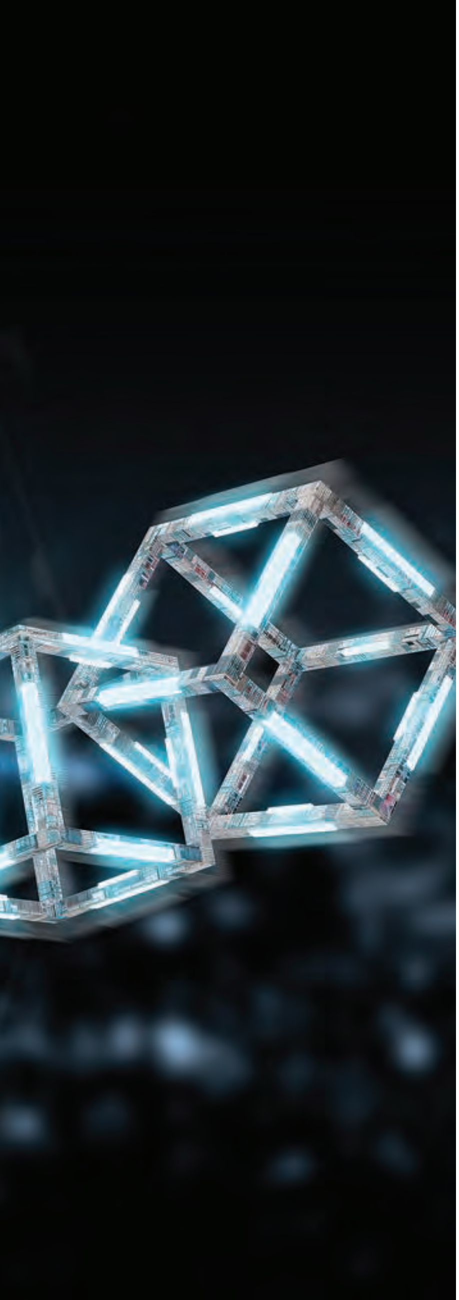

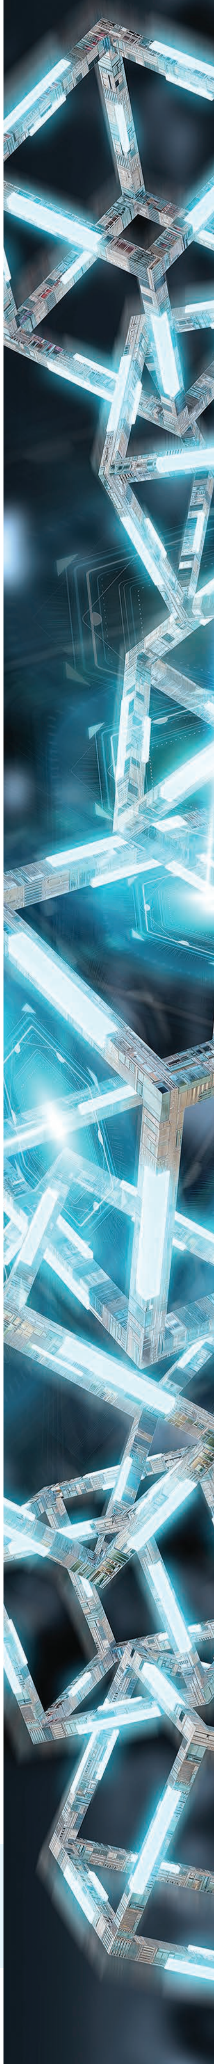

38

A New Era for Web AR with Mobile Edge Computing

Magazine Roundup

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

Computer

The Age of Artificial Emotional Intelligence

Science fiction often portrays future AI technology as having sophisticated emotional intelligence skills to the degree where technology can develop compassion. But where are we today? In this article from the September 2018 issue of *Computer*, the authors provide insight into artificial emotional intelligence (AEI) and present three major areas of emotion—recognition, generation, and augmentation—needed to reach a new emotionally intelligent epoch of AI.

Computing in Science & Engineering

The Heat Equation: High-Performance Scientific Computing Case Study

In recent years, high-performance computing and powerful supercomputers have become staples in many areas of academia and industry. The author of this article from the September/October 2018 issue of *Computing in Science & Engineering* introduces the concept of shared memory

 2469-7087/19/$33.00

programming in the context of solving the heat equation, which will allow the exploration of several finite difference and parallelization schemes.

IEEE Annals of the History of Computing

How Atex Helped an Industry Change the World

Douglas Drane was one of three co-founders of Atex, an early dominant supplier of digital systems for newspaper and magazine publication. In the July–September 2018 issue of *IEEE Annals of the History of Computing*, Drane recounts his memories of Atex's founding and operation until it was bought by Kodak.

IEEE Computer Graphics and Applications

ThunderPunch: A Bare-Hand, Gesture-Based, Large Interactive Display Interface with Upper-Body-Part Detection in a Top View

The authors of this article from the September/October 2018 issue of *IEEE Computer Graphics and Applications* present a new bare-hand gesture interface for large-screen interaction in which multiple users can participate simultaneously and interact with virtual content directly. Unlike the conventional method, which involves positioning the camera in front, the cameras are mounted on the ceiling so that they avoid covering the large screen. To achieve bare-hand interaction in this hardware structure, the authors propose real-time algorithms that detect multiple body poses and recognize punching and touching gestures from top-view depth images. A pointing and touching test shows that the proposed algorithm outperforms other algorithms. In addition, the authors created a game to make the best use of the proposed system.

IEEE Intelligent Systems

Detecting Personal Intake of Medicine from Twitter

Mining social media messages such as tweets, blogs, and Facebook posts for health- and drug-related information has received significant interest in pharmacovigilance research. Social media websites have been used for monitoring drug abuse, searching for adverse reactions to drug usage, and analyzing expression of sentiments related to drugs. Most of these studies are based on aggregated results from a large population rather than specific sets of individuals. To conduct studies at an individual level or in specific groups of people, identifying posts mentioning intake of medicine by the user is necessary. Toward this objective, the authors of this article from the July/August 2018 issue of *IEEE Intelligent Systems* developed a classifier for identifying mentions of personal intake of medicine in tweets. They trained a stacked ensemble of shallow convolutional neural network (CNN) models on an annotated dataset and used random search for tuning the hyper-parameters of the CNN models. The classifier has direct uses in the areas of psychology, health informatics, pharmacovigilance, and affective computing.

IEEE Internet Computing

Real-Time Identity-Deception Detection Techniques for Social Media: Optimizations and Challenges

Identity-deception detection methods have been proposed for social-media platforms with high effectiveness, but their efficiency can vary. Previous literature has not examined the potential of these methods to work as real-time monitoring systems. Such implementations further highlight the challenges of applying computationally intensive methods in online environments that involve large datasets and high speeds of data. This article, which appears in the September/October 2018 issue of *IEEE Internet Computing*, attempts to classify detection methods based on the approaches and identifies factors that, in real-time systems, will impact the effectiveness and efficiency of these methods. Optimizations are proposed that can limit the computational overhead. Further challenges involving real-time identity-deception detection are discussed.

IEEE Micro

Newton: Gravitating Towards the Physical Limits of Crossbar Acceleration

Recent works take advantage of highly parallel analog in-situ computation in memristor crossbars to accelerate the many vector-matrix

multiplication operations in deep neural networks (DNNs). However, these in-situ accelerators have two significant shortcomings: The ADCs account for a large fraction of chip power and area, and these accelerators adopt a homogeneous design in which every resource is provisioned for the worst case. By addressing both problems, a new architecture, called Newton, moves closer to achieving optimal energy per neuron for crossbar accelerators. In this article from the September/October 2018 issue of *IEEE Micro*, the authors introduce new techniques that apply at different levels of the tile hierarchy, some leveraging heterogeneity and others relying on divide-and-conquer numeric algorithms to reduce computations and ADC pressure. For many convolutional-neural-network (CNN) dataflows and structures, Newton achieves a 77-percent decrease in power, 51-percent improvement in energy-efficiency, and 2.1× higher throughput/area, relative to the state-of-the-art In-Situ Analog Arithmetic in Crossbars (ISAAC) accelerator.

IEEE MultiMedia

Digital Twins: The Convergence of Multimedia Technologies

Originally developed to improve manufacturing processes, digital twins are being redefined as digital replications of living and nonliving entities that enable data to be seamlessly transmitted between the physical and virtual worlds. Digital twins facilitate the means to monitor, understand, and optimize the functions of all physical entities and for humans to provide continuous feedback to improve quality of life and well-being. Read more in the April–June 2018 issue of *IEEE MultiMedia*.

IEEE Pervasive Computing

Privacy-Preserving Incentive Mechanisms for Mobile Crowdsensing

The smartphone market has proliferated rapidly in recent years. Ubiquitous smartphones not only possess powerful computational capabilities but also contain various built-in sensors. Along with providing easy access to communication networks such as Wi-Fi and 3G/4G networks, smartphones have become an important information interface between users and environments and potentially constitute a huge mobile sensor network. These advances motivate researchers to propose mobile crowdsensing (MCS) systems that can provide large-scale and complex social or geographical sensing applications. Read more in the July–September 2018 issue of *IEEE Pervasive Computing*.

IEEE Security & Privacy

Privacy and Civilian Drone Use: The Need for Further Regulation

Current US regulation is not equipped to provide explicit privacy protection for drone use in an era of sophisticated audio/video and social media. In 2016, the National Telecommunications and Information Administration recognized this deficit by releasing a set of best practices, which we examine in light of the current privacy concerns with drone use in the US. Read more in the September/October 2018 issue of *IEEE Security & Privacy*.

IEEE Software

Software Engineering Research and Industry: A Symbiotic Relationship to Foster Impact

Software engineering is not only an increasingly challenging endeavor that goes beyond the intellectual capabilities of any single individual engineer but also an intensely human one. Tools and methods to develop software are employed by engineers of varied backgrounds within a large variety of organizations and application domains. As a result, the variation in challenges and practices in system requirements, architecture, and quality assurance is staggering. Human, domain, and organizational factors define the context within which software engineering methodologies and technologies are to be applied. This article, which is part of the September/October 2018 issue of *IEEE Software* on software engineering's 50th anniversary, provides an assessment of the current challenges faced by software engineering research in achieving its potential, a description of the root causes of such challenges, and a proposal for the field to move forward and become more impactful through collaborative research and innovation between public research and industry.

IT Professional

The Evolving Payments Landscape: Technological Innovation in Payment Systems

Established financial institutions and their technology partners are contributing significantly to financial technology innovation in the payments market. This article, which appears in the March/April 2018 issue of *IT Professional*, provides a holistic overview of the technological innovations and challenges in the evolving payments landscape.

ADVERTISER INFORMATION

Advertising Personnel

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

Advertising Sales Representatives (display)

Central, Northwest, Southeast, Far East:
Eric Kincaid
Email: e.kincaid@computer.org
Phone: +1 214 673 3742
Fax: +1 888 886 8599

Northeast, Midwest, Europe, Middle East:
David Schissler
Email: d.schissler@computer.org
Phone: +1 508 394 4026
Fax: +1 508 394 1707

Southwest, California:
Mike Hughes
Email: mikehughes@computer.org
Phone: +1 805 529 6790

Advertising Sales Representative (Classifieds & Jobs Board)

Heather Buonadies
Email: h.buonadies@computer.org
Phone: +1 201 887 1703

Advertising Sales Representative (Jobs Board)

Marie Thompson
Email: marie@4caradio.org
Phone: 714-813-5094

EDITOR'S NOTE

Emerging Blockchain Applications

In the public consciousness, blockchain is almost synonymous with cryptocurrencies like Bitcoin. But many other exciting uses for this emerging technology are in the works. Blockchain utilizes cryptography to log records securely and reliably, leading to potential applications in important areas such as banking, healthcare, and voting.

The January 2019 issue of *ComputingEdge* discusses two new blockchain applications. *IEEE Internet Computing*'s "Blockchain: The Emperor's New PKI?" proposes the idea that blockchain could become the foundation of our future digital identities. Meanwhile, *IEEE Software*'s "Adaptable Blockchain-Based Systems" presents a novel decentralized system for tamper-proof product tracking in complex supply chains.

Other technologies and trends—such as the Internet of Things (IoT), mobile Internet use, and virtual and augmented reality (VR and AR)—are also expanding into new territories. In *Computer*'s "Accurate Indoor Location for the IoT," the authors explain that improved location accuracy on IoT devices is opening doors to possibilities like in-building navigation and automated smart-home operation. "The Last Mile for IoT Privacy," from *IEEE Security & Privacy*, presents a conceptual framework for notifying people when IoT devices are collecting their data.

Mobile Internet trends are pushing boundaries thanks to emerging technologies such as 5G and edge computing. The author of *IT Professional*'s "5G in E-Commerce Activities" details the impact of 5G networks, platforms, and devices on the online shopping experience. The authors of *IEEE Internet Computing*'s "A New Era for Web AR with Mobile Edge Computing" discuss how mobile edge computing on 5G networks can reduce latency, thereby improving web AR performance.

Finally, VR is seeing novel implementations, such as the 360-degree camera presented in *IEEE MultiMedia*'s "360-Degree Virtual-Reality Cameras for the Masses." VR is also playing an increasingly important role in the tourism industry, as shown in *Computing in Science & Engineering*'s "Virtual Travel."

 2469-7087/19/$33.00

IEEE computer society

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.
MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.
COMPUTER SOCIETY WEBSITE: www.computer.org
OMBUDSMAN: Direct unresolved complaints to ombudsman@computer.org.
CHAPTERS: Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.
AVAILABLE INFORMATION: To check membership status, report an address change, or obtain more information on any of the following, email Customer Service at help@computer.org or call +1 714 821 8380 (international) or our toll-free number, +1 800 272 6657 (US):

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

PUBLICATIONS AND ACTIVITIES
***Computer*:** The flagship publication of the IEEE Computer Society, *Computer*, publishes peer-reviewed technical content that covers all aspects of computer science, computer engineering, technology, and applications.
Periodicals: The society publishes 13 magazines, 19 transactions, and one letters. Refer to membership application or request information as noted above.
Conference Proceedings & Books: Conference Publishing Services publishes more than 275 titles every year.
Standards Working Groups: More than 150 groups produce IEEE standards used throughout the world.
Technical Committees: TCs provide professional interaction in more than 30 technical areas and directly influence computer engineering conferences and publications.
Conferences/Education: The society holds about 200 conferences each year and sponsors many educational activities, including computing science accreditation.
Certifications: The society offers two software developer credentials. For more information, visit www.computer.org/certification.

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

MEMBERSHIP & PUBLICATION ORDERS
Phone: +1 800 272 6657 • **Fax:** +1 714 821 4641 • Email: help@computer.org
Asia/Pacific: Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan
Phone: +81 3 3408 3118 • **Fax:** +81 3 3408 3553
Email: tokyo.ofc@computer.org

IEEE

revised 03 July 2018

Blockchain: the Emperor's New PKI?

Hilarie Orman
Purple Streak

I would like to jump on the blockchain bandwagon. I would like to be able to say that blockchain is the solution to the longstanding problem of secure identity on the Internet. I would like to say that everyone in the world will soon have a digital identity. Put yourself on the blockchain and never again ask yourself, "Who am I?" – you are your blockchain address.

> *"All the world's a blockchain,*
> *And all the men and women merely entries in it."*

Certainly, it is important to solve the online identity problem. Identity can be destiny. Today, the mere fact of a person's physical existence is not enough to guarantee some rights. Citizenship, refugee status, right to travel, government benefits, etc., are tied to a provable identity of the right type. Some have gone so far as to say that identity is a human right.[1] Is it feasible to solve this problem through digital technology? And if so, is blockchain part of the solution? And do we get privacy with that?

The digital divide leaves first-world citizens with the problem of having too many online identities to manage, while those on the other side of the chasm, some two billion people, perhaps, not only have no digital identity, but further have no documented identity at all. Everyone in the world needs at least one persistent and secure identity that is easily manageable by the individual.

Public key technology has always held the elusive promise of a universal digital identity system that makes paper obsolete, but attempts to do this have never measured up. Fraught with complex management, lookup, and trust issues, public key has been largely relegated to a way to protect Internet communication channels. Although some countries use digital identification via public key and electronic ID cards, there is no widespread movement towards this kind of solution.

Blockchain technology seems like an odd savior, especially because its main use to date has been for digital cash with anonymity features. A verified identity is nearly the opposite of anonymity, so what does blockchain have to do with it? And how can we have verified identities that preserve privacy? The technology for doing this is in its infancy, and the management problems are largely unexplored, but when has that stopped technology ebullience? Let's look at what identity means and how life might be if lived in the blockchain bubble.

WHITHER COMEST (OR GOEST) IDENTITY?

Identity is a psychological, philosophical, and legal concept that is carried over into the digital world in a variety of ways. For an interesting discussion of desirable principles in the digital realm, see the work of Kim Cameron.[2]

The simple story is that the various authorities that control the resources that we need to participate in life today need to know who we are, which means they want enough information to disambiguate all people from one another. Somehow, there must be a short index from a person into a table of data about many people, potentially all people on earth. Some authorities want to attach attributes to identities, other authorities need to examine those attributes. At the center of it all is a unique human being. In order to avoid some kinds of dystopian nightmares about authoritarians dictatorships, it seems important to let the person, the identity subject, control how and when the attributes are disclosed. For example, people need to communication their immunization status to a health organization, residence to local housing authorities, financial status to lending organizations, etc. A recent trend is to use cryptographic zero-knowledge proofs for this purpose. But this is not new. What is new is using blockchains for bootstrapping the trust process and for the persistence and availability of the information.

We all know that our identity is more than our name and appearance, and that name recognition and personal relationships are not enough to maintain one's status in an impersonal digital world. People need to document their place of origin and residence, their health status, and their economic ties if they are going to participate in the global economy and make use of the rights to which they are entitled.

Today, identity is established in a variety of ways for people and for governments and businesses and almost all organizations. On social media, identity means a username, a communication history, and interactions with some kind of online community. For email, it is simply an email address. For a website it is a URL and perhaps a certificate. For most website interactions, it is a username and a password, perhaps augmented by a cell phone number. For newer mobile devices, various kinds of biometrics identify the device owner. The essential government issued identity documents are still based on paper: driver's licenses, passports, birth certificates, etc. (I told my doctor's office to address me by my date of birth, as it seems to be essential for them to verify it multiple times in any conversation.)

We may be on the verge of a major change, one in which public keys are of central importance, users have total control over disclosure of their identity attributes, and blockchains are the root points of trust for obtaining credentials. Or, we may be on the verge of discovering that blockchains recapitulate the well-known problems of public key infrastructure.

WHAT IS A BLOCKCHAIN?

Blockchains offer the promise of a trustworthy way to record shared data. It is a publicly verifiable ledger that maintains the integrity of the individual entries.

The basic idea of a blockchain is quite simple: it is a shared, replicated log file (sometimes called a ledger). The entries are sequential and time-stamped. A one-way function produces a short bitstring (for example, 512 bits) and depends on every item and its placement in the log. The function has mathematical properties that assure that it would be astronomically difficult to produce a different log with the same output. The output of the function is an abbreviation for the log itself. To add new entries, the function uses its current value and the contents of every new entry to compute a new output. The log maintainer publishes the log and the output value so that independent parties can verify the correspondence.

In itself, a blockchain offers no particular security. To get to an interesting trust model, two additional mechanisms are required. The first is to have a verifiable definition of correctness for each log file entry. This means that there is a correctness condition that is orthogonal to the correctness of the mere sequence of the entries; usually the log file entries are signed by public keys and have a standardized format. Secondly, the Important Idea that makes blockchains useful is to distribute the maintenance of the blockchain to many (possibly distrusting) parties, thereby achieving "distributed consensus" in real time. As long as the majority of the parties have a mutual interest in publishing a consistent view of the blockchain, it is a reliable record of transactions.

BitCoin has the interesting property that any party doing a sufficient amount of work can extend the blockchain. Anyone can become part of the peer-to-peer network and get the current state of the blockchain and list of transactions that are queued for inclusion. After checking that the transactions satisfy the correctness conditions, and then computing the validation function, a party can advertise the transactions and validation output to other nodes. All the other nodes will check the computation and begin publishing the blockchain with the new block of transactions. The salient property of the BitCoin validation function is that it is very difficult to compute but very easy to check. To incentivize validators (miners), the system automatically rewards them when their computation is accepted.

WHAT IS A BLOCKCHAIN IDENTITY?

There are proposals to create blockchain systems and communities that use blockchains in which the entries are somehow related to identification for an individual person or unique entity, such as a DNS name. In this model, the public key for a person or entity is represented in a blockchain entry, where it could be accessed by a relying party. In the simplest case, an entry would be signed by an identity provider of some kind, using its own public key, to assert: "The public key represented in this data item belongs to a person known as John Q. Public." This is similar to existing public key systems, such as X.509v3 and PGP.

Blockchains offer a couple of improvements, though. One is that the blockchain replication by a diverse set of entities ensure accessibility and reliability. Another is the immutability of the blockchain and the timestamps; there can be little argument that an entry was made, and that can be shown for as long as the blockchain exists.

In order to know if blockchains solve other fundamental identity problems, we need to look more deeply into why identity has become so important. Even though we still use passwords and special purpose apps, digital identities based on public keys are the only hope of getting out of the morass of relying on companies and governments to keep their customer or citizen data private. Their inability to do so is demonstrated on a regular basis.

There are two kinds of identity being discussed today. One is for most of the people in countries with good governmental record keeping, reliable power for computing devices, and nearly universal Internet service; the other is for people who have been cut off from government recognition, isolated by virtue of abuse or persecution, who are poor, or who do not have physical safety for possession. This can happen either because the their government is too weak to provide identity services, because they are ignored by their government, because they are victims of human trafficking, because they are refugees, etc.

The digitally dispossessed cannot rely on cellphones, cards with chips, or anything physical to provide their identity. What then identifies them? One way would be to record biometric data such as iris scans, facial features, fingerprints, or DNA. A cryptographic hash of the data could be an anchor point for attaching attributes such as country of residence, immunization status, etc. The blockchain would hold the cryptographic hash, and various authorities could certify possession of attributes. The feasibility of this is uncertain because biometric attributes are "fuzzy," while a cryptographic hash is anything but. Still, one could imagine a solution based on varying degrees of accuracy and multiple measurements to create attribute collections that could map to a unique match.

In an ideal world, none of the identity attributes could be released without permission from the person being identified. But if that person is digitally dispossessed, then she has no way to present a public key to unlock the record; instead, the biometric measurements alone are enough to unlock the information. This may well be the future identification situation not only for the digitally dispossessed, but perhaps for most of the world's population, if their governments insist on it.

There are many in the tech world today who are banking on the assumption that blockchain technology will be the way register identity and to redesign the entire concept of digital identity. The identity owner could get the unfettered freedom to manage access to his personal information, to his credentials, and to his organizational records. A whole economy of identity might arise, a libertarian utopia where the person and his keys are self-determining entities.

WHO IS DOING IT?

A working example, albeit minimal, of how a blockchain can enable a user to control his personal data can be seen in the MIT Media Lab's credentials experiment.[3] This is a blockchain-based system for issuing certification of academic achievement (it is confusing to security-minded folks that they use the term "certificate" to mean "credential"). There is a central authority (The Media Lab) that manages a blockchain. This system encodes information about a student's completion of a course and the student's public key. The signed hash of the data structure is put on the blockchain, and the student receives a copy. The student can prove that he has the certificate by sending his copy of the certificate and the blockchain address to, for example, a potential employer. The employer can see that the hash of the certificate matches the blockchain entry and that it is signed by MIT. This system gives control over the information to the learner, and MIT does not need to communicate with the potential employer. Indeed, taken to the extreme, MIT would not even need to retain any of the information; it belongs to the learner, and with great autonomy comes great responsibility. The learner must not lose his public key pair, and he must not lose track of his signed course completion document.

An additional advantage is that the Media Lab does not need to grant any special access provisions to the student. Not only does the granting institution not have to retain the record, they do not have to protect any sensitive login information for the student.

Those familiar with the trials and travails of managed public key infrastructure (PKI) may note that there are similarities between this approach and traditional ones, and there are some striking shortcuts. In this approach, the learner does not present a public key to the granting institution. Although the experimenters intended to add that in the future, the reality of the situation is that only a very sophisticated user can generate a public key pair, so it is up to the granting institution to generate a key pair for the learner and to communicate it to him securely. Of course, in the future, that will all be taken care of "in the app." Neither does the learner, or anyone else, have a certificate identifying the granting institution, nor any way to determine the legitimacy of the site with the blockchain. When the learner presents the document and its associated blockchain address to a potential employer, the employer can verify that the hash of the document is in the blockchain and signed by the granting institution's key.

While this seems straightforward, one cannot but wonder what advantage it offers over using PKI certificates for the granting institution and the learner. The granting institution can give the signed copy of the course completion document to the student, and then any third party can receive the document from the student and validate the signature.

The Media Lab system does include a provision for that bugaboo of public key systems, the revocation of keys. In this case, the revocation of the course completion document can be achieved by adding an entry to that effect onto the blockchain.

Finally, the Media Lab felt that Merkle trees offered advantages over blockchains.[4] They envision a hybrid system in which users store their identity credentials on a blockchain, but the granting institutions use Merkle trees for storing their endorsed documents for users. Though minimal, this kind of system might be a solution for storing biometric data for the digitally dispossessed. There is still a great challenge to be met in providing user consent, but there is an imminent need to help the world's vulnerable people, and circumstances may overtake technology in this sector.

There are much more ambitious projects that promise users detailed control over their identifying information and credentials. IBM and SecureKey are working with Canadian partners for an identity system based on the Linux Hyperledger software core (one of many open source blockchain projects). Microsoft got on their bandwagon recently with an announcement from their identity division.[5] Another open source blockchain project, Sovrin, is the foundation for "self-sovreign" identity technology being developed by Evernym (https://www.evernym.com).

These systems appear to promote the concept of releasing all or part of the attributes in a personal credential. The canonical example is a driver's license. People use their licenses for a variety of purposes that have nothing to do with driving: age (to get a drink), residence (to vote, to use a public library), date of birth (for health care, sometimes), appearance (to prevent identity fraud). Many people would prefer to release only the minimal necessary information to the requestor. In the last decade or so, cryptographic techniques for achieving this kind of controlled disclosure have been developed in whole or in part: secret handshakes, hidden credentials, oblivious attributes, and the general technique of homomorphic encryption. The developers of these new identity systems would like to build in the privacy protections afforded by controlled release of information.

While this seems laudable, there are many questions to be answered, not the least of which is an overall analysis of security and usability. Some of the advanced cryptographic techniques are impossibly slow, for example. Another risk is that there may be a hidden reliance on a trusted third-party underlying the cryptography. If that third-party is shown to be untrustworthy or corrupted, then the privacy of the system might be undermined in ways that are difficult to detect. The problem of generating secure public keys remains a crucial issue, one that seems to trip up hardware vendors again and again.[6]

More practical issues center on whether or not the user should be able to hide essential information like the expiration date of a driver's license or special notations about wearing glasses. And, how can the user give consent to organ donation in the event of a fatality? Working out each contingency, even for ordinary credentials, may take us some years.

From the point of view of a credential issuing or attribute using organization, there are advantages to being able to reduce the attack surface for privacy, but there are uses for collections of personal data that are at odds with privacy. Organizations need to do analyses of their user base in order to tailor their services. How many users are under the age of 25? How many live in cold climates? Will companies really be motivated to take privacy more seriously in a blockchain world?

Another concern is the complications for the user of managing a large number of secure blockchain "blobs." Many new systems tout the ability of the user to create different "identities" (keys) for different uses, but that gets complicated. Moreover, it remains to be seen if credential issuers will go along with such a model. For example, if someone uses a driver's license to get a special TSA boarding privilege, will TSA issue the credential to an arbitrary key provided by the user, or will they issue it to the same key that is bound to the driver's license?

Realistically, in the fullness of time, all of the PKI management problems will devolve to the blockchain. Companies acquiring other companies will want to re-sign all the credentials issued by subsidiaries, teenagers will want to have credentials that their parents controlled reissued to their "grown up" key, etc. It seems that these issues, and similar ones, might lead to the same kind of inconvenient X.509 certificate chains that bog down PKI; will we have blockchain certificate chains that replicate the same problems?

Public keys are essential for safe identities, but everything hinges on generating "good" keys, ones that do not have hidden pitfalls. That, in turn, depends on the quality of the software and the systems that generate them.

A further risk is that there will be no reason for institutions to stay in the game without substantive advantages. If members drop out of a blockchain, then the guarantees of availability and non-malleability become weakened. The remaining members may decide to abandon the previous blockchain, either by not maintaining it at all and moving to a new blockchain, or by forking the previous blockchain. At some point, the Wayback Machine may be littered with the charred remains of blockchains that fell out of favor.

Nonetheless, almost anyone in the identity game today probably is considering a move to blockchain for the convenience that it could bring in widespread availability and risk reduction.

THE FUTURE

An underlying current in all the blockchain identity fervor is the desire to start over with digital identity. A fresh beginning, unhampered by stodgy standards and the shackles of complex software, a new identity ecosystem is beginning to be born. I think that vision, vague as it is, might be the future of identity for the world. We could have the forced Orwellian government-issued privacy-depriving biometric credentials coexisting with subversive underground identities. free flowing reputation systems, and multiple personality disorder do-it-yourself identities for the people on the fringes.

It is clear that eventually, every human will have some representation as digital attributes. This is good and bad, like civilization itself. As humans, we seem to enjoy building bigger and bigger societies, even though that has a huge effect on how we view ourselves as individuals. It is destiny.

Based on the nonproliferation of PKI, my prediction is that when we do finally reach the point of having digital identities instead of paper documents and login passwords, then they will be built on some kind of distributed ledger, something resembling a blockchain. I only hope that it improves security, rather than making it ever more obscure and tenuous.

REFERENCES

1. *ID2020, identity for the world*, ID 2020 Initiatve, 2017; https://id2020.org.
2. K. Cameron, "The Laws of Identity," Microsoft Corporation, 2005; https://msdn.microsoft.com/en-us/library/ms996456.aspx.
3. "What we learned from designing an academic certificates system on the blockchain," MIT Media Lab, 2016; https://medium.com/mit-media-lab/what-we-learned-from-designing-an-academic-certificates-system-on-the-blockchain-34ba5874f196.
4. R.C. Merkle, "A Digital Signature Based on a Conventional Encryption Function," *Advances in Cryptology — CRYPTO '87*, Springer, 1987.
5. A. Simon, "The Future as We See It," blog, 2018; https://cloudblogs.microsoft.com/enterprisemobility/2018/02/12/decentralized-digital-identities-and-blockchain-the-future-as-we-see-it.
6. N. Heninger et al., "Mining Your Ps and Qs: Detection of Widespread Weak Keys in Network Devices," *Proc. 21st USENIX Security Symposium*, 2012; https://factorable.net/weakkeys12.extended.pdf.

ABOUT THE AUTHOR

Hilarie Orman is a security consultant and president of Purple Streak. Her research interests include applied cryptography, secure operating systems, malware identification, security through semantic computing, and personal data mining. Orman has a BS in mathematics from the Massachusetts Institute of Technology. She's a former chair of the IEEE Computer Society's Technical Committee on Security and Privacy. Contact her at hilarie@purplestreak.com.

This article originally appeared in
IEEE Internet Computing, *vol. 22, no. 2, 2018.*

CALL FOR PAPERS

2019 IEEE World Congress on Services

BigData/CLOUD/EDGE/ICCC/ICIOT/ICWS/SCC

July 8-13, 2019 Milan, Italy

http://conferences.computer.org/services/2019/

2019 IEEE World Congress on Services (SERVICES) will be held on July 8-13, 2019 in Milan, Italy. The Congress is solely sponsored by the IEEE Computer Society under the auspice of the Technical Committee on Services Computing (TCSVC). The scope of the Congress will cover all aspects of services computing and applications, current or emerging. It covers various systems and networking research pertaining to cloud, edge and Internet-of-Things (IoT), as well as technologies for intelligent computing, learning, big data and blockchain applications, while addressing critical issues such as high performance, security, privacy, dependability, trustworthiness, and cost-effectiveness. Authors are invited to prepare early and submit original papers to any of these conferences at *www.easychair.org*. All submitted manuscripts will be peer-reviewed by at least 3 reviewers. Accepted and presented papers will appear in the conference proceedings published by the IEEE Computer Society Press. The Congress will be organized with the following seven affiliated conferences/congresses:

IEEE International Congress on Big Data (BigData Congress)
Big data acquisitions, analyses, storage, and mining for various services and applications

IEEE International Conference on Cloud Computing (CLOUD)
Innovative cloud computing for both high quality infrastructure and mobile services

IEEE International Conference on Edge Computing (EDGE)
High quality services computing between cloud systems and Iot devices

IEEE International Conference on Web Services (ICWS)
Innovative web services for various effective applications

IEEE International Conference on Cognitive Computing (ICCC)
Cognitive computing, learning algorithms for intelligent services and applications

IEEE International Congress on Internet of Things (ICIOT)
Innovative IoT technology for digital world services

IEEE International Conference on Services Computing (SCC)
Intelligent services computing, lifecycles, infrastructure and mobile environments

Key Dates

Early Paper submission due: December 1, 2018
Review comments to authors of early-submission papers: January 15, 2019
Normal Paper submission due: February 4, 2019
Final notification to authors: March 15, 2019.
Camera-ready manuscripts due: April 1, 2019
Congress Date: July 8 – 13, 2019

2019 Congress General Chair
Peter Chen, Carnegie Mellon University, USA

2019 Congress Program Chair-in-Chief
Elisa Bertino, Purdue University, USA

2019 Congress Vice Program Chair-in-Chief
Ernesto Damiani, University of Milan, Italy

Workshop Chairs
Shangguang Wang, BUPT
Stephan Reiff-Marganiec, University of Leicester

Steering Committee
Elisa Bertino, Purdue University, USA
Carl K. Chang, Chair, Iowa State University, USA,
Rong N. Chang, TCSVC Chair, IBM, USA
Peter Chen, Carnegie Mellon University USA
Ernesto Damiani, University of Milan, Italy
Ian Foster, University of Chicago & Argonne National Lab, USA
Dennis Gannon, Indiana University, USA
Michael Goul, Arizona State University, USA
Frank Leymann, University of Stuttgart, Germany
Hong Mei, Beijing Institute of Technology, China
Stephen S. Yau, Arizona State University, USA

CALL FOR WORKSHOP PROPOSALS: See *http://conferences.computer.org/services/2019/* for more information.
Send inquiries to: ieeecs.services@gmail.com

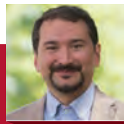

Editor: **Cesare Pautasso**
University of Lugano
c.pautasso@ieee.org

Editor: **Olaf Zimmermann**
University of Applied Sciences
of Eastern Switzerland, Rapperswil
ozimmerm@hsr.ch

Adaptable Blockchain-Based Systems

A Case Study for Product Traceability

Qinghua Lu and Xiwei Xu

From the Editors
Tracing the origin of products across complex supply chains requires a transparent, tamper-proof metadata infrastructure that's not only trusted by all the involved parties but also adaptable to changing environments and regulations. Can such advanced infrastructure be implemented in a decentralized way? Qinghua Lu and Xiwei Xu share their story of developing the originChain system, which leverages emerging blockchain technology to do so. —*Cesare Pautasso and Olaf Zimmermann*

A TRACEABILITY SYSTEM enables tracking products by providing information about them (for example, originality, components, or locations) during production and distribution. Product suppliers and retailers usually require independent *traceability service providers* who are government-certified to inspect the products throughout the supply chain. If everything satisfies the requirements, the traceability service providers issue inspection certificates that verify the products' quality and originality. To expose information and generate certificates, these service providers employ a traceability system.

In this context, security is important for accountability and forensic information. Traceability systems normally store information in conventional databases controlled by the service providers. Such centralized data storage becomes a potential single point of failure and runs the risk of tampering.

One of our partners is an independent third-party traceability service provider certified by the Chinese government. Its system provides traceability information for products imported to China. This system has been integrated with several big Chinese e-commerce websites (for example, JD.COM). Hundreds of product suppliers and retailers use its traceability services to manage their products' traceability information, and millions of product consumers use it to access the information.

Here, we share our experience of building originChain. It restructures the service provider's current traceability system by replacing the central database with a blockchain. (For more on blockchains, see the two sidebars.) OriginChain provides transparent tamper-proof traceability data, enhances the data's availability, and automates regulatory-compliance checking. We implemented originChain and tested it under realistic conditions

2469-7087/19/$33.00 © 2019 IEEE

BLOCKCHAINS AND SMART CONTRACTS

As a data structure, a blockchain is an ordered list of blocks that contain transactions such as monetary transfers and smart-contract creation and invocation.[1] Each block contains a hash of the previous block's representation, thus creating the chain. So, historical transactions in the blockchain can't be deleted or altered without invalidating the chain of hashes. Combined with computational constraints and incentive schemes for block creation, this can prevent the tampering with and revision of information in the blockchain.

Blockchains also provide a general-purpose programmable infrastructure. Programs can be deployed and run on a blockchain; such programs are called *smart contracts*.[2] The result of a smart-contract invocation is stored in public data storage. Smart contracts can express triggers, conditions, and business logic to enable more complex programmable transactions. Hence, smart contracts differ from service contracts in service-oriented computing, which are interfaces between services and consumers so that they can successfully interact. A common simple example of a smart-contract-enabled service is escrow, which holds funds until the obligations defined in the smart contract are fulfilled.

PROPERTIES

Any data in a committed transaction eventually becomes immutable. The immutable chain of cryptographically signed historical transactions provides nonrepudiation of the stored data. Cryptographic tools also support data integrity, the public access provides data transparency, and every participant has potentially the same ability to access and manipulate the blockchain. However, those rights can be weighted by the participants' stake or computational power. To facilitate transactions, the participants rely on the blockchain network itself instead of trusted third-party organizations.

LIMITATIONS

Blockchains lack data privacy; there are no privileged users, and, as we just mentioned, every participant can access all the information on the blockchain. In addition, public blockchains have limits on the amount of data, transaction processing rate, and data transmission latency. *Consortium blockchains*, in which the consensus process is limited to several participants, perform much better. However, developers still must consider these factors when designing systems.

Furthermore, some public blockchains use a *proof-of-work* consensus mechanism that "wastes" significant electricity because it doesn't lead directly to a successful solution. Researchers are developing alternative consensus mechanisms for public blockchains. One example is the *proof-of-stake* mechanism, which isn't computationally expensive. Consortium and private blockchains also often use consensus mechanisms that don't rely on proof of work.

References

1. M. Swan, *Blockchain: Blueprint for a New Economy*, O'Reilly Media, 2015.
2. S. Omohundro, "Cryptocurrencies, Smart Contracts, and Artificial Intelligence," *AI Matters*, vol. 1, no. 2, 2014, pp. 19–21.

employing the users' traceability information. We're planning how to replace the existing system with the restructured one.

Dealing with Traceability

Product suppliers and retailers apply for traceability services for different purposes. Suppliers want to receive certificates to show their products' origin and quality to consumers and to comply with regulations. Retailers want verification of the products' origin and quality.

Each product supplier that uses our partner's services has on average 20 products to be traced. The traceability information's granularity is rather large because it corresponds to product packages rather than individual products. This information's size isn't easy to estimate because many documents currently aren't digitized, such as certificates issued by traceability service providers.

Traceability is flexible. Figure 1 shows a simplified possible process in BPMN (Business Process Model and Notation).[1] In the real world, the sequence of some activities in Figure 1 (such as "examine factory" and "test sample") is dynamic owing to customization of the quality control and inspection processes. The labs that test samples must adapt to the labs' availability and the characteristics of the products, such as powdered milk or clothing materials.

Furthermore, regulatory-compliance checking can change because of new regulations. For example, China's

BLOCKCHAINS IN THE SUPPLY CHAIN

The supply chain is a promising area for applying blockchains.[1] There are blockchain startups in this field. For example, Everledger (www.everledger.io) uses blockchains to track diamonds' features, such as cut and quality, and to help reduce risk and fraud for banks, insurers, and open marketplaces.

Big enterprises are also applying blockchains in supply chains for different domains. For example, in January 2017, Microsoft started the Manifest project through a partnership with Mojix to leverage an Internet-of-Things platform with a blockchain to help factories, distribution centers, and retailers track goods using RFID devices.[2] In May 2017, Manifest grew to 13 partners.[3] BHP Billiton has been exploring blockchain technology to track movements of wellbore rock and fluid samples and secure the real-time data generated during the samples' delivery.[4]

References

1. M. Staples et al., *Risks and Opportunities for Systems Using Blockchain and Smart Contracts*, Commonwealth Scientific and Industrial Research Org., 2017.
2. M. del Castillo, "Microsoft Unveils Project Manifest, a Plan for Blockchain Product Tracking," CoinDesk, 25 Jan. 2017; www.coindesk.com/microsoft-unveils-project-manifest-a-plan-for-product-tracking-via-blockchain.
3. M. del Castillo, "Microsoft's Blockchain Supply Chain Project Grows to 13 Partners," CoinDesk, 3 May 2017; www.coindesk.com/microsofts-blockchain-supply-chain-project-grows-to-13-partners.
4. P. Rizzo, "World's Largest Mining Company to Use Blockchain for Supply Chain," CoinDesk, 23 Sept. 2016; www.coindesk.com/bhp-billiton-blockchain-mining-company-supply-chain.

Food Safety Law, which took effect in October 2015, set out new requirements for formulating national food safety standards and traceability systems (specifying what information should be provided). So, adaptability was one of our main concerns when we designed originChain.

OriginChain's Architecture

Figure 2 illustrates originChain's architecture. OriginChain currently employs a geographically distributed private blockchain at the traceability service provider, which has branch offices in three countries. The plan is to establish a trustworthy traceability platform that covers other organizations, including government-certified labs, big suppliers, and retailers that have long-term relationships with the company (such as e-commerce companies that have already built their reputation among customers). Compared to a public blockchain, such a *consortium blockchain* can perform better and cost less.

Blockchains grow continually because the data and code on them are immutable. So, a major design decision is to choose what data and computation to keep on-chain and off-chain. We discuss this in more detail later.

As Figure 2 shows, product suppliers or retailers manage product or enterprise information through the *product and enterprise management* module. They access the information on the blockchain through a webserver hosted by originChain. In the future, suppliers and big retailers that host a node by themselves will be able to access their own nodes to obtain information on the blockchain.

After the traceability service provider validates an application from a product supplier or retailer on the basis of paperwork (see Figure 1), the two parties sign a legal agreement about which traceability services are covered. OriginChain generates a *smart contract* that represents the legal agreement. (For more information on smart contracts, see the sidebar "Blockchains and Smart Contracts.") The smart contract codifies the combination of services and other conditions defined in the agreement. So, the smart contract can automatically check and enforce these conditions. It also checks whether all the information required by regulation is provided, to enable automated regulatory-compliance checking.

The traceability service provider manages traceability information, certificates, and onsite photos using the *traceability management* module. Because of the blockchain's limited data storage, originChain stores two types of data on-chain as variables of smart contracts:

- the hash of traceability certificates or photos and

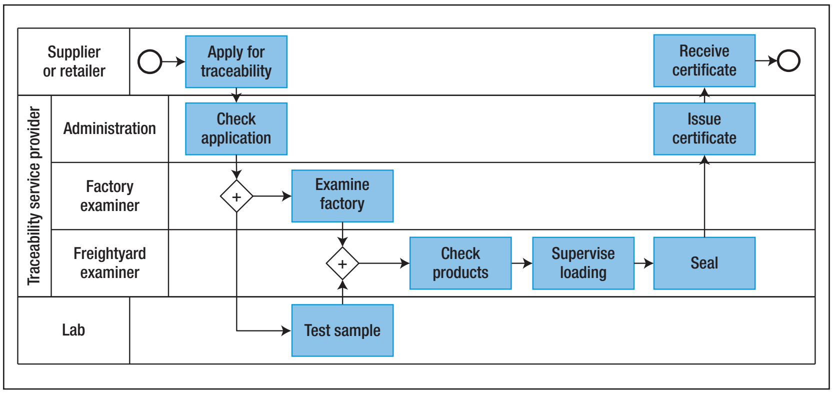

FIGURE 1. The dynamism of the traceability process. The sequence of some activities (such as "examine factory" and "test sample") is dynamic owing to customization of the quality control and inspection processes.

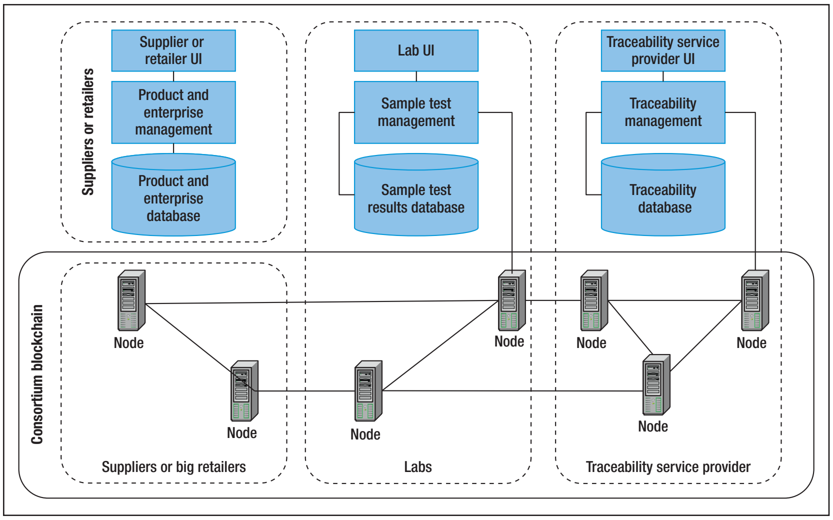

FIGURE 2. OriginChain's architecture. OriginChain currently employs a geographically distributed private blockchain at the traceability service provider company. The plan is to establish a consortium blockchain, which will include other organizations.

- the small amount of traceability information required by the traceability regulation, such as the batch number, traceability results, place of origin, and inspection date.

The raw files of traceability certificates and photos (.pdf or .jpg) and the addresses of the smart contracts are off-chain in a centralized MySQL database hosted by originChain. Other partners can still maintain their own database of product information (for the suppliers or retailers) or other numbers shown in the sample testing (for the labs).

The labs manage sample-testing results through the *sample test management* module. A blockchain's execution environment is self-contained. So, a smart contract can access only information stored in the blockchain. It can't directly access the states of external systems (for example, testing results and product geolocations). Thus, the labs periodically inject the result of sample testing from the external world into the blockchain.

The information of blockchain-layer permission control (for example, permission for content management, for writing smart contracts, or for joining a consortium blockchain) can be on-chain or off-chain. However, an off-chain centralized permission management module could become a single point of failure from both an operational and a management perspective. So, originChain stores the control information, such as permission to join the blockchain network (to own a copy of all the historical transactions). On-chain permission management leverages the blockchain's decentralized nature so that all the participants can access the blockchain.

Figure 3 shows how smart contracts are designed using our blockchain; Figure 4 shows part of the related pseudocode. In originChain, a *factory contract* creates smart contracts. This reduces the complexity of creating customized smart contracts. The factory contract contains code fragments representing different traceability services. The generation of smart contracts requires authority from both the traceability service provider and the supplier or retailer.

When the factory contract is called, it creates two kinds of smart contract: a *registry contract* and *service contract*. The registry contract represents the legal agreement and contains the address of the service contract, which codifies the legal agreement. The service contract could be updated by replacing its address stored in the registry contract with the address of a new version. Possible updates include adding or removing services from the legal agreement after the initial legal agreement is signed, or selecting labs for sample testing on the basis of their availability. The registry contract specifies a list of addresses allowed to update the registry contracts, and a threshold of the minimum number of addresses required to authorize an update.

If a testing sample involves multiple labs for crosschecking, signatures from all the labs are required before the traceability application can undergo further processing. To enable more dynamic lab selection, users can employ an *M*-of-*N* multisignature to define that *M* out of *N* labs are required to authorize the testing results.

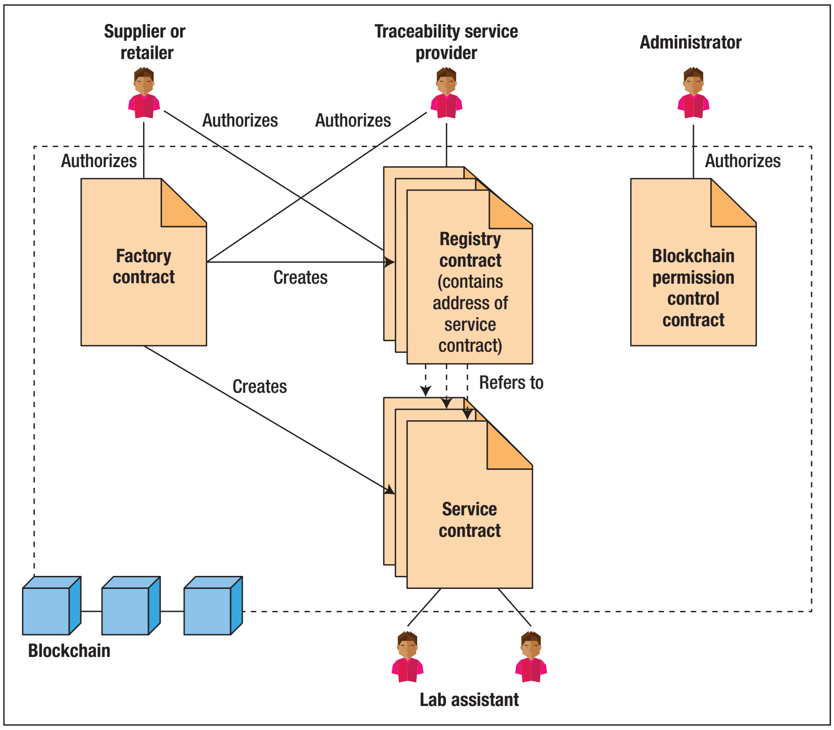

FIGURE 3. Designing smart contracts in originChain. The factory contract reduces the complexity of creating customized smart contracts.

Lessons Learned

Our experiences with originChain led to the following insights.

The Design of Blockchain-Based Systems

Owing to blockchains' unique properties, some design considerations are specific to blockchain-based applications—for example, the consideration of on-chain and off-chain. On the other hand, because smart contracts are programs running on a blockchain, some existing architectural patterns might be applicable to them. From the business process perspective, approaches such as model-driven development[2] and behavior-driven development[1] are also applicable.

The smart contract's structural design has a large impact on the cost if the blockchain is public. The contract's deployment cost depends on its size because the code is stored in the blockchain, which entails data storage fees proportional to the contract's size. So, a structural design with more lines of code costs more money.

A consortium blockchain doesn't have to have a token or currency, so monetary cost isn't an issue there. However, the blockchain's size is still a design concern because it grows with every transaction and because every participant has a replica of the whole blockchain. In addition, a more structural design might affect performance because it might require more transactions.

On-Chain vs. Off-Chain

Regarding what should be on-chain and off-chain, two factors are important: performance and privacy. Performance depends highly on the blockchain's deployment. For example, a consortium blockchain can be configured to perform much better than a public blockchain.

With originChain, owing to the current traceability system's characteristics (for example, it has low writing throughput because of the large granularity of traceability information), a blockchain's limited throughput isn't the main concern. However, as we mentioned before, all the participants can access the data on the blockchain. So, private data (for example, customer information) shouldn't be on-chain. Regarding traceability, large sensitive raw data (for example, traceability certificates and photos) must be immutable. Thus, the raw data is off-chain, whereas its hash is on-chain.

```
contract FactoryContract {
    address[] registryContracts;
    address[] serviceContracts;
    // deploy a new registry contract
    function newRegistryContract() returns(address
newRegistryContract){
        RegistryContract r = new RegistryContract();
        registryContracts.push(r);
        return r;
    }
    // deploy a new service contract
    function newServiceContract() returns(address
newServiceContract){
        ServiceContract s = new ServiceContract();
        serviceContracts.push(s);
        return s;
    }
}

contract RegistryContract{
    address[] serviceContracts;
    address[] authorities;
    uint threshold;
    uint authorityNum;
    function authorizeVotingRight(address authority) {
        authorities.push(authority);
    }
    function setThreshold(uint threshold) {
        threshold = threshold;
    }
    function vote(address authority) {
        authorityNum++;
    }
    function update() returns(address newServiceContract){
        // If there are enough authorities then
        …
        If(authorityNum >= threshold){
            ServiceContract s = new ServiceContract();
            serviceContracts.push(s);
            return s;
        }
    }
}

contract ServiceContract {
    bool testsampleselected;
    bool examinefactoryselected;
    bool superviseloadingselected;
    function ServiceContract (bool testsample, bool examinefactory,
bool superviseloading){…}
    function testSample(){…}
    function examineFactory(){…}
    function superviseLoading(){…}
```

FIGURE 4. Pseudocode for the design of a smart contract in originChain. The generation of smart contracts requires authority from both the traceability service provider and the supplier or retailer.

The Adaptability of Blockchain-Based Systems

Adaptability is a quality attribute required by many industrial projects that are inherently dynamic. For example, changes to the legal agreement or new regulations could necessitate adaptation. Adaptation here means that the smart contract could be updated by a number of authorities above the threshold defined in the factory contract. However, research on blockchain-based systems rarely discusses adaptability.

We view the blockchain as a component of a larger distributed system. In originChain, we implement some of the business logic on-chain as smart contracts. Thus, smart contracts' structural design also affects their updatability and the whole system's adaptability (for example, separation between data and control).

However, if the blockchain is for data storage only, not much can be done to affect the whole system's adaptability. Moving some logic to the blockchain can leverage the trustworthiness (and the interoperability and the transparency of data and operations) that the blockchain provides as a computational platform. In addition, the data in smart contracts is easier to query (directly on the blockchain) than is the data in transactions.[3]

Access Control for Smart Contracts

Smart contracts running on a blockchain can be accessed and called by all the participants. A smart contract, by default, has no owner; once it's deployed, its author no longer has any special privileges on it. Unauthorized users could accidentally trigger a permissionless function. So, smart contracts should have an embedded permission control mechanism to check permission for every

This article originally appeared in IEEE Software, *vol. 34, no. 6, 2017.*

caller that triggers the functions defined in the contracts.

Traceability processes in supply chain management are complex and dynamic because they involve multiple parties. A blockchain provides neutral ground that should help integrate the disparate participants into those processes. Also, the integrity and audit trail in a blockchain ledger should improve transparency and confidence across the processes.

Although joining a consortium blockchain benefits all the relevant stakeholders, adopting a new technique such as a blockchain is always a challenge to traditional industries because of the learning curve and the cost of integrating the blockchain into the existing systems. Negotiating the business details also takes time. In addition, the development of smart contracts must take into account quality attributes such as adaptability.

Data transparency and sharing data with others are main concerns for most companies that provide intermediary services in industries. Overall, blockchains are a good option for providing traceability in supply chain management. Nevertheless, industry needs to take the time to understand their risks and opportunities.

ABOUT THE AUTHORS

QINGHUA LU is an associate professor in the College of Computer and Communication Engineering at China University of Petroleum (East China). Contact her at qinghualu@upc.edu.cn.

XIWEI XU is a research scientist in Data61 at the Commonwealth Scientific and Industrial Research Organisation (CSIRO). Contact her at xiwei.xu@data61.csiro.au.

Acknowledgments

The National Natural Science Foundation of China (grant 61402533) and the Fundamental Research Funds for the Central Universities (grant 16CX02047A) supported this project.

References

1. D. Lübke and T. van Lessen, "Modeling Test Cases in BPMN for Behavior-Driven Development," *IEEE Software*, vol. 33, no. 5, 2016, pp. 15–21.
2. I. Weber et al., "Untrusted Business Process Monitoring and Execution Using Blockchain," *Business Process Management*, LNCS 9850, Springer, 2016, pp. 329–347.
3. X. Xu et al., "A Taxonomy of Blockchain-Based Systems for Architecture Design," *Proc. 2017 IEEE Int'l Conf. Software Architecture* (ICSA 17), 2017, pp. 243–252.

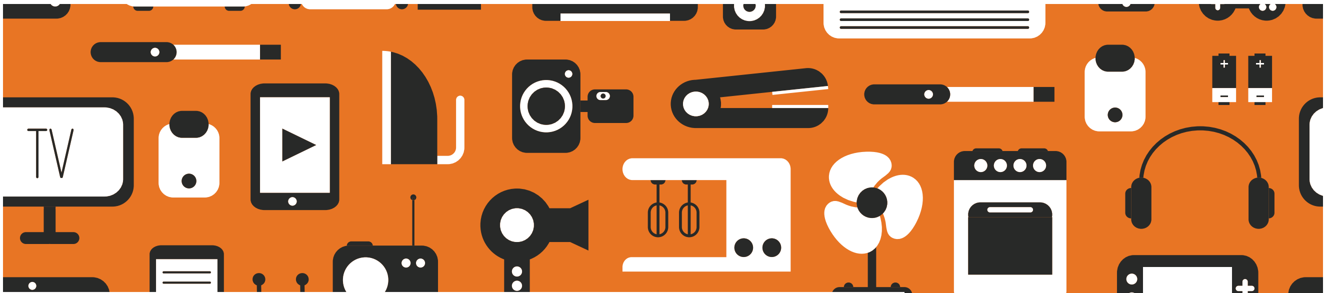

Accurate Indoor Location for the IoT

Roy Want, Wei Wang, and Stan Chesnutt, Google

The combination of IoT and accurate location-finding technologies will enable a wide range of novel context-aware services to automate tasks and support everyday work practice.

A defining characteristic of the Internet of Things (IoT) is that devices are both networked and embedded in the environment. Unlike traditional networked server technology that has very little connection with the environment, IoT devices are designed and positioned so that they can respond to and influence their surroundings. For example, if a user accesses a web server from an office in California, it makes very little difference whether the service is located in the next building or situated in Europe. However, if a user accesses an IoT device that is responsible for measuring the temperature of a room, the sensor and device need to be located in that specific room.

As a result, IoT design is a unique departure from infrastructural networking that tries to abstract away the network topology; instead, the topology and position of IoT nodes are assets in performing the networking task.

IoT devices often contain sensors and actuators. Consumer products for IoT sensors include light sensors, microphones, cameras, security switches, motion detectors, thermostats, barometers, moisture sensors, and when extended to industry, a wide variety of exotic sensors. Actuator outputs include lamps, valves (water sprinklers), speakers, and a variety of displays—again, industry has many additional needs for bespoke actuators. The interrelationship of the position and function of each of these devices allows us to design systems that are often called "smart homes" or "smart buildings."

An additional consideration of spatially arranged infrastructure is that mobile wireless components can move through it. In a smart home, people will be present and the infrastructure is intended to make their lives easier and more comfortable. Devices such as lights, HVAC, and security systems can be turned on or off by the presence and detection of people entering or leaving an area. This is not only convenient, but also addresses human forgetfulness and has cost-saving implications because it reduces energy consumption.

In the modern smart home, voice-driven assistants have become popular (such as Amazon Echo and Google Assistant) and allow devices in smart homes to be activated by casual voice commands. Typically, human speech is often tersely phrased with the assumption that the context is understood. This is challenging for a computer, which doesn't have the shared context to respond

 2469-7087/19/$33.00

EDITOR ROY WANT
Google; roywant@google.com

to these commands. One of the biggest qualifiers of context is a user's position. For example, an ambiguous command such as "turn on this lamp" can be immediately qualified if the system knows the relative position of the user and all the nearby lamps.

Although it might be hard to determine the location of a person, it is now possible to accurately calculate the position of a mobile device using Wi-Fi radios (see www.ieee802.org/11) with positioning support. Since the early Wi-Fi (802.11b) products appeared in 1999, Received Signal Strength Indication (RSSI) has been used to determine the approximate position of a mobile device relative to surrounding access points installed at known positions.[1] More recently, Bluetooth low-energy beacons have been used in a similar way, but at higher density due to their small size and low cost. However, due to multipath fading, the RSSI fall-off from the antenna does not follow a uniform inverse square law and is distorted by the building materials that it travels through. As a result, the position accuracy for Wi-Fi can often have large errors of 10 to 20 meters. Improvements have been made by measuring the RSSI fingerprints for all detectable access points at each square meter of a building, but this requires considerable investment in surveys, which are labor intensive and have to be regularly updated.

A more accurate approach was needed that could localize a device to about 1 square meter (the area occupied by a person), but only required the knowledge of Wi-Fi access point positions.

THE IMPORTANCE OF A STANDARD

In December 2016, the IEEE 802.11 working group ratified amendment 802.11-REVmc[2] for the Wi-Fi standard, which included a new ranging capability based on time of flight (ToF) of Wi-Fi packets. As these packets are transmitted as RF signals that travel at the speed of light *c*, the distance can be calculated by knowing the ToF of a packet from a mobile to an access point and multiplying by the speed *c*. Knowing the distance or range to three or more access points (and using a process called multilateration), the mobile's position can be calculated (see Figure 1).

In recent times, there have been other radio ranging standards (for example, ultrawideband [UWB], which was standardized in IEEE 802.15.4a[3]), but none have had much impact on the mobile industry. Wi-Fi is different because of its wide-scale use in mobile phones, laptops, and tablets throughout the world. And most importantly, Wi-Fi access points' infrastructure has grown as a grassroots effort and is deployed ubiquitously, including in many public spaces.

The spread of Wi-Fi has been enhanced because of three value propositions:

- it has a practical range of about 25m, and as a result the density and related cost of access points to cover a work area are not great;
- its power consumption is low enough that it can be embedded into small mobile devices with acceptable impact on battery life (for smartphones and tablets); and
- the data rate of Wi-Fi products has remained competitive and has typically had a x10 advantage over cellular data, which often has a data limit and results in additional cost if exceeded.

Wi-Fi has an additional characteristic in that it has already been improved several times to meet growing user demands. There is a natural upgrade cycle for Wi-Fi as radio technology improves and as users crave higher data rates to support video and faster downloads. These upgrades also serve as a vector for new features (such as ranging) that can be introduced in each upgrade cycle.

In short, Wi-Fi 802.11-REVmc is ideally suited as a standard to disseminate a new positioning technology with the potential for a wide reach and the opportunity for mobile application developers to build on its novel capabilities.

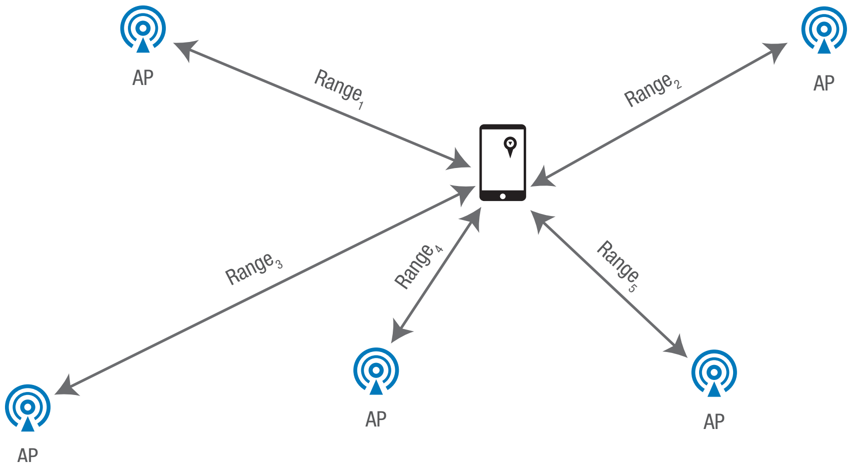

Figure 1. Multilateration, the process of converting ranges to a position. AP: access point.

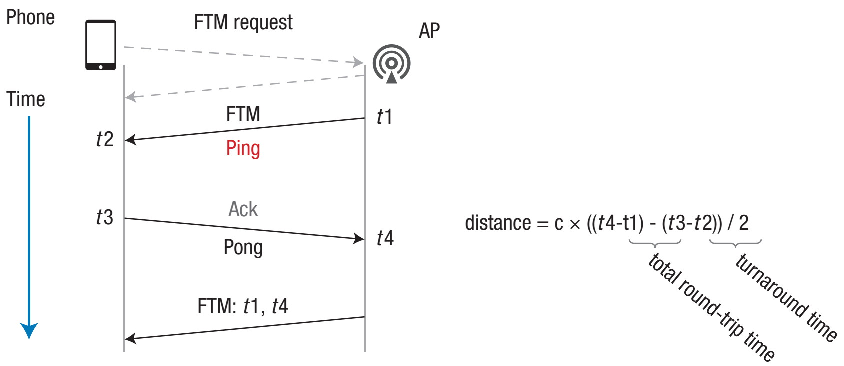

Figure 2. The 802.11mc Fine Timing Measurement (FTM) Protocol.

IEEE 802.11-REVMC AND WI-FI ROUND-TRIP TIME

The core of the 802.11-REVmc protocol is built around a new packet type called a Fine Timing Measurement (FTM) frame. The recipient responds to an FTM with an acknowledgement packet, and the transmission and reception times of each are recorded. It's basically a ping-pong protocol that calculates ToF by measuring the time it takes for a packet to be sent from an access point to a mobile and then back again. This is called the round-trip time (RTT), so the protocol is sometimes referred to as Wi-Fi RTT.

For a mobile device to measure the ToF, the precise times of the packet leaving the access point ($t1$) and arriving at the mobile receiver ($t2$) must be known. Further, the precise times the mobile device transmits the response ($t3$) and the time it is received back at the access point ($t4$) must also be recorded. From the mobile's perspective, it measures only $t2$ and $t3$ as they occur locally at its radio. However, $t1$ and $t4$ are measured at the access point, so for the mobile to calculate the RTT, the access point needs to send $t1$ and $t4$ as an additional message to the mobile (see Figure 2), which is piggybacked on the next FTM packet.

The phone can now calculate the RTT based on the total time $t4 - t1$, subtracting $t3 - t2$ to account for the turnaround time of the acknowledgement packet.

$$\mathrm{RTT} = (t4 - t1) - (t3 - t2)$$

The desired range is determined by multiplying by the speed of light c and dividing by 2, as we are only interested in the distance from mobile to access point and not the round-trip distance.

$$\mathrm{range} = c \times \mathrm{RTT} / 2$$

Some general assumptions are made that the clocks are running at about the same rate on the two devices, and with a modern crystal oscillator with an accuracy of 40 parts per million, this is reasonable. Note that the clocks of each device don't need to be synchronized as the $(t4 - t1) - (t3 - t2)$ terms provide time differences for each separate clock at the beginning and ending of each exchange, and therefore do not require an absolute time reference. Given the typical clock stability and an 80 MHz bandwidth in the 5 GHz Wi-Fi band, the expected accuracy of this protocol is expected to be ~1m, which meets the design goals and considerably outperforms approaches based on RSSI.

INCREASED LOCATION ACCURACY ENABLES NOVEL APPLICATIONS

Important questions for the design of any positioning system are "What is the position system going to be used for?" and "What target accuracy is needed?" If a system has an accuracy of 10m and is improved to 9m or even 8m, does it make the system more useful? Perhaps not that much. However, if you were to improve it to 1m, it would cross a significant human threshold because a person typically occupies about 1 square meter of standing space. It opens up the possibility to place a person next to another person, product, or thing.

For Wi-Fi RTT, 1 to 2m is just less than the width of most corridors or aisles found in a building. This resolution actually enables in-building navigation, for example, directing a person to a product in a store or an office worker to an unfamiliar conference room. Errors two to three times this size would cause the indoor navigation system to place you in the wrong aisle, resulting in a confusing user experience. In the same way that GPS is accurate enough (~5m) to place a car on the correct road, in-building navigation is enabled by a 1 to 2m Wi-Fi location system. In combination with indoor maps and a routing system calculating the shortest path, Wi-Fi RTT can show your progress along a route and provide corrections by rerouting when you deviate from the path.

In the IoT world, ~1 meter accuracy can automatically determine the location of devices such as lights and thermostats, which might not have their exact position recorded. Positioning to ~1m can automate the documentation for an installation process, reducing contractor time and removing transcription errors that inevitably occur in hectic construction environments. Further, it can automatically keep records up to date as devices are replaced and during remodeling.

In daily work practices, users carrying mobile devices with Wi-Fi RTT can be located in a smart home (using an application with permission from the owner) and automate the operation of heating, lighting, door locks, and security systems. The modern smartphone is a multipurpose device, and owners typically carry the device with them most of the time. As home monitoring and control capabilities grow via smartphone apps, we expect that users will make use of these apps in their everyday lives.

There are further advantages for accurate location in peer-to-peer interactions. When two devices are next to each other, there is the opportunity to share information. For example, if a user wishes to share a photo, it might be done with a simple flick gesture on a screen. A phone that can detect the proximity of nearby devices can limit the scope of gesture-based sharing to the closest device. You can imagine many other forms of digital peer-to-peer transfers that are made possible by automatically and accurately detecting close proximity, which includes ticketing and payments.

LOCATION DATA AND PRIVACY

We have already seen that accurate location can bring many advantages to daily work practice and social interactions. However, there are important privacy implications to consider, and great care needs to be taken when implementing and enabling these services for users.

Mobile devices that can repeatedly link a user's identity to the device's location create a location "trail," which could result from location data being stored persistently on the device or uploaded to a server. With nefarious access to this information over time, it's straightforward to determine the habitual whereabouts of the user, which could be used to tip off thieves or others with malicious intent.

Protecting the user's privacy starts with an agreement with the user. First, the software must request explicit permission to establish the user's location and to store that information. Second, persistent location information must be carefully protected after being uploaded to a central server (for example, protection against reuse in unexpected ways and against observation by a rogue employee). This need for protection also applies to raw ranging data, not just position, as it can also be converted to user location by a nefarious party.

Furthermore, an IoT device might have capabilities that can threaten user privacy when coupled with location data. For example, a security camera that uploads video data could also add identity and location information.

The European Union's General Protection Data Regulation (GDPR) establishes clear standards for data retention, storage, and use. These guidelines should be treated as a starting point for data stewardship.

THE FUTURE OF LOCATION TECHNOLOGIES

Going forward, we expect 802.11-REVmc to be adopted by a wide variety of client devices and access points, which will result in many novel location-based applications that can make use of ~1m indoor accuracy. The designers of the first Wi-Fi RTT protocol had many additional features that they wanted to include, but they had to be deferred due to the timing of the first release.

A more extensive set of location protocol features are currently being undertaken by the follow-up task group 802.11az (Next Generation Positioning), whose primary goals are to enable greater scalability and provide for secure operation.

The 802.11az task group is developing a protocol that takes advantage of the new high-efficiency (HE) data standard being developed in parallel in Task Group 802.11ax with four times the data throughput of existing 802.11ac-based systems. In addition to higher-density coding, the radio supports multi-user uplink MIMO (multiple input, multiple output), allowing more spatial parallelism in data streams than before and resulting in more efficient use of the communication channel. Location techniques can take advantage of this mechanism by ranging between an access point and several stations simultaneously, thus supporting a larger number of potential clients.

An additional technique that supports even larger populations of clients is called passive ranging. This mechanism is characterized by mobile devices that only need to listen to the transmissions made by other stations and never need to transmit themselves. The result is that the channel capacity is not affected by the number of users trying to measure their position—furthermore, privacy is guaranteed as a mobile never needs to announce its presence. In that sense it's similar to the operation of GPS, although it differs in the geometric calculation needed to determine a position.

An IoT device might have capabilities that can threaten user privacy when coupled with location data.

In this technique, special types of access points called anchor stations range to nearby access points. If a mobile device listens to the protocol exchanges between anchor stations and access points, it can use the reported timestamps to figure out the time difference of arrival of two packets at its receiver, as if they were sent simultaneously from two known positions. This information allows a mobile to place its position on a hyperbolic curve. When repeated for multiple pairs of anchor stations and access points and at known positions, the mobile can calculate its position by determining the intersection point of these hyperbolic curves.

Authentication and security are

always important—even more so when wireless devices are involved, because it's often hard to determine the entity that a mobile device is communicating with. For example, when you are in an unfamiliar location such as a large airport, are you aware of the access point your smartphone is associated with? 802.11az sets out to address this problem and provides the following mechanisms.

- A mobile device can ensure it's ranging to a trusted infrastructure. In other words, a mobile device should be able to reject potentially false range or positioning information from a nearby rogue access point.
- If a mobile device A is ranging to nearby access points, it should not be possible for another nearby nefarious mobile device B to overhear and decode the ranging packets. Therefore, B should not be able to calculate the position of A.
- Infrastructure that determines the location of a mobile device (assuming a user has given it permission to use that information) should be able to validate that the mobile is not falsifying its position.

Lastly, 802.11az provides for operation in the 60 GHz band, which promises greater range accuracy to a few centimeters with the increased bandwidth. In practice, this means position accuracy could have 10 times the resolution of 802.11-REVmc, which will open up new categories of location-based applications.

Location technology is coming of age at a time when IoT technologies are proliferating. The combination of IoT and accurate location-finding technologies will enable a wide range of novel context-aware services to automate tasks and support everyday work practice. Over the next few years, we expect fine-grained location-based services to become more prevalent and rapidly become a standard feature used by smart infrastructure.

REFERENCES

1. P. Bahl and V.N. Padmanabhan, "RADAR: An In-Building RF-based User Location and Tracking System," *Proc. 19th Ann. Joint Conf. IEEE Computer and Communications Societies* (INFOCOM 2000), 2000, doi:10.1109/INFCOM.2000.832252.
2. *P802.11-REVmc/D8.0, Aug 2016—IEEE Approved Draft Standard for Information Technology— Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks— Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, revision of IEEE Std. 802.11-2012, IEEE, 2016; https://ieeexplore.ieee.org/document/7558098.
3. E. Karapistoli et al., "An Overview of the IEEE 802.15.4a Standard," *IEEE Communications Magazine*, vol. 48, no. 1, 2010, pp. 47–53.

ROY WANT is a research scientist at Google. Contact him at roywant@google.com.

WEI WANG is a software engineer at Google. Contact him at weiwa@google.com.

STAN CHESNUTT is a software engineer at Google. Contact him at email chesnutt@google.com.

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

Editors: Heather Richter Lipford, richter@uncc.edu | Jessica Staddon, jessica.staddon@gmail.com

The Last Mile for IoT Privacy

Richard Chow | Intel Corporation

According to Mark Weiser:[1]

> *The problem, while often couched in terms of privacy, is really one of control. If the computational system is invisible as well as extensive, it becomes hard to know what is controlling what, what is connected to what, where information is flowing, how it is being used ... and what are the consequences of any given action.*

Weiser was discussing ubiquitous computing more than a decade ago, but he might as well have been talking about the Internet of Things (IoT) today. On our computers, we have at least a semblance of control because we can, in principle, determine what applications are running and what data they're collecting. For the IoT, traditional methods of control are largely absent. In fact, there are common cases where people are no longer *users* of an IoT service but rather *subjects* of the service, such as a smart city sound monitor. Another example is public Wi-Fi connectivity. A fraction of people in the vicinity might discover and use the service, but not being the ones who installed the actual access points, most people would be unaware of the service's privacy properties.

One of the IoT's major privacy problems is that users aren't always aware when a device is collecting personal data. IoT devices' ubiquitous nature means that a person can easily not know when sensors are present. A basic privacy tenet, dating back to the Fair Information Practice Principles, states that personal data collection should happen only with appropriate notice. Therefore, one of the goals in this article is to provide a framework—called the privacy stack—for user communication issues and needs regarding IoT privacy (see Figure 1).

The Privacy Stack

Many design questions, such as discovery, usability, and privacy, involve end users. How do users learn about the services in public spaces, their workplaces, or even their own homes? How do they learn about these services' privacy properties and yet not be overwhelmed by the sheer volume of information? I give a couple motivating examples of IoT privacy notifications.

Suppose a mall installs a suite of surveillance cameras in its parking structures for security purposes. Rather than describing the cameras' purpose on physical signage, the mall uses beacons to interact with phones carried by passersby. The beacons emit an ID that can be looked up in the cloud using a standard mobile app. The users can thus be alerted to the surveillance cameras' presence, which not only serves as a privacy notification but also provides them with a sense of security knowing that the area is safe due to the service.

Privacy notifications can also be incorporated into smart home systems. Consider open space voice controllers such as the Amazon

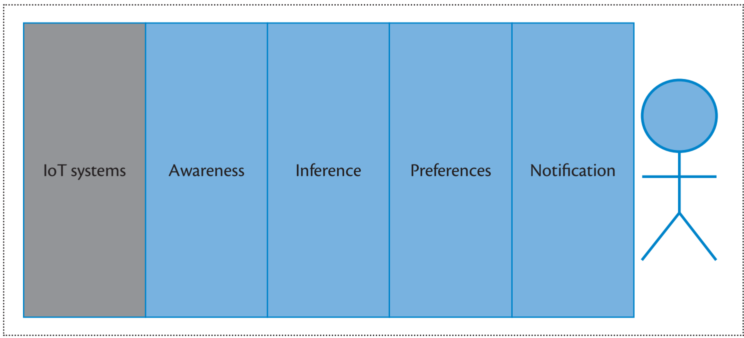

Figure 1. The privacy stack framework bridges from today's Internet of Things (IoT) systems to users.

Echo, which allow voice commands to play music, adjust lighting, and so on. Guests to the home might not be comfortable with a whole-home voice-recording system. Homeowners might ask permission to use such a service if their guests indicate that they'd like to be notified about audio recording in their vicinity. Depending on the guests' preferences, they could be notified as they enter the house, or alternatively, the system could disable the service temporarily or partially. All this might be done silently or explicitly.

In these motivating examples, the users don't understand what services are active in the environment. In the sections that follow, I describe how the examples might be realized through the privacy stack shown in Figure 1, which takes us from today's IoT protocols to human notification. I emphasize that this stack isn't a protocol stack, but simply a conceptual framework.

Awareness

The awareness part of the stack provides for discovery of services' privacy properties by users or users' agents. The difficulty with the IoT is the potentially unobtrusive nature of data collection. One can, in fact, not realize that data is being collected. According to Judith Donath, a fellow at Harvard's Berkman Center for Internet and Society, we must design so that we can know how public or private a space is, to know how to act and how candidly to speak.[2] Along these lines, Jason Hong has issued a challenge to the pervasive research community: can we make it so that when people enter a room, they can reliably identify all of the sensors and dataflows within 30 seconds?[3]

Awareness primarily concerns how IoT services might open communication channels to users and subjects. These channels might be new visual signifiers of data collection,[4] or more traditional network protocols between the IoT service and a user's device. For protocols, multiple industry efforts, including the Open Connectivity Foundation (openconnectivity.org), are working on standardizing IoT interoperability, allowing discovery and communication among devices.

Apple's iBeacon is another example of device discovery (developer.apple.com/ibeacon). A beacon device broadcasts an ID through Bluetooth, and a compatible device can use this ID to retrieve associated information (such as for location-based services). Beacons can also place a person in proximity to devices in the environment.

Thus far, the IoT protocol work has not gone into privacy metadata standardization, a way for services to declare their privacy policies so that they are universally understood. One advantage of doing so would be that, in a world of ubiquitous services, privacy decisions can be made with minimal cognitive burden. But even for the web, we don't have privacy metadata standardization for various reasons. One interesting approach that sidesteps standardization is to build natural language processing tools for privacy policies (for example, see www.signifiers.io). Nevertheless, the risks of data collection with completely unaware subjects are greater for the physical world and IoT, which argues for giving a standard language for IoT privacy another try.

Inferences

Awareness protocols can enable IoT services to declare what they're doing, but what do they declare? It seems simplest to declare what sensor data is collected and what it will be used for, but this might be insufficient. Users have limited understanding of what might be inferred or learned from sensor data, and systems will only get better at learning from sensor data. For instance, device or browser fingerprints are now a fact of life on the web, and it's reasonable to assume that fingerprints are even more prevalent in data measuring in the physical world. A couple of examples are location patterns and, more recently, ambient audio.[5] The inference problem is central to privacy, and yet it's unclear how to declare inferences; for instance, how does the system handle probabilistic inferences and inferences with auxiliary data?

I propose a balanced approach here: users can't be relied on to understand the inferences possible from the data collected, so services must explicitly provide basic inferences. At the same time, users must understand that this set of inferences is continually growing and refined, not only through the never-ending growth of user data but also through the advancement of machine learning techniques. One way for IoT services to provide inferences is through the privacy policy; for instance, an IoT service collecting GPS data would declare that identity is a possible inference from the data, and a device collecting energy usage data would declare behavioral patterns as a possible inference.

An understanding of inferences wouldn't just help users understand what the system is learning about them but would also help systems protect privacy by translating user preferences. For instance, a system might collect video data, timestamps, and Wi-Fi data. A user might be comfortable with sharing video data collected at work but not video data collected at home. Security policies for the raw data itself can be implemented, but only after the system uses the data to infer "work" versus "home." This highlights the gap between low-level raw data and the language of human preferences, which consists of higher-level concepts such as "home." Sensors operate at a different level from what users find meaningful, and inferences help bridge the gap to usable systems.

Internet of Things Privacy Preferences: Current Research

What privacy factors are users most concerned about in an Internet of Things (IoT) context? Here's what some recent research has discovered.

Hosub Lee and Alfred Kobsa approached the problem by studying user reactions to scenarios with varying context elements.[1] For each scenario, they chose from a small set of parameter values in the following categories: who was collecting the data, where the data was being collected, what kind of data was being collected, the reason for collection, and the persistence of collection. An example scenario had "government" for who was collecting and "safety" for the reason for collecting. They found clusters in this context space that predicted user reaction.

Reuben Binns and his colleagues pointed out that privacy decision making is influenced by particular attributes, such as the reputation and size, of the company doing the collecting.[2]

Linda Lee and her colleagues investigated the "what" and "who" factors for wearable devices; the "what" had more effect than the "who" on user reactions, and users were most concerned about the collection of video or photos.[3]

Pardis Emami Naeini and her colleagues provided more statistical evidence for the impact of various factors, such as what and where, and also advanced explanations for the impact, such as perceived benefit.[4]

References

1. H. Lee and A. Kobsa, "Privacy Preference Modeling and Prediction in a Simulated Campuswide IoT Environment," *Proc. IEEE Int'l Conf. Pervasive Computing and Communications* (PerCom 17), 2017, pp. 276–285.
2. R. Binns et al., "My Bank Already Gets This Data: Exposure Minimisation and Company Relationships in Privacy Decision Making," *Proc. CHI Conf. Extended Abstracts on Human Factors in Computing Systems* (CHI EA 17), 2017, pp. 2403–2409; doi.acm.org/10.1145/3027063.3053255.
3. L. Lee et al., "Information Disclosure Concerns in the Age of Wearable Computing," *Proc. NDSS Workshop Usable Security* (USEC 16), 2016; dx.doi.org/10.14722/usec.2016.23006.
4. P. Emami Naeini et al., "Privacy Expectations and Preferences in an IoT World," *Proc. 13th Symp. Usable Privacy and Security* (SOUPS 17), 2017; www.usenix.org/conference/soups2017/technical-sessions/presentation/naeini.

Privacy Preferences

Moving on to the stack's next layer, suppose the system has a good understanding of what data an IoT service is collecting and what inferences might be possible. What does the user or subject actually care about? Whether a particular IoT data collection scenario is considered privacy sensitive depends on the individual, but context is also critical. For instance, audio recording might be fine in a restaurant depending on who's doing the recording (restaurant or friend) and who you're with (work or social). Context also gives clues as to what's surprising and what's expected. Video surveillance in a football stadium might be unremarkable, but perhaps not at a restaurant.

Because so many factors might influence a privacy decision, the central question is: what factors do people really care about in an IoT context? The sidebar provides a sampling of what the current research has found.

This research is just beginning, but the vision is that an understanding of the relevant factors will encourage system creators to present this information to those making privacy decisions. Knowing that particular factors might be important to some users can also influence system design. For instance, an IoT service at a bar might set patrons at ease by informing them that video is being captured by the bar for age-verification purposes only.

Notification

The stack's last layer involves notification of the end user. This notification depends on awareness of the services in the environment and considers inferences and user preferences.

Notification is envisioned to be a relatively rare event, as users won't want to interact with most IoT services when they become ubiquitous; even services that merit a notification the first time might not the second or third time. For instance, drivers might want to be alerted when their car enters a region where it might be tracked but won't want to be alerted every day for the same street camera.

Central questions for notifications include: How does a system present alerts, how does a user express alert preferences, and how are these preferences learned? One notable point is that an alert that's not privately communicated could allow others to draw undesirable inferences about a user's history or preferences. Also, technology and inference algorithms might improve enough such that people who were previously notified about a service might need to be notified again.

Notification represents the last layer of the privacy stack: actual interaction with users. But other forms of interaction are possible. Notification implies interrupting a user involved in another activity, but a more active user might be interested in visualizing the privacy properties for nearby IoT services. How does a user visualize nearby IoT services and their privacy and trust properties?

Implementing the Stack

To illustrate, I describe how the stack might be implemented in the commonly proposed privacy proxy or privacy assistant on a user's device.[6,7] For awareness, the privacy proxy manages protocols that pull (or are pushed) privacy metadata from nearby IoT services. Basic inferences might be included in the privacy metadata, and the privacy proxy might also have the intelligence to make its own inferences from the data. The proxy manages personalized privacy preferences, for instance, allowing users to declare preferences, or inferring them from past privacy decisions or demographics. A privacy proxy builds and maintains a user's privacy preferences and decisions, and thus the proxy can use intelligence to alert the user only when appropriate.

A basic privacy principle is that personal data collection should happen only with appropriate notice and choice. However, this principle's implementation—already difficult for traditional clients—is even more difficult for the IoT because there's no natural communication channel with users. In many cases, users might find themselves in environments where services are running and capturing their data, even though the users haven't installed the services and aren't aware that they're running.

A conceptual privacy stack describes how IoT systems being developed today interact with users. This stack is being actively researched in multiple areas. One area is the user interface, both for setting and learning privacy preferences and for notification and visualization. A second area is the inference problem for IoT data—how to translate raw sensor data into human-understandable inferences. A third area of research is the privacy schema and how to represent the plethora of IoT services and their privacy policies, which is critical for automatic filtering. This is an exciting and active area, combining machine learning, human–computer interaction, and privacy. ■

References

1. M. Wieser et al., "The Origins of Ubiquitous Computing Research at PARC in the Late 1980s," *IBM Syst. J.*, vol. 38, no. 4, 1999, pp. 693–696.
2. "Five Minutes with Judith Donath," The MIT Press, 30 July 2014; mitpress.mit.edu/blog/five-minutes-judith-donath.
3. J. Hong, "The Privacy Landscape for Pervasive and Ubiquitous Computing," *IEEE Pervasive Computing*, vol. 16, no. 3, 2017, pp. 40–48.
4. S. Zimmeck et al., "Automated Analysis of Privacy Requirements for Mobile Apps," *Proc. 24th Network and Distributed System Security Symp.* (NDSS 17), 2017; dx.doi.org/10.14722/ndss.2017.23034.
5. N. Karapanos et al., "Sound-Proof: Usable Two-Factor Authentication Based on Ambient Sound," *Proc. 24th USENIX Conf. Security Symp.* (SEC 15), 2015, pp. 483–498.
6. E. Wang and R. Chow, "What Can I Do Here? IoT Service Discovery in Smart Cities," *Proc. IEEE Int'l Conf. Pervasive Computing and Communication Workshops* (PerCom 16), 2016, pp. 1–6; dx.doi.org/10.1109/PERCOMW.2016.7457097.
7. "Personalized Privacy Assistant Project," Carnegie Mellon University, 2017; www.privacyassistant.org.

Richard Chow is a university research manager and scientist at Intel Corporation. Contact him at richard.chow@intel.com.

This article originally appeared in IEEE Security & Privacy, *vol. 15, no. 6, 2017.*

5G in E-Commerce Activities

Nir Kshetri
University of North Carolina at Greensboro

Editor:
Nir Kshetri, University of North Carolina at Greensboro;
nbkshetr@uncg.edu

5G cellular services are expected to be commercially available in developed countries later this year. These services, on their own or in combination with other technologies, are likely to have a transformative impact on e-commerce activities. The development of 5G networks, platforms, and devices certainly require high investments. There are, however, important economic benefits as well as psychological or intangible benefits associated with 5G. The faster speed of 5G networks and high-resolution screens of 5G-enabled devices might lead to a higher degree of customer willingness to engage in e-commerce activities, more time spent on e-commerce websites, and more purchases online. The features of 5G can also lead to a higher degree of effectiveness of e-commerce vendors' activities such as online advertising. Finally, faster speeds and higher-resolution screens are also associated with a higher degree of enjoyment (psychological or intangible benefits) when consumers engage in e-commerce activities.

Consider online video ads, which are a key component of an e-commerce ecosystem. In 2017, social video ad spending in the US was estimated at over $4 billion dollars or 20 percent of the total social media ad sales.[1] Most video ad tags contain sophisticated tracking codes, making it possible for advertisers to track users' interaction with the ads.[2] Due to their larger file sizes, loading video ads is more difficult and time consuming than loading text or photos. This is a major concern because many viewers strongly dislike ads and other content that do not download quickly. In a survey conducted by Adobe in December 2017, about 80 percent of respondents said that if content takes too long to download, they will stop viewing it or switch to a different device. Likewise, another study conducted in 2017 by video analytics firm Mux found that 85 percent of respondents would stop watching a video if it takes too long to load.[2]

5G offers great promise and potential to address these challenges. With 5G, data transmission and processing speeds will rocket to new levels. 5G will thus help create and deliver effective online video advertising that can attract customers' attention and produce the best results. Likewise, 5G networks ensure that devices do not lose Internet connections when traveling from one location to another.

5G-RELATED PROGRESS SO FAR AND FUTURE PROSPECTS

The 5G technology standard was finalized in December 2017, and saw its first deployment in the 2018 Winter Olympics in South Korea. An estimate by the cellphone trade group GSMA suggests that there will be 1.2 billion 5G connections worldwide by 2025.[3]

Asia and the US are expected to take the lead in 5G, whereas Europe is reported to be lagging.[4] The first commercial 5G projects are expected to be launched in the US in 2018.[5] Japan and South Korea are expected to launch 5G in 2019 and China in 2020. South Korean cellular company KT's peak 5G network speeds were up to 3.5 gigabits per second (GBps) on Samsung tablets.[6] It is worth noting that, in January 2018, South Korea's average mobile broadband download speed (which was the fourth highest in the world) was 133.05 megabits per second (Mbps). This means that peak speeds in 5G networks during the 2018 Winter Olympics were more than 26 times the country's mobile broadband download speed.

During the 2018 Olympics, KT provided 5G networks, Samsung provided 5G tablets, and Intel provided 5G platforms.[7] Samsung's 5G-equipped tablets allowed viewers to switch among multiple cameras placed along the cross-country skiing route so they could track an athlete's real-time location on a 3D map of the entire skiing course. Many other companies are developing products and services that would help enrich the global 5G ecosystem. China-based tech companies ZTE and Huawei announced that they would release 5G smartphones by early 2019.

COMBINING 5G WITH OTHER TECHNOLOGIES TO STIMULATE E-COMMERCE

5G—in combination with other technologies such as artificial intelligence (AI), Internet of Things (IoT), blockchain, augmented reality (AR), and virtual reality (VR)—is likely to be a transformative force in the e-commerce industry and market. Table 1 highlights some examples and benefits of combining 5G with other technologies to enhance e-commerce.

Table 1. Combining 5G with other technologies to enhance e-commerce.

Technology	Examples of uses	Benefits of combining with 5G
Internet of Things (IoT)	Improve consumer experiences, track inventory in real time, manage orders more effectively	Will make it easier to transfer data created by IoT devices
Artificial intelligence (AI)	Order products online, track orders, and perform other e-commerce activities	Will make it possible to access additional information quicker and understand the environment and context better
Blockchain	Smart contracts can be used by online vendors to automate order fulfillment; supply-chain management; B2B e-commerce	Will address security issues and feed the information (for example, from IoT devices) required for a smart contract more efficiently
Augmented reality (AR) and virtual reality (VR)	AR-enabled apps allow a potential customer to virtually place real products in a real setting to provide a clear visualization of the products' use.	5G networks have higher bandwidth, reduced latency, and a higher degree of uniformity to deal with the complex worlds and sophisticated inputs that require processing huge amounts of data.

Internet of Things

A study found that 70 percent of retailers worldwide were ready to adopt IoT to improve consumer experiences.[8] There are a number of ways in which e-commerce activities can benefit from IoT. For instance, IoT makes it easy to track inventory in real time and manage it more effectively. By doing so, human errors can be reduced. IoT can also help minimize waste, control costs, and reduce shortage. For instance, temperature-monitoring sensors can be used to maintain optimal temperatures for perishable products and send alerts when certain conditions are met.

Unsurprisingly, e-commerce companies have made heavy investments in IoT. China-based e-retailer JD.com's 3-System Fridge has sensors on every shelf and an internal camera. It registers the time and date when items are stored inside. The data is fed to a smart screen on the fridge's front side, which sends an alert when an expiration date is near. It can also order the next grocery list from JD.com based on the fridge's contents.[9]

Vast amounts of unstructured data are created by IoT devices, and the amount of data created is growing twice as fast as the available bandwidth.[10] It is estimated that by 2020, a network capacity that is at least 1,000 times the level of 2016 will be needed.[11] The amount of communication that needs to be handled will also increase costs exponentially. The current 4G networks are not capable of handling this growth. 5G networks can play a role in addressing the potential bandwidth deficit.

Virtual and Augmented Reality

VR and AR are likely to emerge as driving forces in the e-commerce industry and market. By wearing a VR headset, a shopper can instantly find herself in a company's virtual shop, where she can "walk" around to explore items exactly as she would in the real shop. For instance, if she wants to know more about a new piece of jewelry in the shop, she can focus her sight on that item and see the relevant information needed to make a purchasing decision. If she wants to buy it, she can make the payment or add it to her cart and look for additional items.[12]

AR applies VR in the real world with live video imagery.[13] For example, a furniture vendor can develop an AR-enabled app that allows a customer to point the mobile camera to the place where she wants to keep the furniture. The app places a 3D model of the furniture as an overlay on the living room's live image. The shopper now has a clear visualization of how the furniture would fit in her living room. In addition, by rotating the camera, she can see where the new furniture fits best.[12] This has already been done by some companies such as Lego and IKEA.[14] Likewise, the home furnishings and decor company Wayfair launched its AR app WayfairView in 2016, which is available in Google Play.

The complexity and richness of the AR and VR worlds require processing a large quantity of data. Current 4G networks standards suffer from some limitations such as those related to bandwidth, latency, and uniformity, especially when the data needs to be fed remotely. In this regard, 5G is likely to unlock the full potential of VR and AR technologies. 5G's significantly faster speeds and lower latency would help overcome these weaknesses.[15] 5G streams' transmission delay is about 1 millisecond, which is much shorter than human beings can notice. 5G's almost zero delay in transmission is thus likely to enrich customers' experiences with AR and VR technologies.

According to TimeTrade research, 85 percent of consumers prefer shopping in physical stores. Many features of brick and mortar stores cannot be replaced by e-commerce.[16] However, AR and VR can create consumer experiences that are close enough to brick and mortar stores.

Retailers can also use AR and VR in concept testing to measure consumers' responses and feedback before implementing their ideas. Such testing can provide retailers with insights into consumers' emotional connection with the products, which can result in consumer adoption and market impact.[17]

Blockchain

Sophisticated applications of blockchain to facilitate e-commerce activities have been or are being developed. JD.com has implemented blockchain in its supply-chain management system and B2B e-commerce. In 2017, the system went live with beef manufacturer Kerchin as its first supply-chain partner. The company announced that JD.com would have more than 10 brands of alcohol, food, tea, and pharmaceutical products on its blockchain by the end of October 2017.[18]

One of the most high-profile future uses of blockchain is likely to be smart contracts. Online vendors can use smart contracts to automate fulfillment of orders for the delivery of digital products. In smart contracts that are executed “above” the blockchain, 5G can play a key role in feeding the information (for example, from IoT devices) more efficiently.

Cyberattacks have been a main concern for the growth of e-commerce. Such concerns are further heightened by the rapid growth of IoT networks, as IoT sensors carry sensitive information about their users. Securing these systems is thus important. Another key use of blockchain would be addressing security issues in 5G. Telecom companies can provide an eSIM (embedded SIM) or an app to a subscriber that creates a unique virtual identity that is encrypted and stored on a blockchain. Subscribers can use the unique IDs for automatic authentication on e-commerce websites. Blockchain also allows for secure P2P network solutions.

AI

AI-enabled devices are already playing important roles in helping consumers in e-commerce activities such as making buying decisions and tracking products. For instance, virtual assistants are transforming the way consumers purchase products online. Amazon’s personal assistant Alexa, for example, has been integrated into Amazon products as well as those from other manufacturers. Customers can use Alexa to find information about local concerts through eBay’s online ticket exchange company StubHub. In addition, they can arrange transportation to the event via Uber and order pre-event dinner from Domino’s. The order status can be tracked in real time.[19] Likewise, consumers can order flower arrangements with 1-800-Flowers.com through Alexa or Google assistant-enabled devices.[20] AI-enabled devices have already started handling unstructured information. For instance, social media platforms use AI to enable facial recognition and photo tagging.[21]

5G will dramatically improve consumer experiences with AI-based devices. With 5G, AI-based devices can access additional structured and unstructured information quicker and understand the environment and context better. Overall, AI-powered services will be more reliable in a wide range of contexts and situations in which they operate.

CONCLUSION

Some of the current challenges in the development of e-commerce can be overcome with the deployment of 5G networks, such as with IoT devices. Addressing the exponential growth in IoT devices will be no small feat. Current 4G networks, however, cannot handle all the data coming from IoT devices. 5G’s higher data transmission and processing speeds will address this concern. Specifically, 5G (in combination with AI, VR, AR, and other technologies) will play a powerful role in transforming the e-commerce industry and market. Such a combination can result in a rich e-commerce ecosystem and a better customer experience.

REFERENCES

1. “Social Video Ad Spend Will Double in 2017,” *BI Intelligence*, 3 April 2017; http://uk.businessinsider.com/social-video-ad-spend-will-double-in-2017-2017-4?r=US&IR=T.
2. R. Benes, “Many Users Will Stop Viewing Content That’s Slow to Load,” *eMarketer*, 9 February 2018; www.emarketer.com/content/half-of-u-s-users-stop-viewing-content-slow-to-load.
3. “10 Things We Learned at Mobile World Congress,” *Rakuten.Today*, 2 March 2018; https://rakuten.today/blog/10-things-learned-mobile-world-congress-2018.html.
4. A. Kharpal, “China 'Has the Edge' in the War for 5G and the US and Europe Could Fall Behind,” *CNBC*, 7 March 2018; www.cnbc.com/2018/03/07/china-has-the-edge-in-the-war-for-5g-us-and-eu-could-fall-behind.html.

5. E. Auchard, "Faster 5G Connections Still Far Off for Most Mobile Users," *CRN*, 4 March 2018; www.crn.com.au/news/faster-5g-connections-still-far-off-for-most-mobile-users-486257.
6. M. Dano, "KT's Millimeter Wave 5G Network Transmitted 3,800 TB of Data during Winter Olympics," *FierceWireless*, 25 February 2018; www.fiercewireless.com/5g/kt-s-millimeter-wave-5g-network-transmitted-3-800-tb-data-during-winter-olympics.
7. E. Ramirez, "In the Race for 5G, South Korea Shows Off its Tech Prowess at the Winter Olympics," *Forbes*, 23 February 2018; www.forbes.com/sites/elaineramirez/2018/02/23/in-the-race-for-5g-south-korea-shows-off-its-tech-prowess-at-the-winter-olympics/#4c1384762853.
8. N. Dave, "4 Major Ways in Which e-Commerce is Benefitting from IoT," *Business 2 Community*, 15 November 2017; www.business2community.com/ecommerce/4-major-ways-e-commerce-benefitting-iot-01957768.
9. E. Lye, "Five Chinese Brands That Will Be Household Names Next Year," *The Drum*, 14 June 2016; www.thedrum.com/opinion/2016/06/14/five-chinese-brands-will-be-household-names-next-year.
10. N. Desai, "IBM Scientists Team with the Weather Company to Bring Edge Computing to Life," *IBM Research*, 15 February 2017; www.ibm.com/blogs/research/2017/02/bringing-edge-computing-to-life.
11. S. Waterman, "Industry to Government: Hands Off IoT Security," *FedScoop*, 17 November 2016; www.fedscoop.com/industry-to-government-hands-off-iot-security.
12. "Retail e-commerce: How AR and VR Can Transform Shopping," *CIO*, 2018; www.cio.in/feature/retail-e-commerce-how-ar-and-vr-can-transform-shopping-0.
13. T. Perdue, "Applications of Augmented Reality," *Lifewire*, 7 July 2017; www.lifewire.com/applications-of-augmented-reality-2495561.
14. "11 Creative Uses of Augmented Reality in Marketing and Advertising," *Forbes*, 13 June 2017; www.forbes.com/sites/forbesagencycouncil/2017/06/13/11-creative-uses-of-augmented-reality-in-marketing-and-advertising/#268a5a3b6b12.
15. J. Mundy, "How Will 5G Boost VR and AR?," *5G.co.uk*, 2018; https://5g.co.uk/guides/5g-virtual-reality-and-augmented-reality.
16. T. Bragg, "Is IoT Technology Transforming Brick and Mortar Retail?," *TechWire Asia*, 13 November 2017; http://techwireasia.com/2017/11/iot-technology-transforming-brick-mortar-retail.
17. S. Tamang, "Will 5G Truly Transform Our Lives?," *Enterprise Innovation*, 30 May 2018; www.enterpriseinnovation.net/article/will-5g-truly-transform-our-lives-1244626393.
18. E. Xiao, "Alibaba, JD Tackle China's Fake Goods Problem with Blockchain," *Tech in Asia*, 16 October 2017; www.techinasia.com/alibaba-jd-ecommerce-giants-fight-fake-goods-blockchain.
19. D. Asling, "19 Powerful Ways to Use Artificial Intelligence in Ecommerce," *Linnworks Blog*, 2018; http://blog.linnworks.com/artificial-intelligence-in-ecommerce.
20. J. Renfrow, "1-800-Flowers.com Offers AI-Powered Shopping for Valentine's Day," *FierceRetail*, 1 February 2018; www.fierceretail.com/digital/1-800-flowers-adds-ai-powered-shopping-for-valentine-s-day.
21. T. Harbet, "How 5G Will Make AI-Powered Devices Smarter," *IQ Intel*, 21 February 2018; https://iq.intel.com/how-5g-will-make-ai-powered-devices-smarter.

ABOUT THE AUTHOR

Nir Kshetri is a professor of management at the Bryan School of Business and Economics at the University of North Carolina at Greensboro. Contact him at nbkshetr@uncg.edu.

This article originally appeared in IT Professional, *vol. 20, no. 4, 2018.*

A New Era for Web AR with Mobile Edge Computing

Xiuquan Qiao
State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications

Pei Ren
Beijing University of Posts and Telecommunications

Schahram Dustdar
TU Wien

Junliang Chen
Beijing University of Posts and Telecommunications

Dedicated device-based and app-based augmented reality (AR) solutions have inherent limitations regarding cross-platform, pervasive AR application provisioning. *Web-based AR* (web AR), a promising lightweight and cross-platform approach to AR, is gaining increasing attention owing to its extensive application areas. However, for computationally intensive AR applications, the weak computational efficiency of current web browsers seriously hampers applications of web AR on a large scale. The "browser + cloud" approach suffers from the high-latency dilemma. Now, with the emerging 5G networks, *mobile edge computing* (MEC) promises to greatly reduce network latency (even to 1 ms) by deploying applications at the network edge closer to users, which provides an opportunity for performance improvement of web AR. In this article, the authors envision that the application of MEC has the potential to bring web AR into a new era. Specifically, an MEC-oriented web AR solution is first proposed, followed by the design and deployment details. The authors also discuss future directions aimed at using MEC to tackle the performance issues of web AR in 3G, 4G, and 5G networks.

Mobile augmented reality (MAR)[1] nowadays greatly attracts both academic and industrial attention, owing to its extensive application areas and the popularity of smartphones and wearable devices. Recently, many MAR use cases have already emerged in the fields of art, education, entertainment, and so on. These use cases typically are divided into two types: wearable-device-based and app-based MAR solutions. However, both approaches suffer from some inherent limitations regarding cross-platform, portability, and cost aspects:

 Published by the IEEE Computer Society

- Wearable-device-based MAR solutions are costly and inconvenient to carry.
- App-based MAR solutions require downloading and installing the specific app in advance and lack cross-platform support (i.e., an AR activity of one app cannot be used in other apps directly).

Therefore, it's necessary to explore a lightweight, low-cost, and cross-platform pervasive MAR solution for mobile users. Owing to web-based AR's cross-platform and no-installation advantages, it is attracting more and more attention and is opening up a new research field of MAR.

Currently, the proposed web AR implementations can be classified primarily as pure front-end solutions using JavaScript at a mobile Web browser, browser-kernel-based extension solutions, or "browser + cloud" solutions. However, all these web AR solutions are also struggling with some problems:

- *Pure front-end solutions*. MAR applications often rely on computationally intensive computer vision algorithms. However, JavaScript cannot provide efficient computational capability for complex matrix computations. Therefore, pure front-end solutions can only use several typical algorithms to recognize some simple markers, and have poor recognition and tracking ability for nature images and real material objects.
- *Browser-kernel-based extension solutions*. To more fully utilize the computational resources of smartphones than the pure front-end solutions, some web browser providers try to extend the browser kernel to enable web AR functions, even though this approach still faces the challenges of cross-platform requirements owing to the lack of web AR standardization between different web browsers. In addition, if the application scenario needs more computing capacity (for example, to recognize more pictures), the computing capacity of end devices is still limited.
- *Browser + cloud solutions*. Considering that MAR applications often involve complex matrix computations with extreme latency requirements, the local computing operations are often limited by hardware capabilities (e.g., CPU, GPU, memory, and battery). To address this issue, cloud computing is currently used to extend the computing ability of end devices. However, this approach also brings about significant communication latency, which often incurs poor quality of experience (QoE), such as in real-time object-tracking scenarios.

On the basis of the above analysis, we need to explore a practical web AR implementation solution to resolve both the computational inefficiency of the pure front-end and browser-kernel-based solutions and the high-latency of browser + cloud solutions. Fortunately, with the emergence of 5G networks, the new *mobile edge computing* (MEC) paradigm[2] aims to provide cloud-computing capabilities at the edge of networks close to mobile users. Thus, MEC can efficiently reduce the communication latency and strengthen the end device's computing capability. We argue that MEC will provide an opportunity for the performance improvement of web AR and open up a new era for web AR applications.

In this article, we first give an overview of the state of the art of web-based MAR solutions and discuss the problems to be addressed. Then, a practical web AR service-provisioning framework leveraging the MEC paradigm is presented, which achieves encouraging performance improvement. Finally, we discuss future research directions of web AR with MEC in current 3G and 4G networks and future 5G networks.

THE STATE OF THE ART

Pure Front-End Solutions

In general, pure front-end web AR solutions utilize different JavaScript libraries to support web AR applications, which can be adopted simply in users' end devices.

AR.js[3] is one of the popular marker-based MAR solutions for efficiently performing AR applications on the Web, including recognition, tracking, and 3D object rendering on any mobile browser with WebRTC and WebGL. JSARToolKit (https://github.com/kig/JSARToolKit) is an AR library written in JavaScript, which allows rendering a 3D model of a detected marker inside a camera feed by computing the distance from the camera to the physical marker in real time. CaffeJS (https://github.com/chaosmail/caffejs) aims to run neural-network functions (e.g., image recognition) by porting Caffe models to a web browser using a modified version of ConvNetJS, which saves a lot of network traffic and server resources.

Browser-Kernel-Based Extension Solutions

Another branch of web AR solutions is browser-kernel-based extension solutions, where the functions of AR applications have an extreme dependency on browsers. Currently, most browser-kernel-based web AR projects are still in closed beta stages.

Google is developing WebARonARKit and WebARonARCore, which are experimental apps for iOS and Android that let developers create AR experiences using web technologies. Mozilla has announced its mixed-reality program: Mozilla WebXR (https://github.com/mozilla/webxr-api), which aims to make it easy for web developers to create web applications that adapt to the capabilities of each platform. The Baidu mobile-browser team has also shared its advancements in web AR and web VR in a recent open class. The Baidu mobile-browser T7 kernel supports the WebVR standard and has completed web AR prototype development (https://github.com/baidu/AR). Argon,[4] an AR-supported web browser, provides a pervasive web AR application deployment platform. Moreover, Argon implements the JavaScript library argon.js, which aims to support web AR applications in any web browser.

MOTIVATION AND PROBLEM STATEMENT

The aforementioned web AR implementations all face inherent limitations. The pure front-end solutions lack computational efficiency, and the browser-kernel-based extension solutions are still in their infancy and have not yet been adopted on a large scale. The cloud-computing paradigm greatly extends the computing ability of the terminal devices and provides an opportunity for the promotion of web AR applications on users' end devices with limited hardware resources.

Benefiting from the cloud-computing paradigm, we have conducted several online real advertising promotion activities using web AR within WeChat, which have gained a great deal of attention. According to our practical project experience, nowadays most users' end devices such as smartphones and tablet PCs have satisfied the deployment requirements of web AR. Technologies that have enabled this situation include

- WebAssembly, a new portable, size- and load-time-efficient format suitable for compilation to the web;
- Web Workers, an HTML5 technique that provides a means to run a script operation in a background thread separate from the main execution thread of a web application; and
- WebRTC, a collection of communications protocols and APIs that enable real-time communication over peer-to-peer connections.

Moreover, all devices equipped with iOS 11 or later versions support WebRTC. All these enabling techniques provide a platform for web AR applications' deployment, enabling the lightweight, cross-platform and pervasive application of web AR.

Here, we detail two influential use cases as shown in Figure 1, describing the challenges we faced during design and development.

Figure 1. Two real use cases of web augmented reality (web AR) for market promotions of China Mobile and FenJiu.

- *China Mobile use case*. A web AR sales promotion for the CMCC A3s smartphone has been proposed for advertising effects. The activity URL link is embedded in China Mobile's official account in WeChat. Users only need to access the activity item in the corresponding official account, by scanning the logo of China Mobile; an augmented 3D model then is rendered on a webpage. In this application scenario, the most important and complex part is image matching. However, the available image-matching algorithms such as SURF (Speeded-Up Robust Features), SIFT (Scale-Invariant Feature Transform), and ASIFT (Affine-SIFT) are too heavy to run on web browsers. The ORB (Oriented FAST and Rotated BRIEF) matching algorithm provides a faster feature extraction method and can be used directly by JavaScript; however, it results in insufficient matching accuracy.
- *FenJiu use case*. To promote the company culture, FenJiu has adopted a web AR technique for its museum. As in the use case mentioned above, users can experience the web AR application by scanning different material objects such as wine bottles and other exhibits in the museum. However, once the image is captured by a user's end device, it needs to be compared with 10 other samples (the number of samples will increase in the future), the results of which are used to identify different exhibits and then present different cultural-promotion pages. Such intensive computing apparently cannot be performed by users' end devices locally. Benefiting from the cloud-computing paradigm, computation tasks, therefore, can be migrated to cloud servers. However, the cloud-computing approach introduces high latency due to long-distance data transmission through the wireless network, which significantly degrades the performance of the web AR effects.

These two use cases placed us in a dilemma:

- MAR applications rely on complex computer vision algorithms, while web browsers always lack sufficient computational capability.
- Although the cloud-computing paradigm extends the computational and storage capability for users' end devices, it causes significant latency in the mobile network. However, MAR applications require low latency.

WEB AR WITH MOBILE EDGE COMPUTING

By offloading computationally intensive workloads to the proximity servers at the edge of the network, the emerging edge-computing paradigm provides a promising solution to meet the ever-increasing computational demands of numerous mobile applications.[5] To compensate for the lack of computing capability and to ease the new constraints caused by cloud computing, especially concerning latency and bandwidth, we propose a web AR service-provisioning framework with MEC, on the basis of our project experiences.

Overview of the MEC Framework for Web AR

The overall framework is first proposed in three perspectives as shown in Figure 2, followed by the details of three processing components on the terminal, the edge cloud server, and the remote cloud server. We conducted the experiments on the basis of our practical projects, and the results fully demonstrate the advantages of MEC for web AR applications.

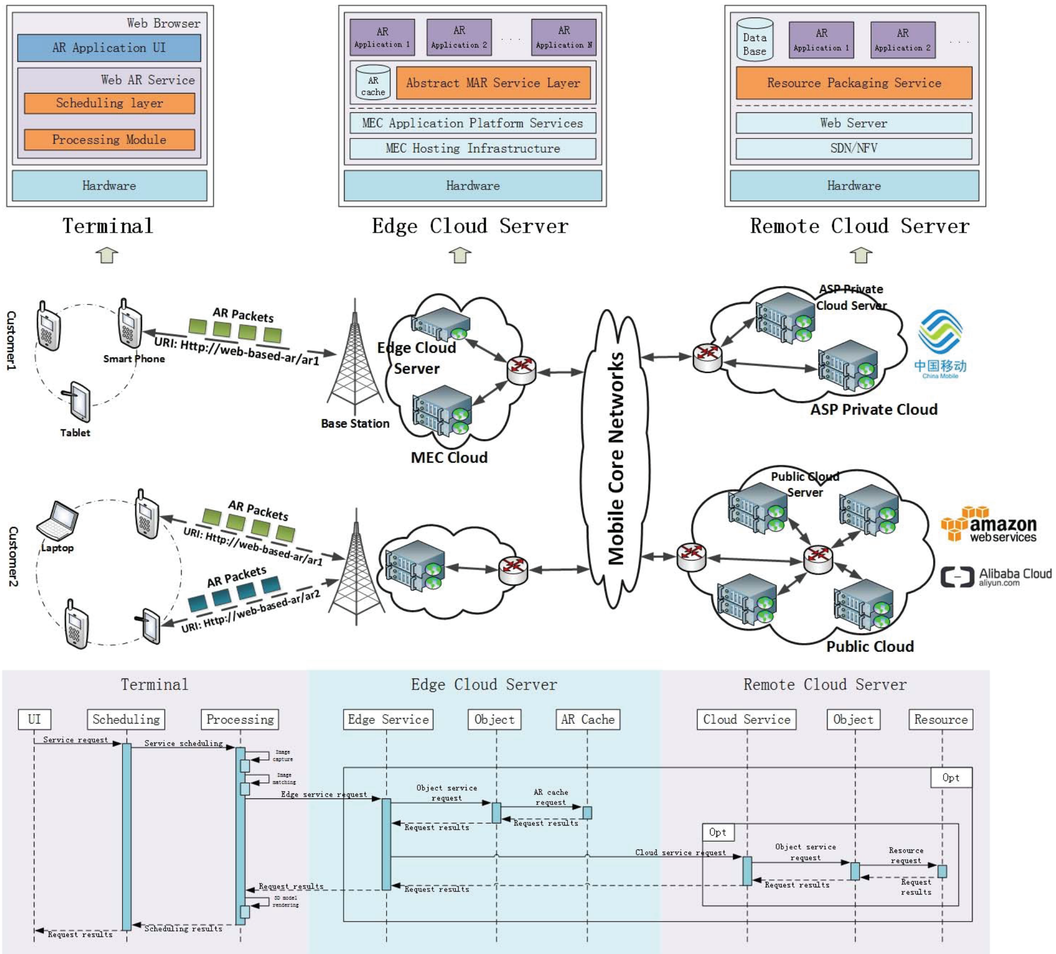

Figure 2. Overview of the framework for web AR with mobile edge computing (MEC).

The Terminal Side

The terminal-side computation platform is responsible mainly for web AR service scheduling and basic processing, which will not pose a severe computational burden, considering the weak computational capability of the web browsers in users' end devices.

The web AR service platform consists mainly of two modules: the *scheduling layer* and *processing module*. The processing module provides the underlying support for the scheduling layer. The scheduling layer processes all the web AR application logic and collaborates with the related service parts.

The processing module in the web AR service layer is composed of three submodules: the image-capture, image-matching, and 3D-model-rendering submodules. The image-capture submodule uses a WebRTC technique to capture an image from a camera and then execute a resizing operation (set by the application service provider in advance) on it, taking into account the communication cost of image transmission. The image-matching submodule performs an image-to-image matching operation, leveraging the lightweight image-matching algorithms such as ORB mentioned above, which is already supported in JSFeat (the JavaScript Computer Vision Library). The 3D-model-rendering submodule leverages a WebGL technique to perform 3D model rendering. Users can interact with the 3D models to acquire more augmented information, which provides a more attractive and friendlier experience compared to traditional ways.

However, for some complex images, the proposed ORB matching algorithm cannot work well. Then, the specific edge cloud services will be invoked in the case where the image-matching result acquired on the terminal side is insufficient for the web AR application.

The Edge Cloud Side

The edge-cloud-side computation platform consists mainly of an *abstract MAR service layer*, which is used to process incoming edge service requests and manage the web AR application objects, including object deployment, object destruction, and other underlying service supports.

The underlying abstract MAR service layer consists of some common web AR functional modules, which aim to provide simpler and faster service to the web AR application instances in the upper layer and reduce the access cost to hardware at the same time, so as to improve the edge server's overall performance. Currently, it includes a performance-monitoring module and image-matching module.

Once the edge cloud server receives a web AR request (i.e., an image-matching request) from the terminal side, it will directly forward that request to the specific web AR application instance. The AR cache in the edge cloud server can be accessed by all web AR applications. However, in the case where the requested web AR application has not been deployed yet, the abstract MAR service layer will send the application deployment request to the remote cloud server. The current performance of the edge server, including the CPU, memory, and storage usage, is also sent to the remote cloud server for the decision making of the web AR application deployment. The remote cloud server—i.e., the application service provider (ASP)—finally determines where the specific application is to be deployed, taking into account the overall cost of deployment and transmission.

After the abstract MAR service layer in the edge server receives the requested application, it will perform deployment. Thus, it provides more efficient servicing of the follow-up web AR requests. Moreover, the processing module in the terminal performs the rendering operation on the specific 3D model according to the received matching result.

The Remote Cloud Side

The remote-cloud-side computation platform aims at providing a more generalized service provision mechanism. One of the most important components in the framework is the *resource packaging service* layer, which acts as a web AR resource manager. When the ASP faces different web AR application requirements, the resource-packaging-service layer combines specific web AR resource components into different web AR applications, which will then be deployed to appropriate edge cloud servers according to the deployment decisions. Meanwhile, different ASPs also have different web AR resources, including various 3D models, image-matching algorithms, etc. The resource-packaging-service layer is also responsible for the management of these contents.

The database is used mainly for the user information store. Moreover, images that match successfully on the edge cloud server will also be transmitted to the remote cloud server, since these types of images are currently not handled by the terminal-side image-matching algorithm. These images can be used to improve the performance of the terminal-side image-matching algorithm in the future.

Performance Evaluation

Theoretically, the introduction of the edge server will bring the following advantages:

- In scenarios that need to match a single image, the matching accuracy can be improved by providing a more optimized algorithm on the edge server side. The edge server has greater computational capability, and the image-matching time can be reduced effectively compared to the same matching algorithm on the terminal side. Although uploading images to the edge server adds extra time, the total processing efficiency still improves. Moreover, the edge server performs better since it is closer to the users. Hence, the network delay is no longer the bottleneck for web AR applications when the MEC paradigm is adopted.
- In the case where multiple images need to be compared, users' end devices apparently cannot meet the performance requirement of the web AR application owing to their limited computational capability. Migrating the service module (i.e., the web AR application) from the remote cloud server to the edge server effectively eases the issue of network delay caused by the frequent image transmissions.

To get performance comparisons, we conducted experiments on a Samsung Note 4 mobile phone, using Chrome under the use cases mentioned above. The mobile phone connected to a Wi-Fi access point on our campus (i.e., BUPT-mobile). By modifying the DNS setting, all the requests were redirected to the server in our lab, which played the role of the edge server. The MEC framework and two web AR applications were deployed in our lab and on Alibaba Cloud, respectively. A modified version of the ASIFT algorithm was employed to provide more accurate and effective image matching. Moreover, the feature points of image samples on the server were computed in advance to achieve faster matching. Four performance metrics are considered: frames per second (fps), latency, the terminal's power consumption, and the average number of matched feature points.

In experiments, the image was first resized to 250 * 250 pixels, about 8.92 Kbytes on average. The round-trip time for terminal-edge and terminal-cloud was 13.787 ms and 39.781 ms, respectively, for 64 bytes. Each ORB image-matching operation cost 394 ms on a user's end device, while only about 30 ms on the server. Hence, offloading computationally intensive tasks to the edge cloud was still a good choice. Meanwhile, edge computing performed better than the cloud paradigm regarding latency and power consumption, as shown in Figure 3, since it benefited from a shorter image transmission distance.

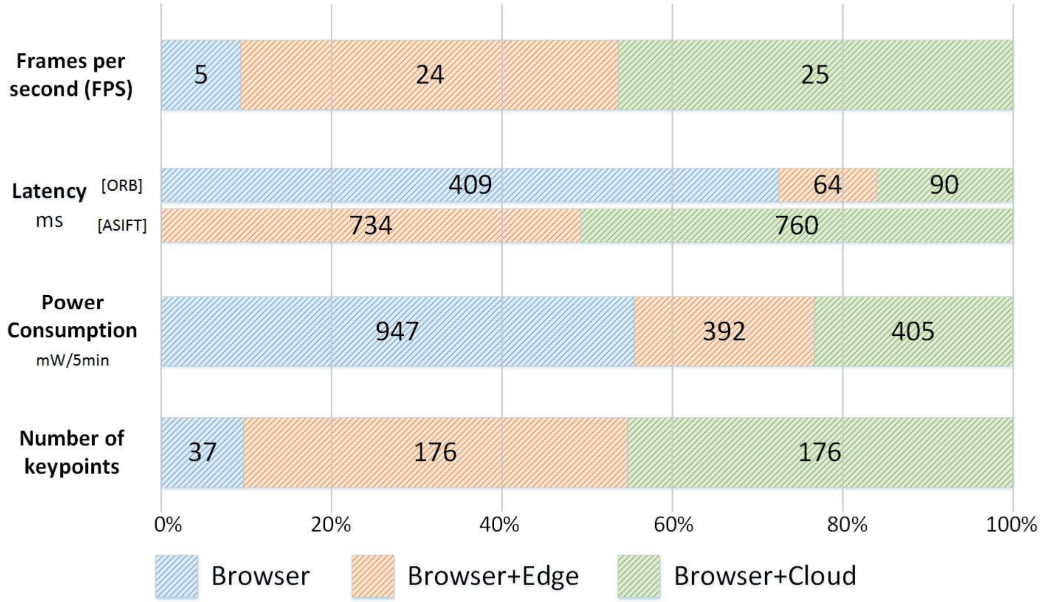

Figure 3. Experiment results in terms of frames per second, latency, the terminal's power consumption, and the average number of matched keypoints.

The experiment results show that MEC indeed provides a new choice for the future deployment of MAR.

FUTURE DIRECTIONS

The practical projects and performance results have given us great inspiration for the future development and deployment of web AR applications. Next, we discuss some important research challenges and directions for web AR and MEC.

Future Directions for Web AR

Web AR is becoming more and more popular owing to its lightweight and cross-platform features. We argue that it will provide a promising future for MAR. However, there is still a lot of work to be done to further promote the application of web AR.

Computational Efficiency for CPU-Hungry Tasks on Mobile Web Browsers

The widely used JavaScript performs poorly for complex computation tasks, such as matrix computation and floating-point computation. Therefore, it is necessary to introduce a more efficient computation paradigm to web browsers to meet the computational-efficiency requirement. Although some advanced techniques such as WebAssembly and Web Workers are already supported in some mainstream web browsers such as Chrome, Firefox, and Safari, there is still room for improvement to achieve efficient computation for end devices.

Lack of Standardization

Browser-kernel-based extension solutions for web AR implementations can make full use of end devices' hardware resources to achieve better performance. We envision that this has more potential as a web AR solution. However, a variety of current browser-kernel-based solutions have serious compatibility problems. Different web AR applications can be used only by their dedicated browser, which significantly limits the popularization of web AR. In addition, during the development of web AR applications, we found that there also exist serious compatibility issues between Web3D rendering technologies (such as three.js) and Web3D models made by different modeling tools such as 3ds Max, Maya, and Blender. To promote web AR applications on a large scale, the issue of compatibility mentioned above will be one of the most critical problems to be solved in the future.

Network Constraints

MAR is heavily dependent on network latency and bandwidth. However, wireless networks have an adverse effect on the performance of web AR applications. Although the current 4G networks already have good performance capability, they still cannot meet the low-latency requirements of new applications such as AR and VR. Technologies such as software-defined networking (SDN), device-to-device (D2D) communication, and mobile crowdsourcing mechanisms[6] offer new ways for wireless-network resource optimization. However, there is still plenty of optimization room to further improve the performance of web AR applications.

Future Directions for MEC

Because MEC is a newly proposed computing paradigm, it offers many opportunities for research, development, and deployment. In this section, we discuss potential research directions for MEC under 3G, 4G, and 5G networks.

MEC in 3G and 4G Networks

MEC has been explored by many carriers, including China Mobile and Huawei. However, it is still at an initial stage. Under the current 3G and 4G networks, many problems and challenges regarding the application, development, and deployment of MEC still exist. For MAR applications, the severe delay issues under wireless networks greatly degrade their performance. Currently, the content delivery network (CDN) closest to users is used mainly to cache the static content resources and has no computing and application deployment functions. Therefore, it's necessary to upgrade the existing CDN as much as possible to support the edge-computing paradigm.

MEC in 5G Networks

5G mobile networks have recently gained momentum in both academia and industry. Benefiting from the inherent features of 5G cellular networks, such as enhanced mobile broadband (eMBB), massive machine type communications (mMTC), and ultra-reliable low-latency communications (URLLC), 5G is expected to enable a host of new applications, such as the Internet of Things (IoT), self-driving cars, AR and VR, and so on. Moreover, MEC is also available in 5G cellular networks to provide computing and storage abilities at the edge of a mobile network.[7]

Compared with current 3G and 4G networks, the issues of delay and bandwidth occupancy get alleviated in 5G networks. The introduction of MEC further optimizes the utilization of network resources and improves the performance of applications such as MAR, since the requests can be responded to faster in edge computing than with remote cloud-computing centers. However, when, where, and how to deploy MEC servers to fully balance the benefits, cost, and efficiency will be an essential issue in 5G networks.

CONCLUSION

MAR provides an attractive visual experience to users. Web AR further offers a lightweight, cross-platform, and pervasive solution on the web. However, it faces more technical challenges in terms of QoE owing to the weak computational efficiency of JavaScript. The emerging MEC paradigm provides an opportunity for the performance improvement of computationally intensive applications such as web AR. In this article, we proposed an MEC-oriented web AR solution. Then, we discussed challenges and future directions for MEC-based web AR under 3G, 4G, and 5G networks. We hope this article can provide insight on future research efforts in the emerging research field of web AR with MEC.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China under grants 61671081 and 61720106007, the Beijing Natural Science Foundation under grant 4172042, and the 111 project (B18008).

REFERENCES

1. M. Billinghurst, A. Clark, G. Lee et al., “A Survey of Augmented Reality,” Foundations and Trends R in Human–Computer Interaction, vol. 8, no. 2-3, 2015, pp. 73–272.
2. A. Ahmed and E. Ahmed, “A Survey on Mobile Edge Computing,” Proc. 10th IEEE International Conference on Intelligent Systems and Control (ISCO), 2016, pp. 1–8.
3. P. Mendigochea, “WebAR: Creating Augmented Reality Experiences on Smart Glasses and Mobile Device Browsers,” Proc. ACM SIGGRAPH Studio, 2017, pp. 10.
4. B. MacIntyre, A. Hill, H. Rouzati, M. Gandy, and B. Davidson, “The Argon AR Web Browser and Standards-based AR Application Environment,” Proc. 10th IEEE International Symposium on Mixed and Augmented Reality (ISMAR), 2011, pp. 65–74.
5. W. S. Shi, S. Dustdar, ”The Promise of Edge Computing, “IEEE Computer Magazine, Vol. 29, No. 5, 2016, pp. 78-81.
6. D. Zhao, X. Y. Li, H. D. Ma, ”How to Crowdsource Tasks Truthfully Without Sacrificing Utility: Online Incentive Mechanisms with Budget Constraint, “Proc. INFOCOM 2014, pp. 1213-1221.
7. T. X. Tran, A. Hajisami, P. Pandey, and D. Pompili, “Collaborative Mobile Edge Computing in 5G Networks: New Paradigms, Scenarios, and Challenges,” IEEE Communications Magazine, vol. 55, no. 4, 2017, pp. 54–61.

Xiuquan Qiao is a full professor in the State Key Laboratory of Networking and Switching Technology at the Beijing University of Posts and Telecommunications. His main research interests focus on 5G networks, augmented reality, and services computing. Contact him at qiaoxq@bupt.edu.cn.

Pei Ren is a PhD candidate from the Beijing University of Posts and Telecommunications. His research interests lie in machine learning and augmented reality. Contact him at renpei@bupt.edu.cn.

Schahram Dustdar is a full professor of computer science (informatics) with a focus on Internet technologies and heads the Distributed Systems Group at TU Wien. He’s a member of Academia Europaea. Contact him at dustdar@infosys.tuwien.ac.at.

Junliang Chen is a professor at the Beijing University of Posts and Telecommunications. His research interests include services computing. He’s a member of the Chinese Academy of Sciences and Chinese Academy of Engineering. Contact him at chjl@bupt.edu.cn.

This article originally appeared in
IEEE Internet Computing, *vol. 22, no. 4, 2018.*

360-Degree Virtual-Reality Cameras for the Masses

Jack Tan
Kandao

Gene Cheung
National Institute of Informatics, Tokyo

Rui Ma
Kandao

Department Editor:
Tao Mei, Microsoft Research Asia

It is human nature to tell stories and share life experiences with others. Methods of storytelling have evolved over time due in large part to technological advances in both communication and experience recording. Camera technologies have evolved to the point where it is now possible to record a 360-degree 3D visual experience using virtual-reality (VR) video cameras. For the first time, people are free of the constraint of fixed camera viewpoint capture, so they can record and share what they see in 360 degrees as they experience it in real life. Due to the unprecedented immersive experience that VR can potentially bring to the masses, investment in VR is growing fast. Facebook officially acquired Oculus VR at a price of $2 billion. In 2015, more than $700 million was injected into VR and AR startups, according to Digi-Capital's Augmented/Virtual Reality Report. In 2016, the report says, that number more than tripled to $2.3 billion (digi-capital.com/news/2017/02/record-2-3-billion-vrar-investment-in-2016).

But if the technology is so promising, why aren't VR videos already ubiquitous? There are several major factors preventing the technology from being widely deployed. A good-quality 3D, 360-degree camera is still too expensive for mass adoption, costing as much as $20,000. Usability is also an issue. For the masses to adopt VR cameras, it is important that users be able to shoot anywhere and anytime. Many of the systems currently on the market are too bulky to be moved easily. They often have to rely on external power sources due to the massive amount of energy consumed by the camera arrays, which hinders portability. For PC-based solutions, the rendering process is often very slow. For example, a solution from a major VR camera vendor takes 12 seconds to render one frame of a 4k x 4k high-quality stereoscopic video. Another challenge is that the increased data dimensionality of VR poses technical challenges for video compression and delivery.

All of the aforementioned factors make the technology challenging to deploy on a large scale. To make the technology more accessible to the public, a device must be affordable, portable, reliable, high quality, and user friendly. In this article, we describe the challenges in meeting these goals and the techniques that Kandao—a VR startup company based in China—used to conquer them when designing its Obsidian cameras.

 2469-7087/19/$33.00 © 2019 IEEE

CAPTURING STEREOSCOPIC PANORAMAS

Capturing a good-quality stereoscopic, 360-degree panorama is a computationally challenging task. The omni-directional stereo (ODS) projection model was initially proposed in 1992 (shown in Figure 1).[1,2] In the original ODS, a slit camera rotates around an inter-pupil-distance (IPD) circle and scans the real world. The image slices are stacked to form a rectangular panoramic image. An inherent drawback of this and other early methods is that they do not work on dynamic scenes with moving objects.

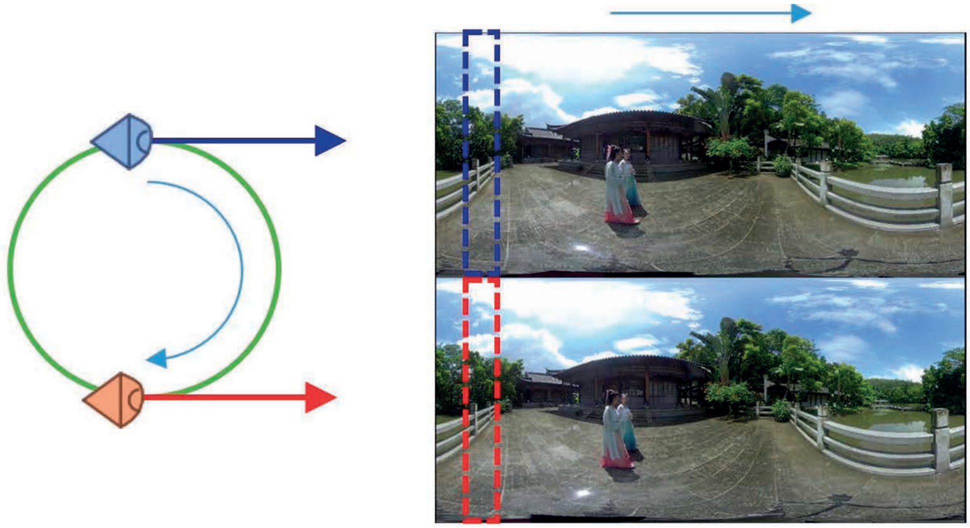

Figure 1. The original ODS uses rotating slit cameras for the left and right eyes.

Recent updates to ODS take a computational photography approach.[3] Instead of actual scanning, the rotating slit camera is synthesized using computational imaging and vision techniques. In this process, an accurate dense pixel-to-pixel correspondence is calculated, and then the depth of each pixel is estimated from correspondence (for example, a′ is the correspondence of a in Figure 1, both of which are images of point A). Each 3D point in the space is then projected to the two slit cameras tangent to the IPD circle (shown in green in Figure 2).

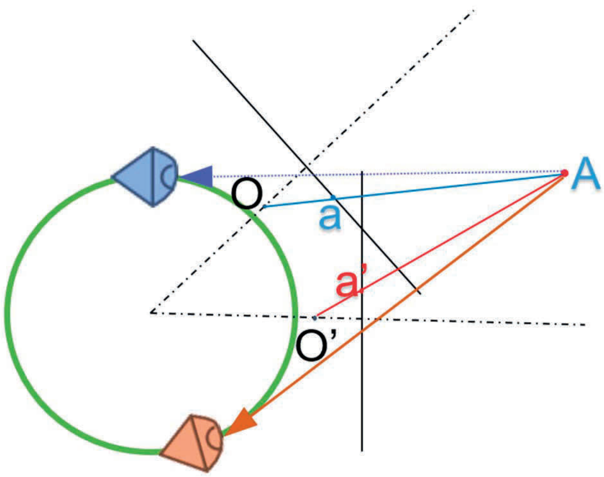

Figure 2. A modern ODS illustration. Point A is captured by two cameras centered at O and O' as a and a′. Conversely, given a and a′, the corresponding point A can be reconstructed. Thus, the images in the left eye and right eye can be reconstructed, as well.

The quality of a stereoscopic panorama relies heavily on the accuracy of underlying pixel correspondence, which mainly depends on two factors: (1) whether correspondent pixels (a and a′ in Figure 2) on different sensors are captured synchronously by multiple cameras and (2) the accuracy of finding actual pixel correspondence post-capture.

The first factor could easily be addressed with a global-shutter imaging sensor, which captures every pixel at exactly the same time. However, one main design goal is affordability. The price of a global-shutter sensor with good imaging quality can be up to several thousand dollars, 10 times that of a good-quality rolling shutter. Unfortunately, building a stereo panorama with a rolling shutter entails formidable challenges for stereo-matching algorithms.

The second factor depends on the accuracy of the dense pixel matching method itself. The underlying technology, optical flow—though extensively studied—is still plagued by the inherent uncertainty stemming from image noise, light changes, lack of texture regions, multiple motions, repeated textures, and other complexities.

KANDAO’S CAMERA SYSTEM DESIGN

Kandao chose to use rolling shutters to bring the price down for consumers. But, to design a device that would be deployable to the general public, Kandao also had to focus on portability, reliability, high image quality, and user friendliness.

Because a VR camera is essentially a multi-camera system composed of multiple individual traditional cameras, Kandao designers set the form factors and weights to be similar to that of off-the-shelf commercial digital single-lens reflex (DSLR) cameras. This small form factor contributes to the device’s portability.

Reliability is also important; no camera in the multi-camera system should break down during shooting. Such a system’s reliability decreases as the number of imaging modules increases. Clearly, portability and reliability both demand a less complex system—one with fewer imaging modules.

Image quality refers to both raw image quality and stitching quality. Raw image quality is traditionally constrained by the lens and the image sensor. With the recent advances in semiconductors and computational photography, image quality is increasing steadily for small-sized imaging modules, as evident in the photo quality of modern smart phones. This enabled Kandao to fit multiple high-quality, compact lenses and sensors onto the same rig.

The more formidable challenge is to achieve good stitching quality. Computing an accurate ODS is much harder with fewer cameras. Kandao solved this problem by incorporating a deep neural network (DNN) specialized in highly accurate optical flow calculation. Solving the ODS synthesis problem with fewer cameras improves system reliability: the fewer components in a design, the more reliable the final product is.

Another important feature added to the design is a rich and convenient interface to further improve the camera’s user friendliness, as shown in Figure 3.

Figure 3. Kandao Obsidian cameras are easy to hold, carry, and interact with.

SOLUTIONS TO ENABLE DESIGN CHOICES

Building an ODS Dataset

A fundamental challenge for a machine-learning project is building a dataset that's right for the problem. A dataset containing ground truth suitable for computational ODS is very hard to construct, because one would need to accurately map every point densely in 3D and generate two slit cameras scanning the scene by rotating around an IPD circle. Even with a dedicated 3D mapping device such as LiDAR, fusing the accurate dense 3D point clouds and visual data is still time-consuming and expensive for static scenes. It is nearly impossible for dynamic scenes.

Collecting real-world ground truth is indeed too impractical, so Kandao took a computational approach. Kandao constructed a synthetic dataset with rich information, including 3D models (mesh and texture), simulated virtual cameras (with real intrinsic and extrinsic parameters), pixel correspondences (optical flow data), depth data, objects labels, and segmentation boundaries. The 3D scenes are built with realistic details such as lighting and reflection. Figure 4 shows an example of the dataset. The top image shows an indoor scene of a sitting room. The middle images show an image captured by a virtual fisheye camera and the associated depth map and label map. The bottom images were captured by other virtual cameras on the camera rig.

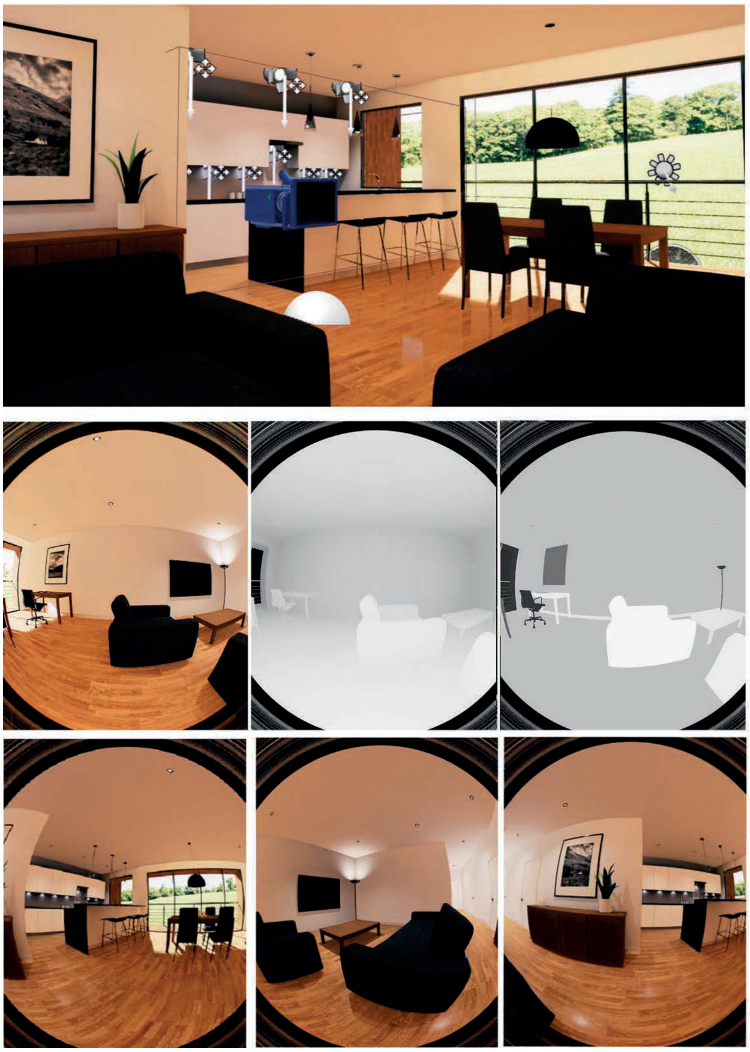

Figure 4. Images of the synthetic dataset built for simulation and training.

Solving ODS with fewer Cameras

The difficulty of building an ODS with fewer cameras is illustrated in Figure 5.

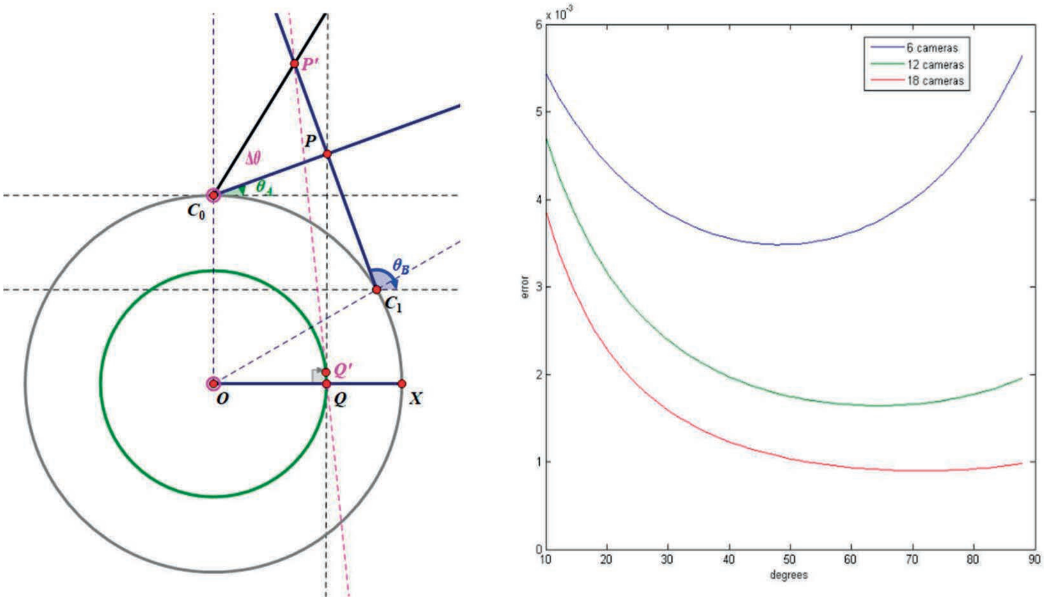

Figure 5. (left) Illustration of ODS projection error caused by error in optical flow. The error Δθ in optical flow causes *Q* to drift to *Q*'. (right) The error (arc length) of *QQ*' due to a unit error in optical flow at different ray positions (angles in degree) with different numbers of imaging modules.

C_0 and C_1 are the optical centers of two neighboring cameras. C_0P and C_1P are the rays coming from a physical point *P*. *Q* is the projected point (right eye) of *P* on the IPD circle. An error Δθ due to inexact match causes *Q* to shift to *Q*' and generates an error in the final stitched ODS panorama. The length of the arc *QQ*' indicates the amount of error on ODS.

A further geometrical analysis reveals that the average error *QQ*' is mainly related to $\angle C_0OC_1$, given that the average error (Δθ) of optical flow is known and form factor of the camera system is fixed. The average error decreases if $\angle C_0OC_1$ decreases; in other words, the number of imaging modules increases. An intuitive example is that when $\angle C_0OC_1$ decreases to zero, the system literally becomes an ideal ODS system with an infinite number of cameras. The right graph in Figure 5 shows the error (arc length) of *QQ*' at different ray positions (angles in degree) with different numbers of imaging modules. A numerical simulation found that to build a six-lens ODS system with a comparable stitching quality (similar amount of ODS ray drift) to a 12-lens or 18-lens design, the average end point error (EPE) of optical flow algorithm needs to be much smaller (34 to 53 percent) than those rigs. To build a high-quality ODS with six lenses, a very accurate optical flow algorithm is necessary.

A Deep-Learning Approach to Solving ODS

Optical flow is the method to compute a dense pixel correspondence, and its performance is crucial to image stitching and depth recovery. Although optical flow algorithms have advanced greatly, they are still error-prone in challenging situations such as lighting change, texture-less regions, multiple motions, and repeated textures. Encouraged by recent successes of deep convolutional neural networks (CNN) in a variety of computer vision tasks, Kandao constructed a robust and compact CNN, which solves the optical flow estimation problem as a supervised deep-learning task. The DNN can fuse multiple cues that many researchers believe are used in human vision, such as the binocular matching metric, monocular texture, and smoothness conformity. In this manner, the CNN results in more stable performance in the presence of different scenes, objects, and lighting conditions.

To handle occlusion, the deep-learning process proposes a novel multi-view depth fusion. To achieve the goal of compactness, recent CNN techniques are adopted (such as dense block[4], correlation volume[5], and differential softmax function[6]), and a number of novel techniques (such as deep cost aggregation and deep flow refinement) are further developed. These techniques contribute to a CNN that is compact enough to run on embedded devices. The proposed neural network is first trained on popular open-source datasets and further fine-tuned on the constructed dataset described previously. Compared with traditional optimization-based optical flow, the deep-learning-based method is 100 times faster. Compared with other deep-learning methods, the new method is more accurate, efficient, and compact. Figure 6 shows a comparison between a state-of-the-art method named FlowNet 2.0[7] and Kandao's new CNN method. The results of the new method show more accurate details than those of FlowNet 2.0.

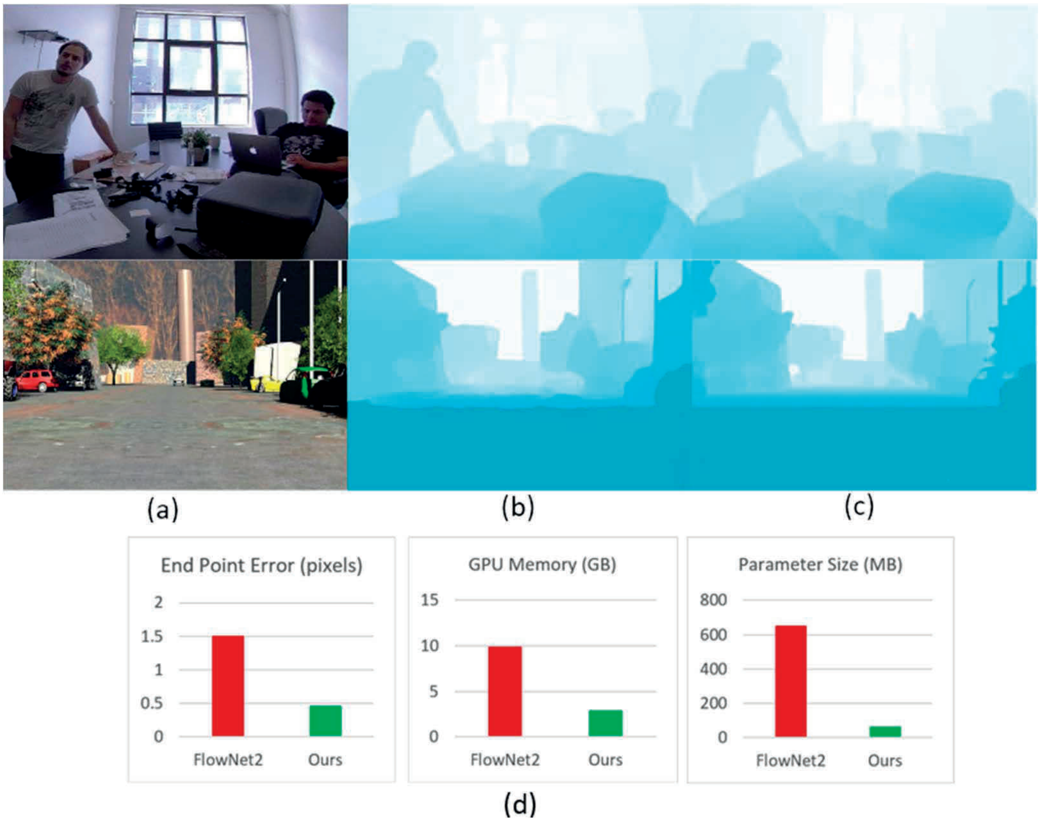

Figure 6. Flow method comparison: (a) input image, (b) FlowNet 2.0 results, (c) results using Kandao's CNN, and (d) a quantitative comparison on various criteria.

Minimizing Rolling-Shutter Effects

Unlike its global counterpart, which records the entire scene in one instant, a rolling-shutter sensor scans across the scene either vertically or horizontally. In other words, each row or column is recorded at a slightly different time. This produces undesirable distortions (such as wobbling, smear, and skew), especially when capturing fast-moving objects or rapid flashes of light—breaking the geometrical constraints that stereo matching algorithms follow.

To use rolling-shutter cameras in 3D geometric computer-vision tasks, the rolling-shutter effect must be carefully accounted for. Simply ignoring this effect and relying on a global-shutter method may lead to erroneous, undesirable, and distorted results. Instead, Kandao used both hardware and software methods to minimize the rolling-shutter effects. Instead of genlocked synchronization, Kandao chose a phase-delayed synchronization method to minimize the time difference over the overlapping area between neighboring cameras. Algorithmically, thanks to the accurate inter-frame pixel correspondences obtained from the aforementioned neural network, it becomes possible to estimate the object trajectory, rewarp each scanline, and reassemble them utilizing classic methods.[8,9]

VR Video Compression

Compressing VR video is challenging due to the large amount of data in the 360 degrees of viewing angles. The state-of-the-art high efficiency video coding (HEVC) standard for video compression can be applied directly for compressing high-resolution VR videos. HEVC provides many advanced features that can effectively handle high-resolution videos. These features include enhanced intra-prediction, hierarchical partitions of coding blocks, and new mechanisms for parallel processing. However, despite its advanced compression performance, using HEVC to compress a full 360-degree VR video would still result in an unacceptably high bitrate at reasonable video quality.

Before compressing VR videos, one must first project the original spherical representation onto a planar surface. A simple and straightforward way is to use an equirectangular projection, which is typically how one would present the surface of the earth on a geographical map. In this projection, the pole regions are stretched, and hence they contain more pixels per unit area compared to the equatorial regions. This unbalance in pixel density results in redundancy at the pole regions. The wasted pixels do not improve visual quality, but instead increase the total amount of data.

An ideal projection would provide uniform pixel density distribution across different viewing angles similar to the original spherical surface. However, it is difficult to design such a projection due to geometrical differences. Many alternative projections can be adopted to reduce the variations in pixel density and get closer to the ideal case. These projections include the cube-map projection, the equal-angular cube-map (EAC) projection,[10] and the barrel projection. Kandao applies the EAC projection in its system because it provides the least variation in pixel density distribution, saving about 25 percent of pixels when compared to the classic equirectangular projection.

Kandao applies the equal-angular cube-map (EAC) projection in its system because it provides the least variation in pixel density distribution, saving about 25 percent of pixels when compared to the classic equirectangular projection.

CONCLUSION

While VR has great potential to deliver a level of visual immersion unseen in previous video capture and display systems, bringing VR to the masses requires a camera that is affordable, lightweight, robust, and easy to use. As described in this article, companies like Kandao have tackled this challenge and are starting to deliver VR cameras to the market that satisfy these multiple, seemingly conflicting goals. Yet numerous technical challenges still remain. One challenge is how to exploit temporal correlation to further improve optimal flow for more temporal-consistent dense pixel correspondence. Another challenge is how to efficiently encode and package VR video into sub-streams for bandwidth-limited network transmission without adversely affecting user experience with round-trip time delay. Many research groups are currently addressing these and other challenges, and the day will soon arrive when VR video becomes truly ubiquitous—easily produced, consumed, and enjoyed by the masses.

REFERENCES

1. H. Ishiguro, M. Yamamoto, and S. Tsuji, "Omni-directional stereo," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, 1992, pp. 257–262; ieeexplore.ieee.org/document/121792/.
2. C. Richardt et al., "Megastereo: Constructing High-Resolution Stereo Panoramas," *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2013; ieeexplore.ieee.org/document/6619010/.
3. R. Anderson et al., "Jump: virtual reality video," *ACM Transactions on Graphics*, vol. 35, no. 6, 2016; dl.acm.org/citation.cfm?id=2980179.2980257.
4. G. Huang et al., "Densely Connected Convolutional Networks," *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2017; ieeexplore.ieee.org/document/8099726/.
5. A. Dosovitskiy et al., "FlowNet: Learning Optical Flow with Convolutional Networks," *IEEE International Conference on Computer Vision* (ICCV), 2015; ieeexplore.ieee.org/document/7410673/.
6. A. Kendall et al., "End-to-End Learning of Geometry and Context for Deep Stereo Regression," *IEEE International Conference on Computer Vision* (ICCV), 2017; ieeexplore.ieee.org/document/8237279/.
7. E. Ilg et al., "FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks," *arxiv.org*, 2016; arxiv.org/abs/1612.01925.
8. S. Baker et al., "Removing rolling shutter wobble," *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2010; ieeexplore.ieee.org/document/5539932/.
9. O. Saurer et al., "Rolling Shutter Stereo," *IEEE International Conference on Computer Vision*, 2013, pp. 465–472; computer.org/csdl/proceedings/iccv/2013/2840/00/2840a465-abs.html.
10. A. Wheeler, "Improving VR videos," *Engineering and Developers Blog*, blog, YouTube, 2017; youtube-eng.googleblog.com/2017/03/improving-vr-videos.html?utm_source=feedburner&utm_medium=feed&utm_campaign=Feed:+YouTubeEngDevBlog+(YouTube+Engineering+and+Developers+Blog).

ABOUT THE AUTHORS

Jack Tan is the cofounder of Kandao and has served as the CTO of the company since 2016. He has a PhD in electronic and electrical engineering from the University of Hong Kong. Before he started Kandao, he worked in roles ranging from engineering to administration in several different companies, including Hong Kong Applied Science and Technology Research Institute, Meizu.com, and Shenzhen Jieshun Science And Technology Industry. Contact him at jack.tance@gmail.com.

Gene Cheung is an associate professor at the National Institute of Informatics in Tokyo, Japan. He has been an adjunct associate professor at the Hong Kong University of Science and Technology (HKUST) since 2015. His research interests include 3D imaging and graph signal processing. He has a PhD in electrical engineering and computer science from the University of California, Berkeley. Contact him at cheung@nii.ac.jp.

Rui Ma is a research fellow at Kandao in Shenzhen, China. His research interests include VR video compression and transmission, panoramic representation, and multi-view navigation. He has a PhD in electronic and computer engineering from HKUST. Contact him at maruidavid@gmail.com.

This article originally appeared in IEEE MultiMedia, *vol. 25, no. 1, 2018.*

Virtual Travel

This article originally appeared in Computing in Science & Engineering, *vol. 20, no. 4, 2018.*

Charles Day
The American Institute of Physics

I spent the years 1988–90 as a postdoc at Japan's Institute of Space and Astronautical Science, whose campus is in Sagamihara, an industrial suburb of Tokyo. During that time, I lived in an apartment above a ramen restaurant called Koguma ("bear cub").

Koguma, it turned out, was one of the best ramen restaurants in Sagamihara. I ate there at least twice a week. After sampling everything on the menu, I favored two dishes: spicy leek ramen in miso broth, and beef ramen with a raw egg in soy broth. My next job was at NASA's Goddard Space Flight Center outside Washington, DC. I didn't eat ramen in DC until 2014, when the ramen restaurant Toki Underground opened on trendy H Street. Now the city has more than a dozen places to savor the toothsome noodle dish.

I applaud the proliferation of one of my favorite foods. On the other hand, it's yet another manifestation of the international homogenization of cuisine and culture. If you walk into a hipster coffee shop in Prague, Sydney, Buenos Aires, or Cape Town, you'll see chalkboard menus, artisanal biscotti, responsibly harvested coffee beans, and bearded baristas. Only the language spoken in the shop varies. Also converging toward a uniform look and feel are luxury beach resorts—at least judging by the interchangeable photos I see in the travel sections of the *New York Times*, *Wall Street Journal*, and *Financial Times*. The Frenchman David Guetta, the Dutchman Nicky Romero, the American Steve Aoki, the Finn Darude, and their ilk all produce blandly perky electronic dance music whose national provenance is impossible to discern.

If things taste, look, and sound the same, why travel? I don't know whether the leaders of Japan's largest airline, ANA, asked themselves that question, but they might have done. In April, ANA announced that it would be investing $10 million to develop a virtual travel system. The initiative is not as fanciful as it seems. The system's main ingredient is a robot avatar through whose senses any number of armchair tourists can experience remote locales. Besides acting as a tourist surrogate, the robot could conceivably perform tasks in hazardous environments. Indeed, one of the project's inspirations was the realization that a remotely controlled robot could have entered the reactor rooms of the Fukushima Daiichi power plant and averted nuclear disaster.

But if homogenization is making real tourism pointless, won't virtual tourism suffer the same fate? After all, virtual reality's principal figure of merit is how closely it matches reality. If I were ANA's CEO, Osamu Shinobe, I'd add fictional or past worlds to the robot avatar's menu of virtual destinations. Wouldn't you like to visit Middle Earth, Tatooine, ancient Rome, and Tang Dynasty China? It would be like Westworld but without rampaging robots and the need to fly.

ABOUT THE AUTHOR

Charles Day is *Physics Today*'s editor in chief. The views in this column are his own and not necessarily those of either *Physics Today* or its publisher, the American Institute of Physics.

 2469-7087/19/$33.00 © 2019 IEEE

BLACKBERRY CORPORATION seeks Sr. Software Test Engineer in San Diego, CA who will be responsible for sw quality assurance & production troubleshooting. Refer to Req#BB3 & mail resume to BlackBerry Corp., P.O. Box 141394, Irving, TX 75014.

QUALITY ASSURANCE SYSTEMS ENGINEER Position available in New York, NY. Manage the testing of advanced financial services applications utilizing tools such as Mocha, PhantomJS, and JavaScript. Create test plans, test data, and review and analyze results. Define test automation strategy, create and execute regression tests, and perform, adhoc, security, user interface, and integration testing. Prioritize quality assurance issues. Develop and maintain quality assurance guidelines and process documentation. Direct applications to: Recruiting Operations, Massachusetts Mutual Life Insurance Company, 1295 State Street, Springfield, MA 01111; Please Reference Job ID: 2362790

SENIOR SOFTWARE ENGINEER, Warren, MI, General Motors. Analyze reqmts, design, dvlp, validate &debug Android apps &embedded software for psgr vehicle infotainment ECUs for speech recognition, audio, media &projection on Android OS platform, using C/C++ &Java programming languages, &Android SDK, Android Studio, Eclipse, &Android Debug Bridge (adb) tools &Serial Terminal protocol. Design &dvlp Android Framework &HAL modules using Java, C &C++ programming languages to support vehicle features incldg speech recognition, audio, media &projection. Dvlp software, debug &fix issues in the Android Framework incldg AudioManager, AudioService, AudioPolicyManager, AudioFlinger, mediaserver &OpenSLES. Dvlp &fix issues in DomainServices for speech recognition, audio &media as part of Android Framework system server. Validate dvlpd software performing unit test using JUnit &GTest to ensure high qlty bug free software. Create &execute unit, integration, system, load &acceptance test plans &scripts. Use repo, git &gerrit configuration management tools for maintaining history &traceability of code changes. Bachelor, Computer Science, Computer Science and Engrg, Computer Engrg, or Electronics Engrg. 60 mos exp as Engineer or Developer, designing &dvlpg embedded software for infotainment ECUs or smartphones on Android OS platform. Mail resume to Ref#2190-201, GM Global Mobility, 300 Renaissance Center, MC:482-C32-C66, Detroit, MI 48265.

BLACKBERRY CORPORATION has the following openings in Mountain View, CA: Principal SW Test Eng (Req#BB4) Des. & dev. sw for secured mobile apps & web-based enterprise services. Principal SW Eng (Req#BB5) Des., dev., troubleshoot & debug sw enhancements & new products. Refer to Req# & mail resume to BlackBerry Corp., P.O. Box 141394, Irving, TX 75014.

SOFTWARE DEVELOPMENT ENGINEER IN TEST sought by Pactera, Inc. in Redmond, WA. Design, develop, test and deploy computer applications, test automation frameworks, databases, and services. Use SSIS technologies to extract & compute data. Des. & dev. Multi. databases using SSAS & OLAP technologies. Req BS+2 yrs exp. To apply send resume to pactera_resumeinc@pactera.com and reference Job# M0511039.50.

TECHNOLOGY

HomeAway.com, Inc.

has openings for following positions in **Austin, TX**:

Software Engineers
(Job ID#: 728.4961)

Design, implement, and debug software for computers including algorithms and data structures.

Staff ERP Database Engineers
(Job ID#: 728.4190)

Implement and enhance the growing Oracle applications footprint for the company.

To apply, send resume to: HomeAway Recruiting, 333 108th Avenue NE, Bellevue, WA 98004. Must reference Job ID#.

SOFTWARE

HomeAway.com, Inc.

has openings for

SOFTWARE ENGINEERS

(Job ID#: 728.5229)
in **Bellevue, WA**:

Design, implement, and debug software for computers including algorithms and data structures.

To apply, send resume to: HomeAway Recruiting, 333 108th Avenue NE, Bellevue, WA 98004. Must reference Job ID#.

SOFTWARE

Hotwire, Inc.

has openings for

SOFTWARE ENGINEERS TEST

(Job ID#: 728.5613)
in **San Francisco, CA**:

Code moderately complex tests to implement a test design, set up test automation and execution framework.

To apply, send resume to: Hotwire Recruiting, 333 108th Avenue NE, Bellevue, WA 98004. Must reference Job ID#.

RESEARCHER, Warren, MI, General Motors. Research, dvlp &implement innovative &creative algorithm solutions to create new psgr vehicle cloud computing system architecture, using multi-modal data fusion &machine learning, for collected vehicle systems &active safety related apps incldg V2V &V2X in Android &Linux platforms, Java programming &MATLAB, Python &SPARK tools. Design control algorithms &software for complex systems. Create &utilize MATLAB models for algorithm definitions &utilize plant models for verification of algorithm &control features. Dvlp cloud computing system using sensing &fusion of data from multi-modal vehicle sensors (speed, tire rotation, acceleration, &yaw rate) &smart phone sensors (gyroscope &inertial) to assess behavior of driver &the context of driving conditions. Research, dvlp &integrate multi-modal vehicle-based sensors incldg DSRC-based V2X. Assess situation awareness &collisions potential with objects. Dvlp reqmts, criteria, &evaluation methods for vehicle-based sensing technologies incldg cameras. PhD., Computer Science, Electrical or Computer Engrg. 6 mos exp as Engineer, researching &dvlpg algorithm solutions to create psgr vehicle cloud computing system architecture, using multi-modal data fusion &machine learning, for collected vehicle systems &active safety related apps incldg V2V &V2X in Android or Linux platforms. Mail resume to Ref#2606, GM Global Mobility, 300 Renaissance Center, MC:482-C32-C66, Detroit, MI 48265.

SOFTWARE ENGINEER, Warren, MI, General Motors. Design &dvlp telematics &infotainment OTA programming master framework for global psgr vehicle platforms using C/C++ on QNX &Linux platforms in Communication Gateway Module (CGM) &Telematics Communication Platform (TCP). Implement UDS diagnostics over Ethernet (DoIP &Socket Adapter) &CAN using C/C++. Design &dvlp ECU updates &debug on target hardware. Perform static code analysis using Coverity &Parasoft. Perform software configuration management using Git, &code review using Gerrit. Integrate &test software on vehicle bench. Provide solutions to Ethernet network configuration incorporating vehicle architecture &vehicle IP network reqmts. Build automation dvlpmt on Jenkins. Mentor new members of the team. Design &implement embedded apps using C, C++ &debug on target hardware. Master, Computer or Electrical Engrg. 12 mos exp as Engineer, Developer, or related, designing &dvlpg telematics &infotainment OTA programming master framework for psgr vehicle platforms using C/C++ on QNX &Linux platforms in CGM or TCP. Mail resume to Ref#3210-204, GM Global Mobility, 300 Renaissance Center, MC:482-C32-C66, Detroit, MI 48265.

FILEMAKER INC. has multiple positions open for the following in Santa Clara, CA: Software Engineer Applications (REQ#AVE36M) Desgn & devlop e-commrce SW. Refer REQ# & mail resume to: FileMaker Inc., ATTN: CP, 5201 Patrick Henry Drive, Santa Clara, CA 95054. FileMaker is an EOE/AA m/f/disability/vets.

VLOCITY seeks Software Engineer, Engineering Productivity in SF, CA to dev internal tools used by co. engineering., testing, & build teams. Ref Job ID: AGM3N5 & send res to L. D. Kingsley at hiring@vlocity.com

TECHNICAL

Oracle America, Inc.

has openings for

TECHNICAL ANALYST

positions in **Bedford, MA**.

Job duties include: Deliver solutions to the Oracle customer base while serving as an advocate for customer needs.

Apply by e-mailing resume to larry.campbell@oracle.com, referencing 385.21488.

Oracle supports workforce diversity.

TECHNOLOGY

Oracle America, Inc.

has openings for

APPLICATION DEVELOPER

positions in **Arlington, VA**.

Job duties include: Analyze, design, develop, troubleshoot and debug software programs for commercial or end-user applications. Write code, complete programming and perform testing and debugging of applications.

Apply by e-mailing resume to will.moore@oracle.com, referencing 385.22307.

Oracle supports workforce diversity.

CONSULTANT

Oracle America, Inc.

has openings for

CONSULTANT

positions in **Atlanta, GA**.

Job duties include: Analyze requirements and deliver functional and technical solutions. Travel to various unanticipated sites throughout the United States required.

Apply by e-mailing resume to prakash.ramanan@oracle.com, referencing 385.21560.

Oracle supports workforce diversity

EPAM SYSTEMS, INC. seeks all levels for the following positions in Newtown, PA and various long term unanticipated worksites. Account Manager, strategize progs to meet cl goals (AM0818); Application Support Analyst, create process docs (ASA0818); Application Support Manager (ASM0818); Business Analyst, create req's & specs (BA0818); Business Analysis Manager (BAM0818); Data Anaylst (DA0818); Database Developer, analyze, design, maintain DBs (DD0818); Delivery Manager, deliver sw, arch. (DM0818); IT Manager, mng local IT prog (ITM0818); Project Administrator (PA0818); Project Manager, mg IT dev projs (PM0818); Software Developer (SD0818); Software Engineer (SE0818); Software Engineering Manager, proj coord. & integration (SEM0818); Software Test Automation Engineer, create auto test plns (STA0818); Software Test Automation Manager (STAM0818); Software Testing Engineer, test & QC sw (STE0818); Software Testing Manager (STM0818); Solution Architect, prod sol dev & arch. (SA0818); Systems Engineer, design sw prod environs (SYE0818); Experience Designer, user research & testing (ED0818); Performance Analyst, test sw performance (PFA0818); Security Systems Engineer, design, improve ITIL/ITSM & SDLC (SCSY0818). Mail resume to Marina Maizet, EPAM, 41 University Drive, Suite 202, Newtown, PA 18940. Must ref job code.

SENIOR EMBEDDED SOFTWARE ENGINEER, Warren, MI, General Motors. Engr, perform, dvlp, using C, C++, &Assembly languages, &execute VIP logic controller for psgr car, truck &SUV Center Stack Module to meet nextgen automotive infotainment systems &in vehicle communications reqmts in compliance with ISO26262 AUTOSAR standards. Deliver audio, tuner, diagnostics, calibrations, camera, vehicle components inter process &intra process communication, &integrate system &hardware peripherals, HMI framework &applications, telematics, vehicle data, navigation, display, chimes, power mode &system states, storage, memory, verified boot, software update functionality to dvlp best in-class embedded software using Android, Java, C/C++, &Kotlin. Engr, design, dvlp, release &continuously improve psgr vehicle ECU &software architecture reqmts &robustness using DOORS for reqmts handling, Vector DaVinci Configurator &Developer, Rhapsody &EA tools for design &development, &Jenkins for automated backend continuous integration. Use Renesas 32bit controller as hardware platform. Bachelor, Computer Science, Computer Engrg, Computer Science &Engrg, or Electrical Engrg. 60 mos exp as Engineer, Developer, or related, delivering audio, calibrations, camera, system process communication, &integrating system &hardware peripherals, HMI framework, navigation, display, storage, &software update functionality to dvlp embedded software using Android, Java, &C/C++. Mail resume to Ref#5891, GM Global Mobility, 300 Renaissance Center, MC:482-C32-C66, Detroit, MI 48265.

SOFTWARE ENGINEER - AUTOMATED DRIVING - ADVANCED PROJECTS, Warren, MI, General Motors. Create, write, debug, &test software using C/C++, Python, MATLAB, Simulink, incldg toolboxes, Integrated Development Environment (IDE), compilers, debuggers, build system, libraries &frameworks, in Linux &Windows to ensure safe &proper functioning of safety critical autonomous vehicles incldg LIDAR, RADAR, high precision camera, &IMU sensors. Work on complex problems involving Coordinate Transformations, Object Perception (includes images, point clouds, segmentation, registration), Mapping &Localization, Path Planning, Obstacle Avoidance, Object Tracking, &Vehicle Control. Use Agile dvlpmt utilizing code repository, code review, work &issue tracking tools including Git/Bitbucket and RTC. Study existing technologies, tools, processes &algorithms. Implement algorithms &software to fit use case &compare results to find best approach to solve problems. Design test strategies, test &debug implemented software. Conduct code review to make sure code follow coding standards. Collect required data for software dvlpmt. Study &create mathematical &machine learning model of system for anlys. Integrate resulting software &model into advanced project dvlpmt vehicles, SIL &HIL simulators. Master, Computer Engrg, Electrical Engrg, Computational Sciences &Robotics, or Computer Science. 12 mos exp as Engineer, creating, debugging &testing software using C/C++ &Python, incldg IDE, compilers, debuggers, build system, libraries &frameworks, in Linux &Windows to ensure safe &proper functioning of safety critical robot or autonomous vehicle with high precision camera, or related. Mail resume to Ref#269, GM Global Mobility, 300 Renaissance Center, MC:482-C32-C66, Detroit, MI 48265.

SOFTWARE

Oracle America, Inc.

has openings for

SOFTWARE DEVELOPER

positions in **Seattle, WA**.

Job duties include: Design, develop, troubleshoot and/or test/QA software. As a member of the software engineering division, apply knowledge of software architecture to perform tasks associated with developing, debugging, or designing software applications or operating systems according to provided design specifications.

Apply by e-mailing resume to suresh.pasala@oracle.com, referencing 385.22229.
Oracle supports workforce diversity.

TECHNICAL

Oracle America, Inc.

has openings for

TECHNICAL ANALYST

positions in **Orlando, FL.**

Job duties include: Deliver solutions to the Oracle customer base while serving as an advocate for customer needs. Offer strategic technical support to assure the highest level of customer satisfaction.

Apply by e-mailing resume to shantanu.sen@oracle.com, referencing 385.225444.
Oracle supports workforce diversity.

TECHNOLOGY

LinkedIn Corp.

has openings in our **Mountain View, CA** location for:

Software Engineer (All Levels/Types) (SWE1218MV) Design, develop & integrate cutting-edge software technologies.

LinkedIn Corp. has openings in our **Sunnyvale, CA** location for:

Software Engineer (All Levels/Types) (SWE1218SV) Design, develop & integrate cutting-edge software technologies; **Data Scientist (6597.943)** Work with a team of high-performing analytics, data science professionals, & cross-functional teams to identify business opportunities & optimize product performance or go-to-market strategy; **Data Scientist (6597.2702)** Manipulate massive-scale structured & unstructured data to drive actionable insights. **Systems Architect, Global Sales Comp. Reporting & Insights (6597.2975)** Work with business stakeholders & IT to plan, design, code, test, & roll-out an enterprise scale Data Warehouse & Business Intelligence Solution; **Site Reliability Engineer (6597.2543)** Ensure that LinkedIn's complex, web-scale systems are healthy, monitored, automated, & designed to scale.

LinkedIn Corp. has openings in our **San Francisco, CA** location for:

Software Engineer (All Levels/Types) (SWE1218SF) Design, develop & integrate cutting-edge software technologies; **Staff Network Engineer (6597.1289)** Design, engineer, & deploy massive, & scalable datacenter & large-scale backbone & edge-pop infrastructure.

LinkedIn Corp. has openings in our **New York, NY** location for:

Software Engineer (All Levels/Types) (SWE1218NY) Design, develop & integrate cutting-edge software technologies.

Please email resume to: 6597@linkedin.com. Must ref. job code above when applying.

TECHNOLOGY

Oracle America, Inc.

has openings for

APPLICATIONS DEVELOPER

positions in **Arlington, VA**.

Job duties include: Analyze, design, develop, troubleshoot and debug software programs for commercial or end-user applications. May telecommute from home.

Apply by e-mailing resume to joshua.lurz@oracle.com, referencing 385.22342.

Oracle supports workforce diversity.