

CONFIG.SYS

Marc Langheinrich , Università della Svizzera Italiana, 6900, Lugano, Switzerland

My smartphone allows me to configure not only a different ringtone for every caller, but also for every different notification of every different app. I can set times when to be silent and when to ring, or dynamically switch between these two modes based on where I am or what is scheduled in my calendar. Configuration heaven or hell?

RECENTLY BROKE MY PHONE SCREEN, and since no single repair center I contacted was able to fix it anymore (citing lack of spare parts), I had to get a new phone. Being the technology aficionado that I am, I did not necessarily mind. However, it was only after I had transferred my SIM card and booted up the new phone that I realized I was now missing several years of customizations. In fact, it was only after I started using the new phone that I realized that I had customized my old one that heavily!

KING OF CONVENIENCE

The two dominant mobile phone platforms—iOS and Android—handle customization quite differently. Android has always been about customization. From very early on, any app could provide almost any functionality of the OS by simply registering its capabilities to handle certain tasks, e.g., one could select Skype to make calls instead of your standard dialer. Similarly, Android would let one swap out wallpapers and icons, and create widgets, folders, and any number of home screens. iOS, on the other hand, got the ability to change to a custom wallpaper and create folders only in 2010 (iOS 4), while widgets had to wait another 10 years (2020, iOS 14). Many key applications (e.g., dialer, messaging) are still not customizable on iOS even today.

When I moved to my new phone, I was able to recreate much of the basic customizations I had on my old phone with the help of a “copy” app the manufacturer provided. But this really only copied over the previous apps (often causing more harm than benefit, as this can be done more reliably by simply logging into Google’s Play Store on the new phone) as well as the

basic layout of my home screen and the various folders I created. Clearly, this did save some work but missed the many detailed tweaks I had put into place over the last few years.

One set of tweaks concerned notifications. Current Android versions allow one to customize the type of notification (i.e., which ringtone and/or buzz) of each class of notifications from each app installed on the phone! Having quickly installed dozens, if not hundreds of app, and each featuring at least 2–3 classes of notifications, this quickly comes out as several hundred settings! Obviously, I had originally set these up not in a day but over the course of many months, typically in the context of some mishap (e.g., a too-loud notification or an unnecessary interruption) or prompted by some need (e.g., to not miss a message from a particular contact). None of these settings was copied over by the “copy” app, and a lack of an easily exportable list of these customizations made it impossible to methodically go through them and recreate them on my new phone. I thus again started the arduous task of individual customizations prompted by need, on a case-by-case basis. In fact, it is not even clear to me if copying over all of the previous settings would have been desirable, as notification needs change and some of my prior settings might have long since warranted a review.

Another layer of complexity comes from apps such as “Tasker”—a “customization on steroids” app that allows one to set up complex “if-then” triggers, similar to the widely popular IFTTT web service (which of course is also available on Android). I had previously created a (relatively small) set of rules that would mute my phone based on my calendar data, e.g., when I was in a meeting or giving a lecture. While Tasker allows for exporting rules, it still took me a considerable amount of work to recreate this on my new phone, as some of the references (calendar names, etc.) changed between the phone manufacturers’ Android implementations, hence all rules had to be manually updated.

Clearly, none of these customizations are essential to the basic functions of my new phone, and probably not even to its secondary functions (e.g., reading and sending mail, browsing the web), yet the sheer number of settings I had to check and readjust means I still did not fully recreate my prior smartphone experience, four months in. In “The paradox of choice,” American psychologist Barry Schwartz actually questions if having the ability to choose from so many options is a good thing, citing empirical evidence that having to make too many decisions, even if largely irrelevant (e.g., what ringtone for which contact), will ultimately lead to stress and anxiety. Maybe iOS had the better approach all along?

HOME SWEET HOME

The ability to configure and ultimately adapt technology to our needs is of course key to the vision of pervasive and ubiquitous computing. Nowhere is this probably more apparent than in the field of Smart Homes. Smart homes need by default a configuration, even if in a relatively simple “home automation” setting. Much of this might be standard behavior that is set up by the manufacturer of the home automation system (e.g., blinds opening or closing based on light levels). But thinking back on the short vignette in Mark Weiser’s original article, “The Computer for the 21st Century,” in which the alarm clock asks Sal if she wanted coffee, one can easily imagine how this might quickly become annoying in case Sal would be a tea drinker. How would the alarm clock know? Clearly, Sal would need to customize this somewhere. What if she preferred coffee on weekdays, and tea on weekends? Already, the set of options for this seemingly simple interaction (in which the alarm clock only understands “yes” and “no”) seem amazingly complex—and this is only about one beverage!

Maybe the smart home could learn this from Sal’s behavior, yet clearly one would not want to have to wait for the system to collect, say, a few hundred data points! While deep learning has revolutionized machine learning, its reliance on large amounts of data might require us to find a way to “jump-start” these systems based on some general understanding of people’s breakfast routines, weekend/workday periodicity, etc.

BETTER SAFE BUT SORRY?

I do not live in a Smart Home (yet), but I recently decided to install a “firewall appliance” in my home network, given the increasing number of “Smart Devices” (e.g., oven, dishwasher, and laundry) connected to my home Wi-Fi. Writing this, I just realized that

I might already live in a Smart Home after all! I would still claim I do not, as none of its lights or blinds can be controlled, but I might be simply getting the definition of what a Smart Home is wrong.

While I do teach a course on information security in our Master’s program, the idea of running a professional-grade firewall did not exactly fill me with joy. As any system administrator can probably testify, end-users can quickly become vicious when basic Internet services do not work, or key emails get falsely identified as spam. I thus opted for a “dumbed-down” commercial solution with a user-friendly mobile app that allows me to quickly create basic rules, device groups, and custom subnets.

It might surprise no one that the first weeks and months were rough. It took hours of tinkering for every single change that I wanted to try until I had sufficiently understood the configuration options I had to put in place—even if I only had a (seemingly) simple effect in mind. To add to the challenge, I could of course only experiment with these settings when the network was not in use by the family—which ruled out scheduled online meetings (lots of these, at the oddest hours!) or when the kids wanted to stream something. While I have since arrived at a stable configuration, I am still far from my initial plans for securing my home network. Yet any improvement in security is probably already a good step, and I am rewarded whenever I inspect the log of blocked port scans, which have definitely increased lately.

Securing a home network is clearly an exercise in selecting from many, many options, but it is of course decisively different from the “Paradox of Choice” that Schwartz describes in his book. The challenge here is the complexity of the system that needs configuration, the many moving parts (i.e., devices) that are typically not well understood, and which might react in unforeseen ways once a certain configuration is made.

While automated configurations with the help of AI can of course be seen as the ultimate solution to this problem as well, a more realistic approach might be the standardization efforts currently underway in the Smart Home space, in particular the “Matter” standard. Matter promises to not only improve interoperability between smart home products but to also improve security and privacy. Apart from the obvious (e.g., mandating secure communication between devices), Matter, for example, employs a blockchain-based ledger to verify if a newly joining device has been certified and runs the latest firmware. While all of this might not alleviate the need for me to run a firewall appliance, it would still greatly simplify the task of securing my home network against “rogue” IoT devices.

CONSUMER CHOICE

Whether it is smartphones or home networks—the abundance of settings that can (should? must?) be made by end-users—both tinkerers like me and those less so inclined—seems ever-increasing. A case in point: I have barely setup 20% of what I could possibly configure in my car, and it is not even an electric one (which are obviously computers on wheels). I certainly have “customized” my mirrors and seat position—though these need reconfiguration anytime my wife takes it for a drive—but I have also paired my phone and programmed my home address into the Satnav.

Yet, I perpetually seem to postpone the remaining 80% (Speed dial? Favorite radio stations? Optional online services? Google car home screen?) to “another time.” Similarly, I still have not bothered to learn about the many improvements in my recently upgraded Windows 11 system (Focus sessions? Desktops? Widgets? Keyboard shortcuts?), relying instead on habits formed many years ago when I still enjoyed “playing” with my Operating System.

User attention is one of the scarcest resources of our age, so it comes as no surprise that spending anything over a few minutes to configure a system only appeals to a very small crowd of enthusiasts. Yet so far there seems no significant countertrend to the ever-increasing number of configuration choices that appear in our lives.

I have written here before about the work by Dourish and Bell on “Mess and Mythology in Ubiquitous Computing,” and how their work puts “the finger on the wound” of the messiness and haphazardness of modern technology that requires us to constantly tinker with it—or be subject to the defaults and limitations put in place by whoever set up the initial system. It is unclear to me in which direction we must push our efforts to mitigate this ever-growing need to attend to our “smart” systems. Do we need more guided setup schemes? Do we need more built-in deep learning that can configure services automatically? Would I even want this? Would this even work?

I am still amazed at how often my phone’s auto-completion feature suggests totally out-of-context words, even though I have been using it for years and must have amassed a sizable “corpus” in my account (then again, maybe all this is only stored locally and thus got wiped out when I switched phones?). Similarly, the automatically suggested destinations in Google Maps that appear whenever I fire up Google Auto are almost never where I want to drive in that moment, even though Maps has been tracking my movements now for over 10 years. Did I mention that

almost none of the music nor the movies recommended to me by my respective streaming services hit home? Manual configuration might still be needed for a while.

IN THIS ISSUE

The title of this special issue is “Mental State, Mood, and Emotion,” and Guest Editors Gavin Doherty, Mirco Musolesi, Akane Sano, and Thomas Vaessen present no fewer than seven articles that discuss novel approaches, opportunities, and challenges for developing pervasive computing technology for mental health and emotion. You can find more details in their Guest Editors’ Introduction later in this issue.

We also have an additional feature article in this issue. In “A Cloud-Based Deep Learning Framework for Remote Detection of Diabetic Foot Ulcers,” authors Bill Cassidy, Neil D. Reeves, Joseph M. Pappachan, Naseer Ahmad, Samantha Haycocks, David Gillespie, and Moi Hoon Yap present a mobile cloud-based framework to support remote medical diagnoses of foot ulcers. While their work includes a straightforward use of deep learning techniques, the authors also focus on the importance of making their system usable. Their discussion of these additional challenges nicely illustrates the need for a wide evaluation scope of such systems.

We also have three departments in this issue: Conferences, Smart Homes, and Wearable Computing!

In our Conferences department, authors Rahul Majethia, Shreya Ghosh, Hannah Nolasco, Farhana Shahid, Varun Viswanath, Ibrahim Shehi Shehu, and Yiran Zhao describe their experiences in attending the 2021 edition of Ubicomp/ISWC—a fully virtual conference held in September 2021. If you missed attending the event, the article should provide you with a rich set of pointers to interesting follow-up reading! This article is complemented by our Wearable Computing department, which has authors Daniel Roggen, Katia Vega, Cindy Hsin-Liu Kao, Troy Nachtigall, Jamie Ward, and Michael Beigl present the “second half” of Ubicomp/ISWC, i.e., the 25th edition of the International Symposium on Wearable Computers.

Our Smart Homes department finally sees authors Sean Barker and Dylan Parsons write about “Smart Homes or Real Homes: Building a Smarter Grid with ‘Dumb’ Houses.” Their article ties in nicely with my own experience on living in a “dumb house” that, while already containing a multitude of internet-enabled appliances, would provide challenging to connect to the smart grid.

TEAM UPDATES

In this issue, we say good-bye to four Editorial Board members: Nicholas Lane, Mirco Musolesi, Heather Patterson, and Vassilis Kostakos.

Both Mirco and Nicholas joined the board in 2018 as AEICs and have served two terms (though Nic had to step back from the AEIC role a year earlier due to other commitments). Both have been a tremendous help for me in my EIC role, not only handling our review process but also being available to handle several of our highly successful special issues (Mirco handled no fewer than three of these)! It has been a pleasure working with both! Vassilis and Heather also joined at the same time. Heather has furthermore

been running the “Social Impact” department until 2020. We are still in the process of reappointing a worthy substitution for her and I am confident that we can soon start publishing this important department again! A big thank you to both as well!

MARC LANGHEINRICH is a professor in the Faculty of Informatics at Università della Svizzera Italiana (USI), 6900, Lugano, Switzerland, where he heads the Research Group for Ubiquitous Computing. His main research interest are usable privacy, pervasive displays, and ubiquitous computing. Langheinrich received the Ph.D. degree in computer science from ETH Zürich. Contact him at langheinrich@ieee.org.