Column: From the Editor-in-Chief

Long Live the IoT

Marc Langheinrich
Università della Svizzera italiana

***Abstract*—Connected devices are increasingly popular (think "smart fridge"), yet it is unclear how "smart," IoT-filled environments will fare in terms of longevity. What if the company that runs a device's underlying cloud service goes out of business? Will its connected car still run? Will its smart TV still turn on? Will we still be able to turn on our lights?**

I HAVE A stereo system at home that hails from the mid-nineties. While the tape deck died last year, all other parts continue to serve me admirably after over 25 years of service. While I retired the FM receiver a few years back (mostly because it would not fit into my sideboard), this, too, would in principle still[†] work. The hundreds of CDs that I collected in the last 30 years still play without problems, as do my 20 years old Mini Discs. HiFi bliss!

A few years back, I "upgraded" my trusty stereo to also play my ever-growing collection of digital music, hosted on a local NAS. All I needed was a small device called the "Chromecast Audio"—which allowed me to "cast" pretty much any type of audio from my phone to my stereo system. Unfortunately, and after only three years of service, my Chromecast stopped working on January 1st this year. When looking for buying a replacement, I learned that Google had discontinued the product some 12 months ago—by now all remaining stock has been sold and prices on eBay have skyrocket (unused ones now go for almost 10 times the original retail price). While several alternative systems seem to exist, none offers the simplicity of the Chromecast approach at the low price point of the original.

While the problem with my Chromecast Audio is most likely a hardware fault, its "death" comes at a time that the Web is abuzz with news that Sonos—a maker of (relatively expensive) so-called "connected" speakers that allow you to stream digital music from a variety of sources, and to synchronize playback across multiple such speakers across an entire home—decided to discontinue[‡] support for some of their oldest products (produced between 2005 and 2011).

There is a great deal to unpack here, but the first thing to observe is that in this case, the hardware of these "discontinued" devices is perfectly fine. The speakers of my trusty old stereo have been working for more than 25 years now, so it should be no surprise if a 2011 Sonos speaker would equally continue to do so for,

[†]At least for a few more years—Switzerland will retire analog FM broadcasts by 2024, replacing it with a digital format (DAB+).

[‡]See https://support.sonos.com/s/article/4798

Digital Object Identifier 10.1109/MPRV.2020.2985136
Date of current version 14 May 2020.

Published by the IEEE Computer Society

say, another 15 to 20 years, if not decades longer. But the software that drives them will give out way earlier.

To be fair, the fact that Sonos until recently was supporting a device that was produced some 15 years ago is quite commendable. Most "smart" devices in our day and age cease to receive software updates much earlier. And Sonos went to great length to explain customers that their old speakers would of course "continue to work"—they simply would not "get new features" anymore. Nor would they get any security updates. And neither would any other components of your Sonos family (even if they are brand new), as long as those "legacy devices" would be present in your home.

DISCONNECTED DEVICES?

Sonos speakers are far from being the only devices that stop working because of software. Anything that relies on a cloud service for its functionality is bound to become obsolete the moment the company behind it ceases operations or decides to stop supporting their "legacy" products. While we have become accustomed to three-year old smartphones being considered "legacy" (i.e., not receiving any software—and hence security—updates anymore), this becomes a lot harder to stomach as devices that have traditionally had a lifespan of dozens of years—washing machines, fridges, cars—now increasingly rely on cloud services for their functionality.

Clearly, the biggest impact of the IoT is probably in the industrial space, rather than with consumers directly, yet in both domains manufacturers can quickly underestimate the cost of having to support a product not only by offering spare parts, but also by operating servers and supporting an ageing code base. Similarly, customer may focus too much on a product's features and thus may end up ignoring questions around the manufacturer's ability to support the device in the years to come.

However, discontinuing support for a device does not have to be its death knell. When smartwatch manufacturer Pebble, Inc., ceased operations in December 2016, they released their source code to the public, thus allowing a community of developer to continue supporting the hardware. Clearly, not all IoT devices in a home will have such a strong community behind it, but thanks to the "long tail" of the Web, even obscure products may enjoy the support of a few dedicated developers to keep them alive.

DEVICE OR SERVICE?

End-of-life issues are of course well-known from software. Just a few weeks ago, on January 14, Microsoft ended its support for Windows 7—an operating system that it released more than 10 years ago. RedHat's CentOS distributions are usually seeing 10 years of support, too, while Ubuntu designates "LTS" editions roughly every 2 years, which then receive updates for 5 years. While there is in principle no reason why software should not continue to run long after its support has ended, anything that is connected to the network is bound to need continuous updates in order to patch critical security flaws that are often discovered years after software has been released (e.g., OpenSSLs "Heartbleed" bug went undiscovered for some two years even though many updates where shipped in this time).

Hardware had traditionally a much longer lifetime, with "end-of-life" being defined only by the lack of spare parts—either original parts supplied by the manufacturer, parts created by a third party, or by simply scavenging them from another (broken) device. Miniaturization has taken its toll on these timelines as well, as parts became increasingly specialized (which minimizes possible reuse across devices), integrated (effectively requiring the exchange of the entire device, as no "parts" can be exchanged), or "locked-down" (i.e., disallowing third-party parts to be used instead).

While IoT devices are usually less costly than, say, a washing machine or even a car, having to replace an entire ecosystem (because a proprietary platform is not supported anymore, or their manufacturer simply goes bankrupt) can easily incur similar costs. The IoT space with its many startups is especially prone to this "sudden-death" syndrome of manufacturers. However, even well-established vendors can quickly decide to abandon a platform, as early adopters of Intel's x86-based "Edison" IoT platform had to realize a few years ago. A similar case was Best Buy's

"Insignia" range of connected devices—smart outlets, switches, and security cameras—, which stopped working in November last year when Best Buy decided to shut down their back-end systems.

A somewhat famous case of "we are shutting down the servers for your IoT system" was the 2014 acquisition of IoT-startup Revolv, Inc., by Nest (which eventually was acquired by Alphabet, Google's parent company). In 2016, Nest decided to shut down the servers that allowed consumers to operate their Revolv Smart Hub. While a subsequent FTC inquiry cleared Nest of any "unilateral breach of consumer contracts," it did prompt the agency to draw up a set of questions to IoT device manufacturers. A few key points are worthwhile repeating here (see www.ftc.gov/news-events/blogs/business-blog/2016/07/what-happens-when-sun-sets-smart-product for the full list).

- Are you selling a device, a service, or both?
- Are consumers getting a fixed-term rental or subscription, or are they getting something they will own and can rely on for the life of the device?
- Would reasonable consumers expect to be able to keep using the device—and have it be fully functional—if the company, even many years later, rides off into the sunset?
- What did you tell consumers at the outset—or what would they otherwise expect—about the security you would provide for the life of the device?

If you are thinking of launching a new IoT service or platform, understanding how to communicate some of these key issues to consumers will go a long way in ensuring that there are fewer surprises when your startup finally does get acquired (fingers crossed)! At the same time, as a community we should start to actively think on how we can improve the interoperability between different IoT architectures in order to increase the chances that a discontinued IoT device will get a second wind. Similarly, work on "cloudless IoT" should help us minimize the reliance of such systems on Internet services. Stay tuned for our upcoming October issue on "Edge Computing" for the latest updates in this space!

PS, and in case you have not followed the latest development in the Sonos case: after the very public backlash, Sonos CEO Patrick Spence announced that the company would continue to provide security updates for their legacy products, and would also try to find a way to "split" a home system such that newer models could continue to receive new features while the older devices in a home would not. Safe for now?

IN THIS ISSUE

The title of this special issue is "Attention – Yours vs. Theirs," and Guest Editors Joe Paradiso and Dan Siewiorek selected four articles that offer insights into key challenges in our interactions with pervasive computers, such as: "How can a computer quantify how much attention a user is paying to its interface?," "How can a computer know a user's current mental capacity for paying attention?," and "How can we help users paying attention to important issues (e.g., security aspects)?"—you find more details in their Guest Editors' Introduction later in this issue.

A special Spotlight article by Gregory Abowd argues for a research push towards *computational materials*—everyday materials (e.g., carpets, asphalt, and paint) that truly have "computing" woven into them. Based on a range of projects in this area conducted at Georgia Institute of Technology, Abowd sets out a set of challenges and research directions. If you are working in this space, note our upcoming special issue on this subject—submission deadline is December 1st! You can find more details in the corresponding call for papers, available from our homepage at www.computer.org/csdl/magazine/pc.

We also have three departments in this issue: Education and Training, Smart Homes, and Wearable Computing.

The Education and Training department features a report by Christian Jansen and colleagues from Utrecht University titled "Multidisciplinary and Interdisciplinary Teaching in the Utrecht AI Program: Why and How?" Their article nicely summarizes the challenges inherent in multidisciplinary programs, as evidenced also by the previous reports by Andrew Kun and colleagues on Ubiquitous Computing education. Jensen *et al.* approach to benefit from the diversity of their

student population is inspiring and can certainly serve as a blueprint for non-AI programs that face similar problems.

In our Smart Homes department, department editors A. J. Brush and Jeannie Albrecht, together with Robert Miller, highlight smart home research from the September 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing, as well as smart home industry updates from the CES conference in January 2020.

Finally, our Wearable Computing department features a contribution by Florian Müller, Sebastian Gunther, and Max Mühlhäuser on "Around-body Interaction." Given the increasing prevalence of head-mounted displays (HMD), the authors provide an overview on how the skills and dexterity of our upper and lower limbs, acquired and trained in interacting with the physical world, can be transferred to HMD interaction.

TEAM UPDATES

In this issue, we say good-bye to our long-term Editorial Board member and department Co-Editor A. J. Brush. A. J. served for three terms on the board and was instrumental in establishing and running our Smart Homes department. I am really thankful for all her contributions to IEEE Pervasive, and glad that the department—with Editors Jeannie Albrecht and Mike Hazas—will continue to be in amazing hands!

At the same time, I am excited to welcome no less than two of A. J.'s colleagues at Microsoft Research to the Board: Dr. John Krumm and Dr. Gabe Cohn! John is a Principal Researcher in the AI lab. His research focuses on understanding peoples' location and how to use that information to benefit the user. Gabe is a Principal Researcher in Microsoft's Health Futures group, working on innovative sensing systems for novel cardiovascular medical devices. John received the Ph.D. degree from Carnegie Mellon University, Pittsburgh, PA, USA, whereas Gabe received the Ph.D. degree from the University of Washington, Seattle, WA, USA. You can contact John and Gabe at jckrumm@microsoft.com and gabe@microsoft.com, respectively.

Finally, a special "Thank You!" to all our reviewers in 2019—the updated list can now be found from our homepage at www.computer.org/csdl/magazine/pc (look for the "2019 Reviewer Thanks" item in the right-hand sidebar).

Marc Langheinrich is currently a Professor with the Faculty of Informatics, Università della Svizzera Italiana, Lugano, Switzerland, where he heads the Research Group for Ubiquitous Computing. His main research interests are usable privacy, pervasive displays, and ubiquitous computing. He received the Ph.D. degree in computer science from ETH Zürich, Zürich, Switzerland. Contact him at langheinrich@ieee.org.