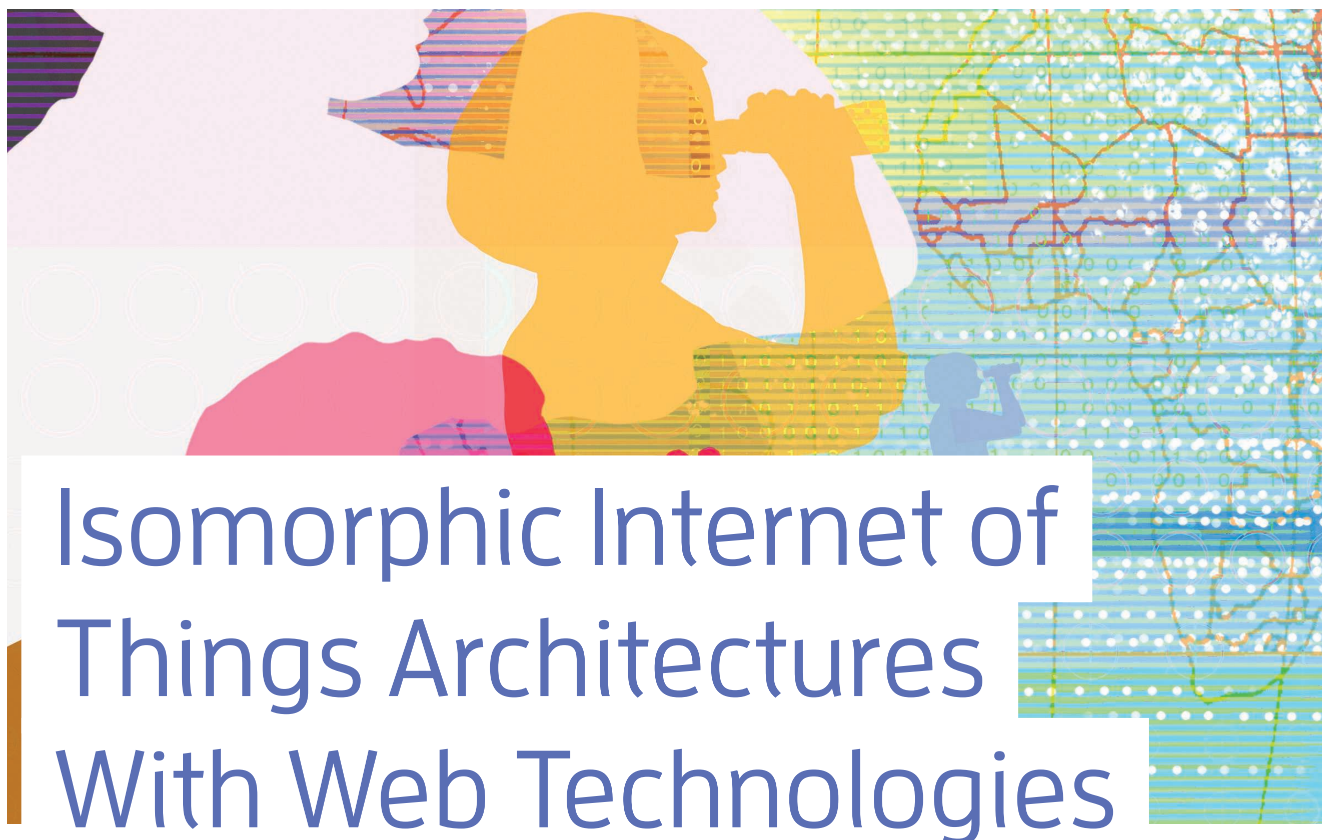

Isomorphic Internet of Things Architectures With Web Technologies

Tommi Mikkonen, University of Helsinki
Cesare Pautasso, Università della Svizzera italiana
Antero Taivalsaari, Nokia Bell Labs

Internet of Things development needs isomorphic software architectures, in which every kind of device can be programmed with a consistent set of implementation technologies, allowing applications and their components to be statically deployed or dynamically migrated without having to change their shape.

Recent years have witnessed an avalanche of digitalization technologies. Processing capabilities have grown dramatically, cloud computing has become a commodity, data science has blossomed due to increasing amounts of data, and artificial intelligence (AI) and machine learning (ML) have emerged as everyday technologies even in devices with limited capabilities, such as mobile phones. These changes are leading us to a "programmable world,"[17] where everyday things around us will become connected and programmable.

A hallmark of the trend toward the programmable world is the Internet of Things (IoT) development. A typical IoT architecture comprises a number of components, including

- sensors and actuators that are at the edge of the network
- gateways that connect them to the Internet

Digital Object Identifier 10.1109/MC.2021.3074258
Date of current version: 1 July 2021

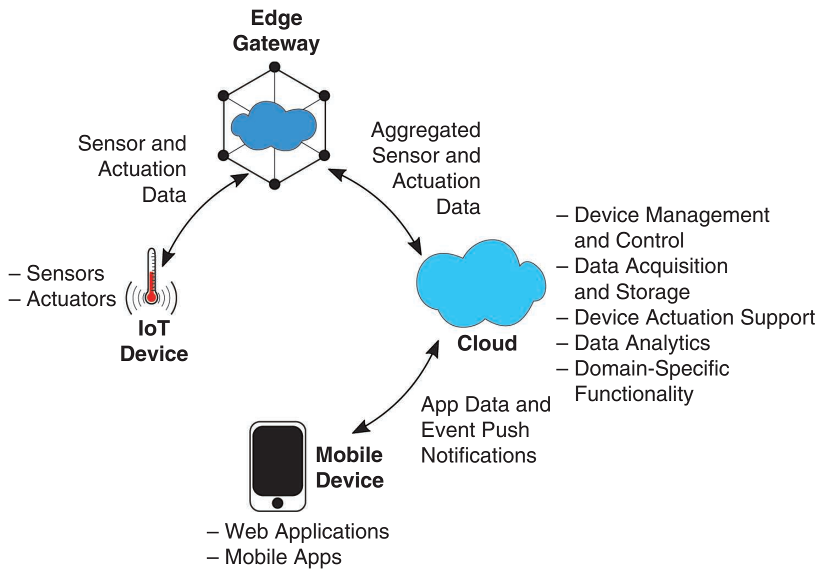

FIGURE 1. A typical IoT system end-to-end architecture.

Hardware Component		Operating System	Virtualization	Programming Language
	Cloud	Linux	Virtual Machine Container	JavaScript Python ...
	Edge Gateway	Linux Android	Virtual Machine Container	JavaScript Java ...
	Mobile Device	Android iOS	App Sandbox	Swift Kotlin
	IoT Device	Zephyr OS FreeRTOS QNX VxWorks	None	C C++

FIGURE 2. An example of the current platform diversity in the context of IoT systems.

- cloud services that offer large amounts of storage
- end-user applications that enable access to data, sensors, and actuators
- scalable analytics facilities that are deployed in the cloud.[14]

A typical IoT system end-to-end architecture is illustrated in Figure 1. Today, a wide variety of implementation technologies are used for developing different parts of the end-to-end IoT system (see Figure 2). This results in diverging development and deployment practices as well as higher integration costs.

In this article, we argue that a unifying software layer is needed to manage the complexity of IoT development and liberate developers from the highly fragmented IoT architectures of today. The work presented in this article is a continuation of a series of vision articles that describe liquid, multidevice software architectures[15] and the programmable world concept.[13] In the present article, we push the envelope toward isomorphic IoT systems, following the same line of thought and motivation.

ON ISOMORPHIC SOFTWARE

Isomorphic means "with the same shape." The word *isomorphism* is derived from the ancient Greek: is *ισοζ*, or *isos*, which signifies "equal," and *μορφη*, or *morphe*, which represents "form" or "shape." Isomorphism is a popular, well-established concept in mathematics. However, in the context of software development, the idea emerged relatively recently. For instance, for web applications, isomorphism refers to the ability to run the same code on both the back end (cloud) and front end (web

browser). More broadly, isomorphic software architectures feature software components that do not have to be modified ("change their shape") when running across the different hardware or software components of the system; some examples of isomorphism in the context of software systems are listed in Table 1.

In principle, writing software for isomorphic architectures is fundamentally simpler since the same code can run everywhere. Because the underlying technologies are handled uniformly, developers do not have to master different development technologies, and, thus, complexity is tamed considerably.

Several different levels of isomorphism can be identified. At the first level, isomorphism refers to the consistent use of the same development technologies across the different computational elements in the entire system. In contrast with such static development-level isomorphism, in dynamic isomorphism, a common runtime engine or virtualization solution is used so that the same code can run in different computational elements without recompilation. In an even more advanced system, the dynamic migration of code from one computational element to another is enabled.

In the case of IoT applications, the same (that is, isomorphic) software can ideally be deployed throughout the end-to-end system to run on edge devices, gateways, mobile clients, and cloud services. However, as we discuss, current IoT systems are a far cry from this ambition. Today, IoT application developers must be aware about the deployment context for their code, and they must be familiar with many different programming languages as well as virtual runtime environments and communication protocols (Figure 2). This platform diversity can make it impossible, for example, to redeploy components from the edge to the cloud without a complete rewrite.

While isomorphic architectures will make it easier, faster, and potentially cheaper to develop IoT applications and systems, we predict that they will also enable new kinds of dynamic applications that take advantage of the possibility to dynamically

TABLE 1. Examples of isomorphic software.

Technology	Description of isomorphic features
Java (1995)	The "Write once, run everywhere" slogan popularized by the Java platform[1] captures the essence of static isomorphism in software. In Java, the concept meant that it is possible to run the same software on different computer architectures and operating systems using a virtual machine.
Squeak Smalltalk (1996)	Virtual machines for the Squeak Smalltalk system are available for many operating systems and hardware platforms, making it possible to run bit-identical images across all.[4]
Unity (2005)	The Unity 3D development platform was born within the gaming domain, but it has recently branched out to the cinematics, automotive, and architecture domains. Applications written for Unity can run across 25 different platforms, including gaming consoles, but also mobile devices, virtual reality headsets, and smart TVs.
Lively Kernel (2007)	Lively Kernel is a web framework where applications are composed with JavaScript, and the code can be run on either the client or server side.[5]
Isomorphic web apps (2013)	The term *isomorphic web app* was introduced in the context of web applications in mid-2010s, referring to the ability to allocate a part of a web application's functionality either on the server or client.[12] While the term was new, the same idea has been used in the context of the web previously, for example, in the Lively Kernel mentioned earlier.
Universal Windows Platform (2015)	Within the Microsoft ecosystem, this platform enables developers to write and run the same software on computers and tablets running Windows 10, Xbox One gaming machines, and HoloLens devices.
Liquid web apps (2015)	Liquid web applications[11] allow the migration of their user interface (UI) components on the fly, allowing users to flexibly use applications on different devices and screens. The main focus in this work is on user experience: how to seamlessly move, clone, and adapt UI components and entire user experiences from one device to another.

redeploy and migrate application components from the edge to the cloud (and vice versa).

CHALLENGES IN IoT DEVELOPMENT

Diversity of programming models

Nowadays, the vast majority of software developers have been trained to do either mobile development or web development.[18] Many of these developers tend to assume that their skills are directly applicable to IoT development. However, IoT systems have many characteristics that do not apply to mobile or web applications. IoT developers must consider several factors that are unfamiliar to most application developers. Such factors include

- multidevice programming
- the heterogeneity and diversity of devices
- intermittent, potentially unreliable connectivity
- the distributed, always-on nature of the overall system
- the general need to write software in a highly fault-tolerant and defensive manner.

Moreover, a typical IoT application is continuous and reactive. On the basis of observed sensor readings, computations get triggered (and retriggered) and, eventually, result in various actionable events. The systems are essentially asynchronous, parallel, and distributed. These qualities alone make IoT applications very different from traditional PC, mobile, or web applications, in which software is typically written for a single client that may communicate with a single back-end server.

In general, IoT devices are bringing back the need for embedded software development skills and education. Software development for IoT devices is very similar to "classic" embedded systems development, as they both require small-memory and energy-aware software development skills. This is an especially relevant note from an education viewpoint since, in the past 10–15 years, many universities—at least in Northern Europe—have scaled back their courses on embedded systems and control theory, focusing on presumably more modern and desirable areas, such as web and mobile software development, instead. Recent Developer Economics survey reports strongly confirm the focus on higher-level programming skills.[18]

While IoT device development is bringing back the need for embedded software, at the other end of the spectrum of IoT end-to-end systems, cloud development relies heavily on multiple layers of virtualization. In a modern microservice-based software architecture, the built-in assumption is that all of the microservices must be turned into Docker containers.

Furthermore, those Docker containers are then assumed to be run in a Kubernetes cluster. This has to be done in spite of the fact that the underlying components are commonly written in Python or JavaScript/Node.js (thus requiring a virtual machine language runtime), and they typically run in a virtual machine rented from third-party cloud service providers, such as Amazon or Microsoft.

Dockerization and the use of Kubernetes effectively means that modern software systems commonly use a minimum of four virtualization layers, even for the simplest of cloud components. The additional virtualization layers often offer little additional value, add overhead to the development process, slow down application execution, and make debugging of the system more difficult. Nevertheless, the use of virtualized software environments is rapidly spreading to IoT edge systems as well, including gateway development.

Virtualization layers add complexity to just about every step of the development process. Dealing with this complexity necessitates a lot of boilerplate software that presumably helps but often distracts developers from focusing on the essentials of the applications. Despite the extensive use of virtualization, IoT systems still suffer from a rigid and fragmented architecture in which tasks cannot be reallocated easily from one computational element to another.

INDUSTRIAL EXAMPLE

Let us present a brief industrial IoT system example to illustrate the current diversity. In this system, a company has developed an industrial measurement and tracking solution that consists of a large number of devices custom built for different measurement and tracking tasks.

Examples include devices for tracking air quality (temperature, humidity, air pressure, and indoor air pollution) and movement (based on both inertial measurement units and indoor localization technologies) as well as those for measuring ambient noise and luminosity (including infrared light level). Devices are connected to the network either via short-area radio, such as Wi-Fi, or low-power wide-area network solutions, such as narrowband-IoT (NB-IoT) or LTE-M. Data uploading and actuation are

performed using the Message Queuing Telemetry Transport (MQTT) protocol (https://mqtt.org/).

In this case study system, sampling and data-upload rates are relatively moderate. On average, data uploading is performed only every few minutes. This simplifies the implementation of the cloud back end quite considerably since there is no requirement for continuous data streaming or extremely low latencies. Figure 3 provides an overview of our case study system.

The measurement devices are built on top of available off-the-shelf hardware. For Wi-Fi-based devices, the popular ESP32 (https://www.espressif.com/en/products/socs/esp32) platform was used. For NB-IoT/LTE-M-based devices, Nordic Semiconductor's nRF91 (https://www.nordicsemi.com/Products/Low-power-cellular-IoT/nRF9160) was chosen. The development language for both of these device platforms is C, but application programming interfaces (APIs), libraries, and tools vary considerably since a different

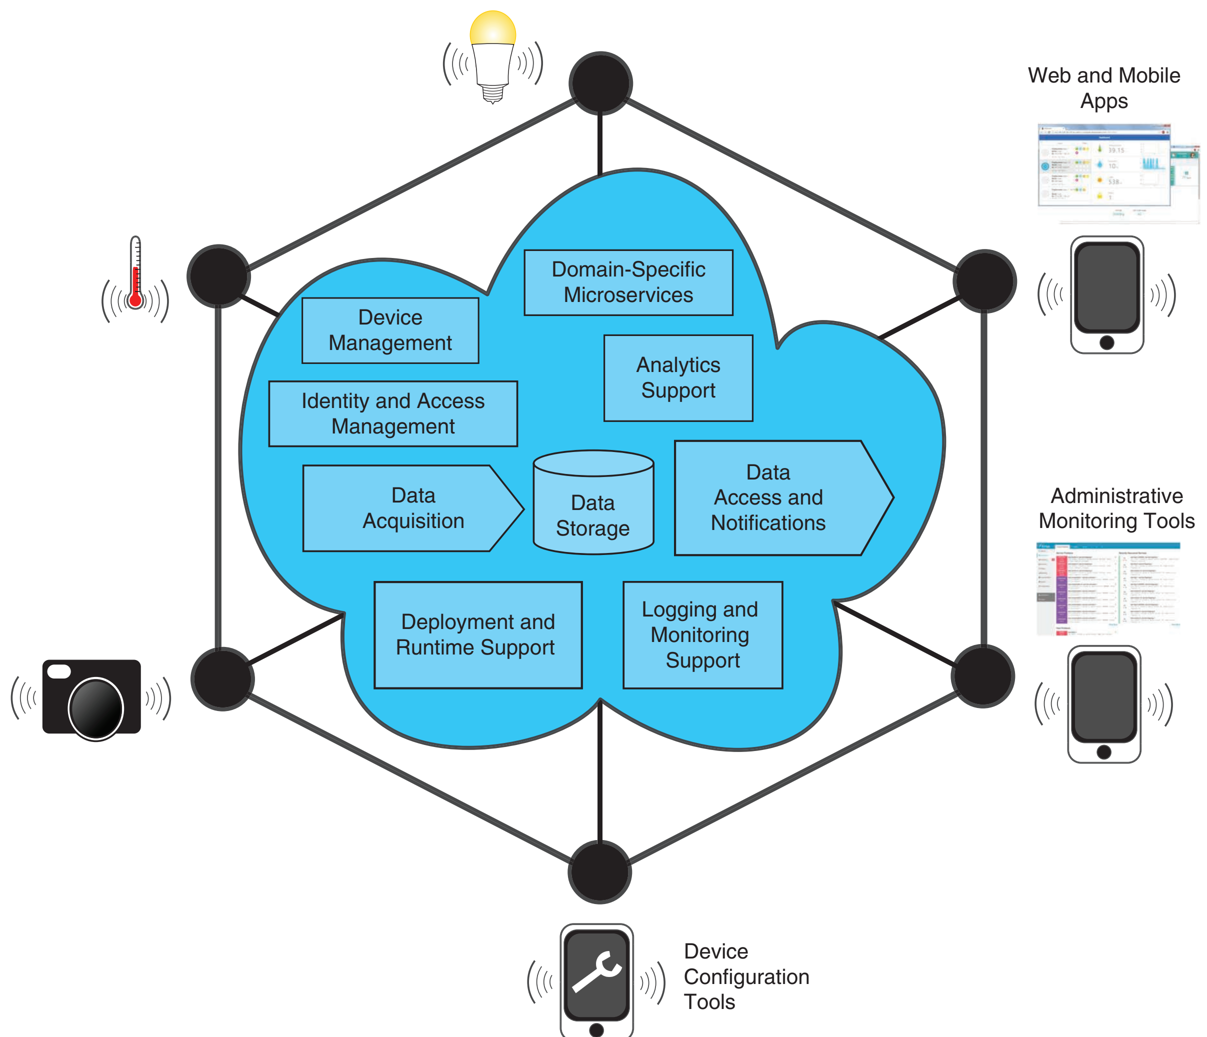

FIGURE 3. An overview of our case study system, including its key subsystems and related applications.

real-time operating system is used in each device platform.

Since data uploading is performed over Wi-Fi or cellular, this use case does not require any custom-built gateway devices running dedicated protocol translation stacks. However, for device configuration purposes, an Android mobile app—written in Java utilizing the Android libraries—was developed as well.

A lot of focus in the development effort was placed on developing the cloud back end. The logical components of the back end are depicted in Figure 3, and the majority were implemented using open source components.

For security perimeter/reverse proxy implementation, NGINX (https://nginx.org/) was chosen. For data acquisition (collection of sensing data from devices), both Apache Kafka (https://kafka.apache.org/) and RabbitMQ (https://www.rabbitmq.com/) are used. For logging and system monitoring, Grafana (https://grafana.com/), Graphite (https://graphiteapp.org/), and Icinga (https://www.icinga.com/) were picked. Data analytics capabilities were originally implemented using Apache Storm (http://storm.apache.org/), but these were later replaced with Apache Spark (https://spark.apache.org/). Domain-specific microservices were all implemented in Node.js.

The entire back end is Dockerized; for instance, each of the microservices runs in its own Docker container. Ansible (https://www.ansible.com/) and OpenStack (https://www.openstack.org/) were utilized in the original deployment, but, later, the entire system was migrated to run in a Kubernetes cluster (https://kubernetes.io/).

In addition to the devices and cloud back end, some additional web and mobile applications were developed for data visualization purposes as well as system administration and monitoring. In web application development, an earlier version of Angular.js (https://angularjs.org/) was used, whereas mobile apps were written in Java/Android Studio.

As can be determined from this discussion, the development of the entire end-to-end system required a very broad palette of technologies ranging from embedded, mobile, and web application technologies to a spectrum of popular cloud back-end implementation components. Given the breadth of the technologies, it would be almost impossible for an individual developer or a small startup company to master all of the necessary technologies to develop the entire system. Furthermore, because of the selected technologies, each of the components is rather tightly coupled with a specific computational element in the end-to-end system.

ADDITIONAL CATALYSTS FOR CHANGE

Intelligence at the edge

In "classic" IoT systems, such as our case study system, the majority of computation and analytics are performed in the cloud in a centralized fashion. However, in recent years, there has been a noticeable trend in IoT system development to move intelligence closer to the edge.

Historically, the computing capacity, memory, and storage of edge devices were limited. Due to the increasing computational capabilities of edge devices and requirements for lower latencies, though, intelligence in a modern end-to-end computing system is gradually moving toward the edge, first to gateways and then to devices. This includes both generic software functions, and—more importantly—time-critical AI/ML features for processing data available in the edge with minimal latency. The requirement to run advanced AI/ML and analytics algorithms in the edge increases the demand for consistent programming technologies across the end-to-end system.

The increasingly dynamic nature of IoT systems

In IoT systems that consist of a massive number of devices overall, device topologies can be expected to be highly dynamic and ephemeral. This dynamism calls for technologies that can cope with dynamically changing "swarms" of devices and their evolving responsibilities. The increasingly dynamic nature is not only related to software features but also to AI/ML capabilities, where reinforcement learning can introduce unexpected situations—something that worked yesterday might not work today, and vice versa. Furthermore, since such features are wrapped in software components, it is often expected that they can be relocated to the best-suited context for execution.

To simplify development, deployment, and long-term use, we expect that future IoT systems will need to support a very flexible allocation of responsibilities so that the roles of devices can evolve over time. This calls for a platform in which different computational entities can run the same code.

PREDICTING THE RISE OF ISOMORPHIC IoT SYSTEMS

With the ever-increasing complexity, dynamism, and sheer amount of software, we are on a dangerous trajectory at the moment. The rigid system

architectures, broad spectrum of technologies, abundant use of cargo-cult reuse (picking certain implementation technologies and methods simply because others have done so), inconsiderate use of virtualization, and highly virtualized deployment and package management approaches are leading us to IoT systems that become increasingly difficult to manage. It is time for a change.

Going forward, we need technologies that liberate us from rigid task allocation and support the use of consistent implementation technologies. Dynamic component deployment should be supported, but in a fashion that emerges from application needs—especially in relation to performance and reliability—and not due to constraints imposed by the dominant development platforms and tools. Moreover, various features (especially AI/ML capabilities) may require the flexible migration of code and models from the cloud to the edge (and vice versa), depending on the data availability and required response times.

Our prediction is that these demands will eventually lead us to isomorphic IoT system architectures, in the spirit of isomorphic web applications.[12] *Isomorphic web applications* commonly refers to the ability to use the same development technologies and code between the front and back ends. Just as we are currently witnessing in the IoT area, isomorphic web applications emerged from an initially fragmented technological landscape in which the development technologies for the web browser and server were entirely different.

Many web developers surely still remember the era when back-end functionality was written in PHP or Perl, resulting in a deep divide between the programming languages used on the client and server sides. Once the use of JavaScript spread to the back end,[16] it became gradually possible to run the same code on both sides as long as the code would rely on compatible library dependencies and comply with different sandboxing restrictions.

In an isomorphic IoT system, devices, gateways, and cloud back-end features and front-end applications will be written using the same technologies and, ideally, be able to run the same software components, allowing the flexible migration of code among components in the overall system. Instead of having to learn many incompatible software development platforms, in an isomorphic architecture, one base technology will suffice and be able to cover all aspects of end-to-end development; the same tools can then be used to compose the software across all of the computing units (Figure 4).

The two key technical elements that are needed for implementing such systems are a uniform API for accessing features of different subsystems and a common runtime that is fast but small enough for embedded devices yet powerful enough to implement lightweight containers to deploy applications everywhere. In addition, an orchestrator function (such as those described by Kurzyniec et al.[7] and Mäkitalo et al.[8]) is needed that will guide the deployment and potential migration of the different subsystems.

More broadly, the "holy grail" in the IoT area is a common programmable world API that would cover device discovery, data acquisition, data access, device actuation, device management, code updates, debugging, and other relevant topics in a universal fashion—thus working universally across devices from different domains, manufacturers, and the necessary security mechanisms. While it is debatable whether there will ever be a single API to cover IoT devices from entirely different domains, it is safe to bet that, in five to 10 years, IoT devices and their APIs will have converged significantly. It

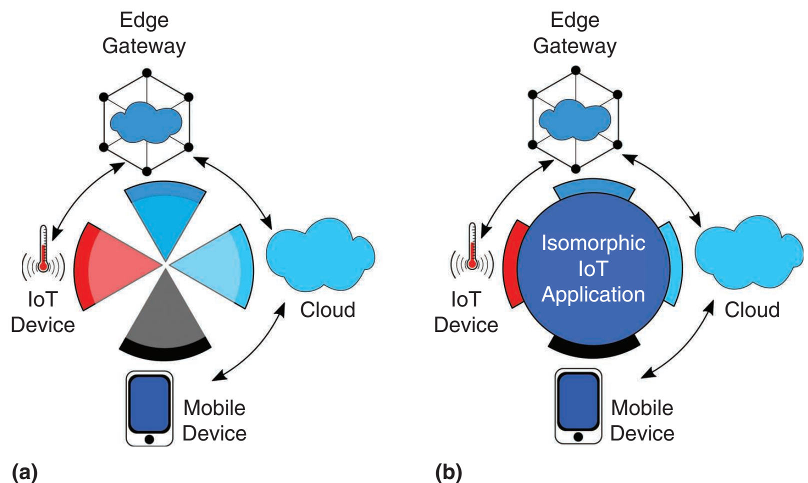

FIGURE 4. The (a) classic versus (b) isomorphic IoT architectures.

is also very likely that the necessary infrastructure will grow around the already existing IP networking and web infrastructure.

ISOMORPHIC IoT WITH WEB TECHNOLOGIES

In seeking concrete technology candidates for implementing isomorphism, we have turned to web standards since they have played a unifying role in many other contexts. For uniform APIs in the isomorphic IoT system context, the most prominent candidate today is the Web of Things (WoT), a set of standards for solving the interoperability issues of different IoT platforms and application domains.[20] In essence, the WoT makes each "thing" part of the Web by giving it a uniform resource identifier that can be used for communicating with it. The communication with each thing is supported by a common data model and uniform API that is recognized by every thing.

For the isomorphic IoT runtime, the web provides two prominent options: JavaScript/ECMAScript[3] and WebAssembly (WASM).[19] The former is the de facto language for web applications both for the web browser and cloud back end (Node.js); it is currently the most viable option for implementing static isomorphism, that is, to allow the use of the same programming language throughout the end-to-end system.

The latter is a binary instruction format to be executed on a stack-based virtual machine that can leverage contemporary hardware[2,6]; we see WASM as the best option for providing support for dynamic isomorphism, that is, the ability to use a common runtime that is powerful but small enough to also fit in low-end IoT devices (Figure 5). Note that these options are not mutually exclusive; that is, it would be possible to implement an architecture in which WASM is used as the unifying runtime, but JavaScript is used as the programming language throughout the end-to-end system.

Both options have their pros and cons in the context of isomorphic IoT applications. JavaScript offers massive library support [more than a million Node Package Manager (NPM) modules], a large number of developers familiar with the language, and high-performance virtual machines. However, for isomorphic applications, the dynamic nature of JavaScript may

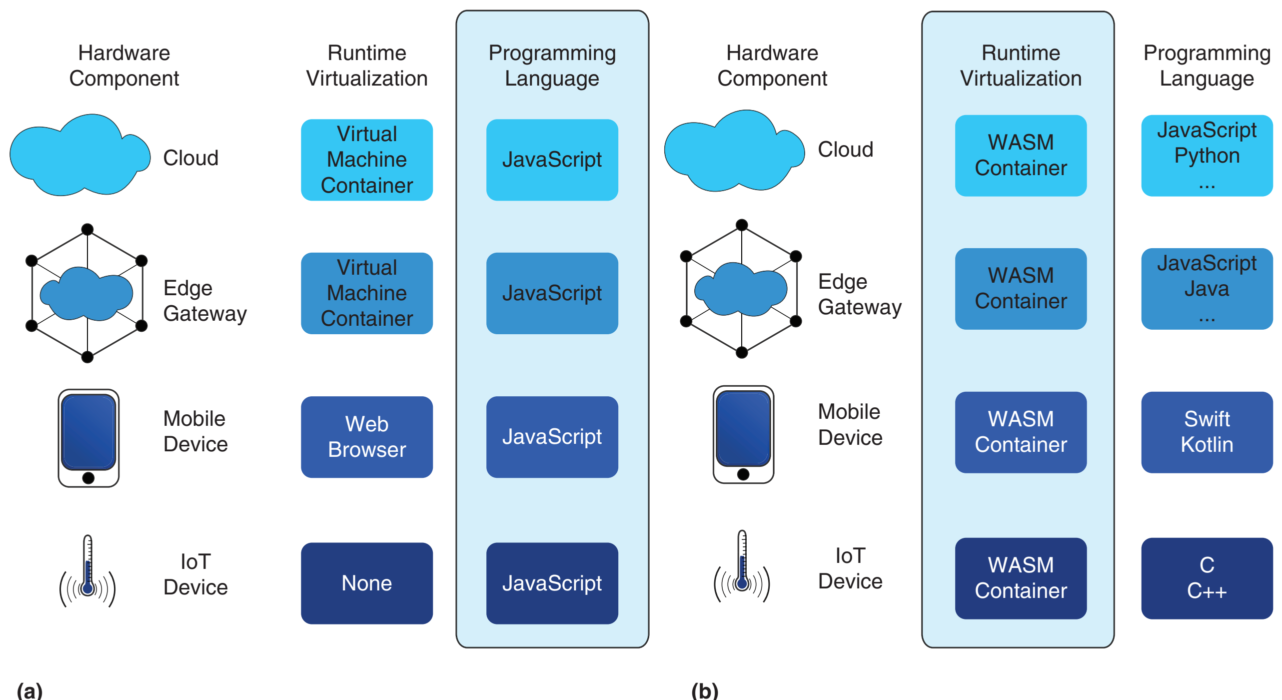

FIGURE 5. Using web technologies to implement (a) static and (b) dynamic isomorphism—the potential options. Notice that the options are not mutually exclusive.

ABOUT THE AUTHORS

TOMMI MIKKONEN is a professor of software engineering at the University of Helsinki, Helsinki, 00014, Finland. His research interests include web engineering, the Internet of Things, and software architectures. Mikkonen received a Ph.D. in computer science from Tampere University of Technology, Finland. Contact him at tommi.mikkonen@helsinki.fi.

CESARE PAUTASSO is a professor at the Software Institute at the Università della Svizzera italiana, Lugano, 6900, Switzerland. His research interests include Web engineering, liquid software architectures, and application programming interface analytics. Pautasso received a Ph.D. in computer science from ETH Zurich, Switzerland. He is a Senior Member of IEEE. Contact him at c.pautasso@ieee.org.

ANTERO TAIVALSAARI is a Bell Labs fellow at Nokia Bell Labs, Tampere, 33100, Finland. His current research interests include the Programmable World, streaming data Internet of Things systems, and the foundations of software and web engineering. Taivalsaari received a Ph.D. in information sciences from the University of Jyväskylä, Finland. Contact him at antero.taivalsaari@nokia-bell-labs.com.

require additional support for packaging the applications into containers.

In contrast, WASM programs are organized into modules, which are the unit of deployment, loading, and compilation; thus, they seem like natural candidates for building lightweight containers.[9] WASM programs can be written in a variety of programming languages and then compiled to WASM for execution. However, the technology is still relatively immature outside the realm of web browsers. Both technologies can be used to realize a model in which new applications are initialized in the locations where they are needed as well as the vision of migratory, liquid applications.[15]

Ultimately, the definition of a common, isomorphic IoT platform is about standardization. While researchers can make relevant contributions and proposals, this area requires collaboration from major industry players to get together and agree on common principles and practices. Alternatively—or in addition—de facto standards will surely be established by those companies that manage to create highly successful businesses around their IoT solutions.

Finally, while predicting the rise of isomorphic software, it is important to note that not all software needs to be isomorphic. For instance, low-cost IoT devices, such as ambient temperature or air quality sensors, are often implemented with "bare metal" solutions without including any kind of an operating system in the device. Although hardware capabilities are increasing very rapidly (just 15 years ago, who would have thought microcontrollers would be based on 32-bit architectures or have megabytes of storage memory?), we do not foresee such devices including support for containers or advanced virtualization capabilities in the next several years. In the same vein, visualization and cloud components that have been developed for monitoring the overall system state usually do not need to be transferable to run in edge devices. Even though advances in hardware development will probably eventually enable the use of virtualization literally in all types of devices, ultimately, these choices will still have to be based on rationally justified use cases rather than blindly trying to make code executable everywhere.

According to a popular saying—often attributed to Alan Kay—in software systems development, "Simple things should be simple, and complex things should be possible." Unfortunately, in modern software development, simplicity seems to be a lost virtue.[10] Instead, modern software systems are characterized by the plentiful use of virtualization, the abundant use of third-party software components from unknown sources, and a cornucopia of overlapping implementation technologies for different parts of the end-to-end system.

When targeting IoT systems, there is currently very little coherence in the development or deployment practices at the level of end-to-end systems. Furthermore, deployment in the large introduces new challenges, especially when one should routinely manage up to millions of devices in a consistent fashion. At the moment, we are still far away from Wasik's prediction: "In the programmable world, all our objects will act as one."[17] We really should not

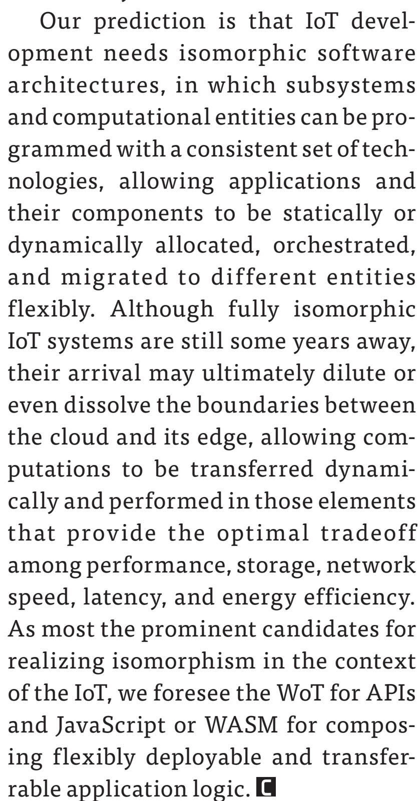

continue programming, installing, and maintaining large-scale IoT systems with the medley of technologies that is in use today.

Our prediction is that IoT development needs isomorphic software architectures, in which subsystems and computational entities can be programmed with a consistent set of technologies, allowing applications and their components to be statically or dynamically allocated, orchestrated, and migrated to different entities flexibly. Although fully isomorphic IoT systems are still some years away, their arrival may ultimately dilute or even dissolve the boundaries between the cloud and its edge, allowing computations to be transferred dynamically and performed in those elements that provide the optimal tradeoff among performance, storage, network speed, latency, and energy efficiency. As most the prominent candidates for realizing isomorphism in the context of the IoT, we foresee the WoT for APIs and JavaScript or WASM for composing flexibly deployable and transferable application logic. **C**

REFERENCES

1. K. Arnold, J. Gosling, and D. Holmes, *Java Programming Language*, 4th ed. Reading, MA: Addison-Wesley, 2005. doi: https://dl.acm.org/doi/book/10.5555/1051069.
2. D. Bryant, "WebAssembly outside the browser: A new foundation for pervasive computing," in *Proc. Keynote at ICWE'20*, Helsinki, Finland, June 9–12, 2020.
3. *ECMAScript 2020 Language Specification*, Standard ECMA-262, June 2020. https://www.ecma-international.org/publications/standards/Ecma-262.htm (accessed Mar. 5, 2021).
4. D. Ingalls, T. Kaehler, J. Maloney, S. Wallace, and A. Kay. "Back to the future: The story of squeak, a practical smalltalk written in itself," in *Proc. 12th ACM SIGPLAN Conf. Object-Oriented Programming, Syst., Languages, Appl.*, 1997, pp. 318–326.
5. D. Ingalls et al., "A world of active objects for work and play: The first ten years of lively," in *Proc. ACM Int. Symp. New Ideas, New Paradigms, and Reflections Programming and Softw.*, 2016, pp. 238–249.
6. M. Jacobsson and J. Willén, "Virtual machine execution for wearables based on WebAssembly," in *Proc. EAI Int. Conf. Body Area Netw.*, 2018, pp. 381–389.
7. D. Kurzyniec, T. Wrzosek, D. Drzewiecki, and V. Sunderam, "Towards self-organizing distributed computing frameworks: The H_2O approach," *Parallel Process. Lett.*, vol. 13, no. 2, pp. 273–290, 2003. doi: 10.1142/S0129626403001276.
8. N. Mäkitalo et al., "Action-oriented programming model: Collective executions and interactions in the fog," *J. Syst. Softw.*, vol. 157, p. 110391, Nov. 2019. doi: 10.1016/j.jss.2019.110391.
9. N. Mäkitalo et al., "WebAssembly modules as lightweight containers for liquid IoT applications," in *Proc. Int. Conf. Web Eng.*, 2021, pp. 328–336.
10. T. Margaria and M. Hinchey, "Simplicity in IT: The power of less," *Computer*, vol. 46, no. 11, pp. 23–25, 2013. doi: 10.1109/MC.2013.397.
11. T. Mikkonen, K. Systä, and C. Pautasso, "Towards liquid web applications," in *Proc. Int. Conf. Web Eng.*, 2015, pp. 134–143.
12. J. Strimpel and M. Najim, *Building Isomorphic JavaScript Apps: From Concept to Implementation to Real-World Solutions*. Sebastopol, CA: O'Reilly Media, 2016.
13. A. Taivalsaari and T. Mikkonen, "A roadmap to the programmable world: Software challenges in the IoT era," *IEEE Softw.*, vol. 34, no. 1, pp. 72–80, 2017. doi: 10.1109/MS.2017.26.
14. A. Taivalsaari and T. Mikkonen, "On the development of IoT systems," in *Proc. 3rd Int. Conf. Fog Mobile Edge Comput. (FMEC)*, 2018, pp. 13–19. doi: 10.1109/FMEC.2018.8364039.
15. A. Taivalsaari, T. Mikkonen, and K. Systä, "Liquid software manifesto: The era of multiple device ownership and its implications for software architecture," in *Proc. IEEE 38th Annu. Comput. Softw. Appl. Conf.*, 2014, pp. 338–343.
16. S. Tilkov and S. Vinoski, "Node.js: Using JavaScript to build high-performance network programs," *IEEE Internet Comput.*, vol. 14, no. 6, pp. 80–83, 2010. doi: 10.1109/MIC.2010.145.
17. B. Wasik, "In the programmable world, all our objects will act as one," Wired, 2013. Accessed: Oct. 13, 2020. [Online]. Available: http://www.wired. com/2013/05/internet-of-things-2/
18. M. Wilcox, S. Schuermans, and C. Voskoglou, "Developer economics: State of the developer nation," London, U.K.: VisionMobile Ltd., Tech. Rep., 2016. [Online]. Available: https://www.developereconomics.com/resources/reports/state-ofthe-developer-nation-q1-2016
19. A. Rossberg, Ed., "WebAssembly Core Specification," Version 1.1, Web Assembly Community Group, San Francisco, May 13, 2021. https://webassembly.github.io/spec/core/_download/WebAssembly.pdf (accessed Mar. 5, 2021).
20. "Web of Things (WoT) architecture," World Wide Web Consortium, 2020. https://www.w3.org/TR/wot-architecture/Overview.html (accessed Mar. 5, 2021).