Column: From the Editor-in-Chief

First Contact!

Marc Langheinrich
Università della Svizzera italiana

***Abstract*—By the time you are holding this issue in your hands (or accessed it online), your smartphone will most likely already have downloaded the new OS extensions by Google and Apple to enable contact tracing via BLE. A watershed moment for... Public health? Governmental location tracking? Privacy by design? Cross-platform interoperability?**

DP-3T. PEPP-PT. TCN. PACT. GAEN. A whole new set of acronyms are about to become part of our everyday watercooler chitchats. Of course, I mean the Teams channel named after it—not the real thing, given that many of us will continue with their home office routine for quite some time still.

After the lockdowns of March and April, May is to be the month when things are supposed to start their slow way back to normal—at least as much "normal" as we can expect in times of a pandemic. In much of Europe, shops and restaurants are slowly opening, whereas citizens are reminded to keep practicing social distancing any time they are in public. Will this be enough to avoid a second wave of infections? No one knows, hence many hope that mobile technology can help. On the one hand, our mobile phones are to help by detecting novel outbreaks (symptom tracking apps), potentially with the help of wearable devices (e.g., smart watches) that measure physiological data. Maybe more importantly though, our mobile phones are to track potential chains of infections (contact tracing apps) to allow states to better control the risk of another outbreak.*

TRACK AND TRACE

The goal of any Covid-19 contact tracing application is to limit the spread of the SARS-CoV-2 virus. Once someone is diagnosed with Covid-19, tracing data allows authorities to alert those that have been in close contact with the person that they may have been exposed to the virus, and thus should either self-isolate and/or get themselves tested. As measuring actual contact with the virus (e.g., via droplets) is not possible, proximity data allows one to approximate the likelihood that an infected person has spread the virus.

At first glance, the location data available on mobile phones offers just this information: using the absolute location of each mobile phone (e.g., based on GPS and/or WiFi fingerprinting) one can easily determine when and for how long

Digital Object Identifier 10.1109/MPRV.2020.3002116
Date of current version 30 July 2020.

*Obviously, combining both detection and tracking into the same app (e.g., as done in the Austrian "Stopp Corona" app) can improve the overall effectiveness of the process.

Published by the IEEE Computer Society

two people were in close vicinity. Little-known location-library provider X-mode shot to international fame after it posted an animation on Twitter[**] that showed how spring-break visitors to a Fort Lauderdale beach later spread out across the entire continental United States. Tracing-pioneer South Korea is using location data in its national contact tracing app, as do the governments of India and Iceland, as well as several U.S. states, such as North Dakota[***] and Utah.

Having detailed location data on both confirmed and potential cases vastly improves a government's capabilities of responding to a pandemic, e.g., by enforcing localized rather than national lockdowns, or by quickly identifying hotspots (e.g., a particular shop or restaurant). It also allows epidemiologists to better understand how the virus spreads—in almost real time. The obvious downside of such data is citizen privacy, even if the app uses only pseudonymous identifiers, as deanonymization is usually easy with secondary data sources. A second shortcoming of location-based tracing is the lack of accuracy, especially in built-up urban areas, where location data can easily be dozens of meters off. While many phone network providers have offered cellular location data to national governments, this (anonymous) data has mostly been used to assess lockdown adherence, rather than perform contact tracing.

A better way to measure the distance between two devices (and hence their owners) is Bluetooth Low Energy (BLE), which offers distance measurements using signal strength data from its radio module. The basic idea is that each phone periodically emits *beacons* that can be received by close-by phones, which can then calculate the distance to the emitting phone from the signal strength. Many location-based tracing apps mentioned above in fact use BLE to detect contact, whereas location data are used to understand where in the country the contact happened.

Yet, even the higher precision of BLE leaves much uncertainty when it comes to judging if two people where close enough to potentially spread the virus. First, radio waves travel through walls, hence two people sitting back-to-back with, say, a thin plywood wall between them would likely register as a "close contact" on BLE. Similarly, radio waves can be absorbed (e.g., by human bodies) or reflected (e.g., by nearby walls), thus making the signal appear weaker (and thus more distant) or stronger (and thus closer) than in reality. Using a range of real-world scenarios (e.g., shopping in a supermarket, sitting in a train, or around a meeting table) Leith and Farrell[†] recently confirmed that phone placement (e.g., back pocket vs. bag) and dense environments (e.g., in a supermarket or train) significantly affect the BLE-measured distance.

CONT(R)ACTED!

With the lack of measurement reliability in mind, most contact tracing apps attempt to minimize false positives by setting relatively high "exposure" standards. The Swiss contact tracing app "SwissCovid"—one of the first apps to use the soon-to-be-released Google/Apple Exposure Notification protocol (GAEN)[‡]—e.g., only records a "contact" when detecting a distance of less than 2 m for at least 15 min (cumulative across an entire day though). While this may miss a substantial number of transmissions, it makes alerts that the app gives significantly more trusted.

Yet, how do these apps go from these "contacts" to an eventual "you are possibly infected" alert? The basic challenge with SARS-CoV-2 is that one can be infected (and infectious!) without showing any symptoms for several days. By the time one is actually feeling sick and going to the doctor, and eventually getting test results back, a week or more might easily pass. The app thus needs to go back in time and alert everyone the person has been in "contact" with during this time

A location-based tracking system where each user transmits their location to a central server obviously can calculate "contacts" in the back-end once a user has been identified as infected. A proximity-based app (e.g., using BLE) instead

[**] https://twitter.com/TectonixGEO/status/1242628347034767361

[***] North Dakota is planning to shift its Care19 app to a decentralized model (based on Google and Apples new OS extension).

[†] Leith, Farrell (2020). "Coronavirus Contact Tracing: Evaluating The Potential Of Using Bluetooth Received Signal Strength For Proximity Detection". Tech Report, Trinity College Dublin. May 6, 2020. https://www.scss.tcd.ie/Doug.Leith/pubs/bluetooth_rssi_study.pdf

[‡] See https://www.apple.com/covid19/contacttracing/

creates lists of "contacts" in a decentralized fashion—these then must be compared with those IDs that have been diagnosed with Covid-19. In order to avoid that third parties can easily track the BLE beacons from such an app (e.g., a coffee shop simply listening to beacons in order to identify repeat customers), beacons are usually updated periodically (e.g., every 15 min). Here, each random-looking beacon is cryptographically derived from a single "master ID" that is generated upon installation of the app.

GRAND CENTRAL

While proximity-based tracing apps are inherently decentralized, the matching of "infected" IDs to contacts can be done either decentralized or centralized. In a centralized setting, such as favored by the pan-European PEPP-PT[§] initiative (with backing, e.g., from France and the UK), a central server stores each app's "master ID"—allowing it to easily derive any beacon ID that the app has sent out in the past. An "infected" app thus only needs to upload all of its "contacts" that it saw during its last few days[§§]—the server can use the "master IDs" it has stored to compute the necessary matches, and then inform those who have been in contact with the infected person.

In the decentralized setting, as advocated by the Swiss DP-3T protocol[§§§] (which in turn inspired GAEN), each mobile phone keeps their "master ID" to themselves. In case of an infection, the "infected" app uploads its *own* beacons that it sent out in the last 2–5 days. As the server does not know who has been in contact with these beacons, it instead adds these to an "infected beacons" list it keeps for other participants to download. Apps in a decentralized setting periodically poll this list of "infected beacons" from the central server and compare them to the beacons they encountered. Users thus detect locally if they have been exposed to an infected person—the server never knows who has been in contact with each other.

Proponents of the decentralized approach argue that trusting a central server with all participants' "master ID" would allow governments to track citizens and their social networks, even with no location data and personal data collected (e.g., by correlating with additional databases). Those that favor a centralized approach instead argue that it is more efficient[¶] (imagine millions of clients constantly having to poll huge "infected beacons" list) and, in fact, more privacy-friendly, as data about infected people is only sent to those who were actually in contact with them (remember that the decentralized approach distributes the lists of "infected" beacons to all participants). The authors of the Singaporean (centralized) contact tracing app BlueTrace—one of the first contact tracing apps in the world—furthermore argue that a centralized approach allows for human contact tracers to get in touch with potentially exposed users, a fact they consider vastly superior to any automated notification system when it comes to limiting false positives.

BACK TO LIFE, BACK TO REALITY

As we move into June, contact tracing apps will become a reality in much of the world. A recent MIT Technology Review Online article by O'Neill *et al.*[¶¶] attempts to track the slew of apps already released or soon to be available: it counts no less than 25 apps worldwide. There is tremendous hope among officials that these apps will eventually help turn the tide and allow the world to avoid a second wave of infections. However, whether these efforts will payoff will depend on three key factors: reliability, user acceptance, and interoperability.

Interoperability: While current lockdown regimes have brought international travel to an almost standstill, Europeans are cautiously beginning to make holiday plans for the

[§]PEPP-PT stands for "Pan-European Privacy-Preserving Proximity Tracing". See https://www.pepp-pt.org/

[§§]The SwissCovid app currently assumes a contagious window of 5 days, while the Austrian StopCovid app uses 54 h

[§§§]DP-3T stands for "Distributed Privacy-Preserving Proximity Tracing". See https://github.com/DP-3T/

[¶]The May-25 White Paper by the DP-3T project estimates some 4-25 MB of daily downloads in their most privacy-friendly scenario (providing unlinkability) while less than a MB in a more optimized (but still pretty privacy-friendly) version that uses intermediate daily keys to generate the beacons. Here, the infected app only has to share the one daily key from when the app owner was considered infectious (say, the key used 5 days ago) - other apps can then derive all subsequent daily keys (and all beacons) from this one key.

[¶¶]O'Neill, Ryan-Mosley, and Johnson (2020). "A flood of coronavirus apps are tracking us. Now it's time to keep track of them." MIT Technology Review, May 7, 2020. See https://www.technologyreview.com/2020/05/07/1000961/launching-mittr-covid-tracing-tracker/

upcoming summer months again. With half-a-dozen national tracing apps being deployed across Europe soon, it is imperative that these allow for some level of interoperability, otherwise foreign visitors would neither be able to trigger a warning should they later be diagnosed with Covid-19, nor could they be notified if they were in contact with an infected person. While centralized tracing systems such as PEPP-PT have interoperability high up on their feature list (it is called "Pan-European" for a reason), the recent push toward decentralized solutions all over Europe has seen initial work[¶¶¶] to also allow decentralized systems, e.g., based on GAEN, work across national notification systems.

User Acceptance: An early model[#] by epidemiologist Christophe Fraser and colleagues at the University of Oxford suggested that contact tracing apps could "stop the epidemic if some 56% of the population (or 80% of all smartphone users) use the app."[##] This seems unattainable in most, if not all countries in the world, given that even tracing pioneers, such as Singapore or Norway, have only seen app uptake of around 20–25%.[###] Privacy issues may certainly play a role in this, but equally significant seem to be the implications of receiving an "you have been exposed"-alert. Current guidelines in Switzerland, e.g., where a national app should be rolled out in a few weeks, suggest that if you get alerted by your app, you simply "self-isolate" rather than getting tested (as tests are still being reserved to those showing symptoms). Yet, self-isolation does not mean you are sick—you will need to work out for yourself how to convince your employer that you will not come to work, as there is certainly no right to get paid during self-isolation (unless you can work from home, that is). If such alerts are furthermore relatively frequent, people may soon decide to uninstall the app to avoid such moral dilemmas.

Reliability: Given the challenge of reliably measuring distance using BLE (and even more so using location data), any contact tracing app will need to be very conservative in its assessment what constitutes a "contact." It remains to be seen how effective these apps can be given that there is increasing evidence of having a very high variability in infectiousness: some "superspreaders" are able to infect hundreds in a single day, whereas others live in close vicinity, e.g., with their family, without ever infecting any of them. Proponents of a centralized model argue that centralization makes fine-tuning (by epidemiologists) much easier and more flexible, even adjusting to regional differences, whereas a decentralized approach would have to run a one-size-fits-all model that may be too sensitive in some areas while too coarse in others. The Austrian "Stopp Corona" app is (still) using the Google Nearby protocol to find close-by devices, which uses ultrasonic chirps for detection. This has the added benefit that these signals do not travel through walls, contrary to BLE. However, it is unclear if by switching to the soon-to-be-release GAEN protocol the need for Google Nearby (and hence the "limitation" of this detection method) will go away.

UNKNOWN QUANTITIES

Expectations on science and technology run high these days, be it for medical research to find a reliable antibody test or even a vaccine; for epidemiologists to better understand how the virus spreads; or for technologists to detect and track infections. While it is probably unavoidable that the public has high expectations, one must make sure to clearly communicate the limits of these "solutions." Proximity tracing apps are a helpful and needed tool to limit the spread of Covid-19. Yet, they are far from being a panacea. Even decentralized solutions like DP-3T will, e.g., most likely not convince privacy-concerned citizens. Also, the reality of Android fragmentation probably means that the much-needed OS updates from Google (which form the basis for efficient proximity detection across the two dominant mobile phone platforms) will initially arrive only to the miniscule fraction of Pixel devices, but not for,

[¶¶¶] Luckas et al, (2020): Interoperability of decentralized proximity tracing systems across regions. Draft Version 2.2, May 15, 2020. See https://drive.google.com/file/d/1mGfE7rMKNmc51TG4ceE9PHEggN8rHOXk

[#] See https://www.research.ox.ac.uk/Article/2020-04-16-digital-contact-tracing-can-slow-or-even-stop-coronavirus-transmission-and-ease-us-out-of-lockdown

[##] Note that epidemiologists have since criticized this 56% figure as too high of a bar. Marcel Salathé of EPFL has stated that " As soon as you have double digits, I think the effect is already quite substantial."

[###] Findlay, Palma and Milne (2020) "Coronavirus contact-tracing apps struggle to make an impact". Financial Times Online, May 18, 2020. See https://www.ft.com/content/21e438a6-32f2-43b9-b843-61b819a427aa.

say, Samsung, Motorola, or OnePlus phones. The much-needed interoperability, even if accomplished soon for, say, all GAEN-based solutions (e.g., Switzerland, Austria, and Germany), will still leave many countries with incompatible tracing solutions (e.g., UK and France)—how will this affect international travel?

With all these challenges looming, it is still heartening to see how fast the tech community has made this much progress. And while none of the tracing apps are perfect, many of the rigorous security and privacy requirements set forth by the TCN Coalition[¤] are being followed by the majority of the apps, such as Server Privacy, Receiver Privacy, and Reporter Privacy.

At the same time, this might also be the perfect moment to rethink the role of privacy protection in society. As great as it is having a GDPR[¤¤] around when so much new tracking technology gets deployed, one must always realize that just as many other rights, privacy *does* have its limits. With the current rush to implement Covid-19 tracing solutions as decentralized and anonymous as possible, we should not lose sight of the fact that privacy is only the means to an end. Both societal and individual concerns, such as health and safety, are often balanced against privacy (and rightly so). There has probably never been a time when the benefits of "big data" have been more obvious—we should make sure that in our drive to protect individual privacy, we do not forget to thoroughly consider the potential benefits (both for individuals and society) of individual data disclosures.

IN THIS ISSUE

Guest Editors Oliver Amft, Jesus Favela, Stephen Intille, Vassilis Kostakos, and Mirco Musolesi report on the latest developments in "Personalized Pervasive Health." They compiled two theme articles and a spotlight article to illustrate the variety of approaches to improve personal health through pervasive technology—you will find more details in their Guest Editors' Introduction later in this issue.

If this is not enough, we even have our regular "Pervasive Health" department in this issue! Lorraine R. Buis and Jina Huh-Yoo describe "Common Shortcomings in Applying User-Centered Design for Digital Health" and argue for a more rigorous application of User-Centered Design when developing Pervasive Health solutions. Three further departments in this issue are Education and Training, Maker Tech, and Wearable Computing.

In our Education and Training department, and fully in the spirit of Covid-19-induced social distancing, Niels Henze, Valentin Schwind, Katrin Wolf, Martin Kocur, and Albrecht Schmid share their experiences with "Preparing an Online Lecture That We Wouldn't Hate to Attend." It is a fascinating account how teaming up in times of crisis can benefit everyone—I am sure their insights will persist even after the current lockdowns have long ended.

Our Maker Tech department has Frikk H. Fossdal and Jens Dyvik present their "Fabricatable Machines" toolkit for designing and making "robust computer-controlled machines." It is a bold view into a future where our CAD programs understand the entire "supply-chain" of a conceptual design, thus allowing one to directly fabricate the machine's individual parts on a variety of fabrication tools.

Our Wearable Computing department features an article by Ismini Psychoula, Liming Chen, and Oliver Amft on "Privacy Risk Awareness in Wearables and the Internet of Things." The authors propose a framework to assess the privacy risk of sensing-based pervasive health applications, based on the concept of a "Trusted Mediator" - a private cloudlet that allows for the storage of personal data under the user's control, and which sensitizes or withholds personal data according to the policy and trust information.

This issue also contains two feature articles—both with Covid-19 related topics, no less: ICU ventilators and mobile phone proximity! In "Leveraging IoTs and Machine Learning for Patient Diagnosis and Ventilation Management in the Intensive Care Unit," Gregory B. Rehm, Sang Hoon Woo, Xin Luigi Chen, Brooks T. Kuhn, Irene

[¤]TCN Coalition (2020): The TCN Protocol. Visited on March 13, 2020. See https://github.com/TCNCoalition/TCN. See also their "Contact Tracing Bill of Rights", available at https://exposurenotification.org/

[¤¤]GDPR stands for the "General Data Protection Regulation"—a pan-European privacy law that went into force across EU member states in May 2018. See https://ec.europa.eu/info/law/law-topic/data-protection_en

Cortes-Puch, Nicholas R. Anderson, Jason Y. Adams, and Chen-Nee Chuah report on their experiences creating a clinical decision support system based on IoT sensing devices—specifically monitoring ventilator units in an Intensive Care Unit.

In "Growing Apart: How Smart Devices Impact the Proximity of Users to Their Smartphones," Jung Wook Park, Hayley I. Evans, Hue Watson, Gregory D. Abowd, and Rosa I. Arriaga question the assumption that people hardly part with their mobile phones (which is good for these tracing apps I guess). Stipulating that the plethora of other smart devices (e.g., smart watches) will allow people to substitute their phones more flexibly with other devices, they ran a multiweek study and online survey, and found that, compared to previous studies, participants increasingly relied on tablets and smartwatches, allowing them to keep their phone in another room while still being "connected" (as long as they take their phones with them when they leave the house, our tracing apps should still work though).

TEAM UPDATES

In this issue, we say good-bye to our Editorial Board members Jennifer Mankoff and Friedemann Mattern. I thank both Jen and Friedemann for their many contributions to the magazine and hope that we will be able to draw on their continued support when it comes to promoting our publication. Thank you!

At the same time, I am delighted that long-term department editor Jeannie Albrecht accepted our offer to join the Editorial Board! Jeannie is professor of computer science at Williams College in Williamstown, MA. Her research focuses on computer forensics, mobile application management, and energy management systems in smart homes. She has been editing the Smart Homes department (together with Mike Hazas and AJ Brush) since the beginning of my EIC tenure, and I am really looking forward to having her work even more closely with the team when it comes to setting the direction of the magazine! You can reach Jeannie at jeannie@cs.williams.edu.

Our second addition to the Editorial Board is Stefan Schneegass. Stefan is assistant professor of computer science at the University of Duisburg-Essen. His research focuses on the crossroads of human-computer interaction and ubiquitous computing. This includes novel interaction techniques for mobile, wearable, and ubiquitous devices as well as novel approaches for implicit authentication using such devices. You can contact Stefan at stefan.schneegass@uni-due.de. I am excited to have Stefan on-board!

Finally, we also have a new Associate Editor-in-Chief (AEIC)! Two-term Editorial Board member Junehwa Song accepted our invitation to replace Steve Hodges, who stepped down in January this year, as the sixth AEIC on the team. Junehwa is a professor in the School of Computing at the Korea Advanced Institute of Science and Technology, where he works in wearable computing, social computing, and context-and interaction-aware platforms and applications. You can contact Junehwa at junesong@nclab.kaist.ac.kr.

The current AEICs and their areas of expertise are thus as follows.

- Stephen Intille: Pervasive Health (Departments Editor).
- Fahim Kawsar: Smart Homes, Enterprises, and Cities (Community Efforts).
- Nicholas Lane: Pervasive and Mobile AI.
- Mirco Musolesi: Mobile and Sensing Systems.
- Florian Michahelles: Internet of Things.
- Junehwa Song: Social and Culture Computing.

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