Safer Together: TraceTogether and the future of digital contact tracing

A pioneer of digital contact tracing discusses the considerations that went into building the world's first national contact tracing app and ponders the future of the technology.

By Jason Bay

DOI: 10.1145/3522695
OPEN ACCESS

"Has this been done before in any other country?" I scanned the audience in the room and considered the question. Taking a deep breath, I replied, "We are aware of research projects that have explored using Bluetooth to perform contact tracing. But, no, I don't think a national effort using this technology has ever been undertaken."

You would have heard a pin drop, but for the synchronized scratching in notebooks of the gathered journalists — both local and international. I gulped. That was the moment that everything changed, when I realized what we had built over eight weeks was more than just a tool for Singapore's pandemic response.

Over the next month, as my inbox exploded with inquiries and appeals to help build similar apps for other countries, I came to realize TraceTogether and the potential of accurate and reliable digital contact tracing was more than an innovative application of technology. It represented hope and the promise of greater agency.

Launch day was supposed to have been a celebratory occasion, after 56 long days and nights of development. It wasn't. Ten percent of Singapore's population downloaded the app on that first day, busting our most optimistic projections. My team and I began to realize we had unleashed something larger than we could have imagined. Expectations were heavy—not on us alone, but also on the technology we had midwifed. As a pioneering adaptation of consumer technology for public good in an unprecedented global crisis, there was no playbook or template. Much remained unknown about the SARS-CoV-2 virus. Digital contact tracing was unproven. Would it be given the chance to learn, improve, and demonstrate its value to medical science and public health? Or would it be sunk—perhaps forever—by the weight of the unrealistic hopes that had been pinned on it?

USING TECHNOLOGY FOR CONTACT TRACING

To combat outbreaks of infectious diseases such as COVID-19, public health practitioners and contact tracers interview patients to build up a history of their activities and social interactions to identify persons they may have come into contact with—and possibly spread the disease pathogen to.

This allows them to reach out preemptively to provide medical guidance, and prevent or reduce further spread of the disease in the community. The process is time intensive and can be quite detailed. Given the private nature of the information being


SG UNITED
This belongs to
SKY
Hotline - 6973 6511
Property of Singapore Government

Figure 1. Mobile device test farm, comprising handsets procured or on loan from manufacturers.

Image courtesy of Team TraceTogether

divulged, a great deal of trust between contact tracers and patients is necessary. Nevertheless, contact tracing works—and has worked for decades—especially against diseases for which there is no readily available cure.

Unfortunately, our memories are fallible—can you remember what you had for lunch last Tuesday? Much of the difficulty in managing an outbreak in its initial stages arises from insufficiently exhaustive or timely contact tracing, resulting in undetected chains of transmission.

The idea behind TraceTogether is simple. Having the ability to reliably log such encounters would be a boon to both infectious disease specialists and society. TraceTogether and other Bluetooth-based proximity recording systems work by exchanging messages or "chirps" between participating devices at regular intervals. These messages or "chirps" and the obfuscated identifiers they contain are then logged and can be used to recreate incidences of close contact exposure if a user is subsequently diagnosed.

When Singapore reported its first case of COVID-19 in January 2020, my team and I resurrected an idea from 2015, when the Middle East Respiratory Syndrome (MERS) outbreak was spreading, and set to work building a prototype contact tracing app.

A MULTIDISCIPLINARY ENDEAVOR

We soon realized that there were several unique challenges involved in building such an app. After all, we were proposing to use Bluetooth in a manner that it was not explicitly designed for.

Although my team had extensive software development experience, we did not have deep radiofrequency (RF) engineering expertise—with the sole exception of my college training as an electrical engineer. Thankfully, we were able to tap on both public and private sector expertise in cryptography, security, privacy, and hardware sensors. In designing our protocol, we also reached out to teams who had piloted similar systems to understand their considerations—especially around protecting identities and managing consent. But the key gap in our knowledge was in the Bluetooth radio space.

WORKING WITH BLUETOOTH

Bluetooth is more than two decades old. While it has been used for peer-to-peer communication and for location-sensitive applications, it has not been used on a large scale for distance and range estimation.

Today's Bluetooth-based digital contact tracing apps rely on signal strength to estimate distance between devices (and by extension their respective users). While most Bluetooth hardware can return a relative signal strength indicator (RSSI) value to apps in approximate dBm units, we must work around several issues affecting the accuracy of ranging.

Differences in transmission power across devices. Relying on received signal strength to estimate distance is analogous to listening to the volume of a person's voice and inferring one's distance from the speaker. It is assumed the loudness of the speaker is known, within a certain margin of error. While this may be a reasonable assumption for people's voices, it is not true for mobile devices. Differences in hardware implementations across mobile handsets result in significant variation in the transmission power—up to three orders of magnitude or a thousandfold difference.

To address this variability, we had to test and characterize the transmission signal strength of numerous mobile handsets on an ongoing basis. This meant building up a "farm" of more than 100 mobile test handsets, representing more than 95% of the handset ecosystem in Singapore. We also secured the use of an anechoic electromagnetic test chamber to perform these calibration measurements in free-space, controlled setups without interference from reflections or other devices. To the best of my knowledge, prior to TraceTogether, there had been no attempts to catalog transmission signal strength across mobile devices on a similar scale.

These measurements allowed us to compensate for the differences in transmitter power, to better estimate distance between devices.

Environmental factors affecting signal propagation. At the same time, even if the signal strength of a transmitter is known, RSSI is not a good estimator for distance for several reasons.

Firstly, a key simplifying assumption in using Bluetooth for ranging is that the inverse square law holds, that is, the propagation of a radio signal is isotropic—symmetrical in all directions. However, this assumption often does not hold.

Metallic structures can sometimes function as waveguides, propagating radio waves in specific axes with greater intensity than would be expected.

Differences in handset design and materials also mean that transmitted signals do not emanate with equal intensity in all directions.

Second, we also have to consider uneven attenuation. Radio waves reduce in strength when passing through objects, such as human bodies, bags, and walls. This can result in signal strength diminishing by a factor of 100 or so.

To compound this, we also have to accept the possibility that RF signals will penetrate materials that can effectively block the spread of aerosols or droplets that can carry SARS-CoV-2 virions—such as glass panels or walls made of radiotranslucent material—giving rise to false positive estimates of proximity.

Third, there are also multipath effects to contend with. Ground surfaces, as well as other surfaces like walls and ceilings can also give rise to reflections that interfere both constructively and destructively with the received signal.

Given the difficulties in compensating for the effect of these factors, it became important to take regular and frequent (at least once a minute) samples of RSSI in hopes of mitigating these environmental factors and correct for different handset characteristics. We did this during post-processing and aggregation of the individual chirps into a contact event.

MODELING EFFECTIVENESS

Ultimately, the single most important factor determining the effectiveness of a digital contact tracing system is not the technology itself. After all, with enough data, the precision and recall of such a system can be quantified and adjusted for. Parameters such as the thresholds for both duration of exposure and distance estimates can and must be calibrated to bias the system towards either greater precision (fewer false positives) or greater recall (fewer false negatives) to suit the context in which it will be operated. TraceTogether was meant to work hand in glove with public health, rather than as a parallel, independent effort.

Adoption and effectiveness. C*eteris paribus*, adoption is probably the most crucial driver of system effectiveness, because the probability that a contact tracing app can detect interactions between any two randomly chosen individuals is limited by the square of the adoption rate of the app. When one in three individuals use a contact tracing app, the "coverage probability" is 1/9; when two in three (or twice as many) individuals use the app, the coverage probability increases fourfold to 4/9. When we think of success measures for a typical app or product, the more effective the app is at addressing its users' needs, the more likely people are to download and use the app. We might then rely on adoption and usage statistics to measure the success of the app.

However, with contact tracing apps, the causality between adoption and effectiveness is reversed; effectiveness does not drive adoption—adoption drives effectiveness. The more people use the system, the more effective the system is for everyone.

An app that sees 30–40% adoption (which is typical for a national digital contact tracing or exposure notification system today) will only capture somewhere between 9-16% of pairwise interactions in the general population—and a similarly low proportion of contacts that should be notified of exposure; this belies the view that digital contact tracing systems can replace more traditional contact tracing processes.

Using TraceTogether would help us all protect ourselves, protect our loved ones, and protect our community.

Figure 2, Performing characterization testing of handsets in an anechoic electromagnetic test chamber.

Image courtesy of Team TraceTogether

TRUST AND ADOPTION

This then is the central challenge of contact tracing apps: We have to persuade people to use an app that—for the most part—does not yield any immediate tangible benefit to the user. Any benefit in the form of early notification of exposure to the coronavirus occurs as an incidental outcome of another user's use of the app.

What does seem to drive adoption though, is an appreciation of both the benefits to the individual user and to the community at large. In our initial product marketing, we focused on a three-pronged message: Using TraceTogether would help us all protect ourselves, protect our loved ones, and protect our community.

Trust in institutions, trust in technology. Ultimately, faith and trust in the institutions and organizations deploying and using the digital contact tracing system is as important, if not more so, than an understanding and intrinsic trust in the technology itself. After all, people and institutions have motives; technological solutions do not.

Comparisons of adoption rates for contact tracing apps and systems globally appear to bear this out. At a digital contact tracing conference hosted by the Massachusetts Institute of Technology in October 2021, participants representing a cross-section of the stakeholders in digital contact tracing noted governmental and public health authority endorsement of and support for the use of digital contact tracing tools made a significant difference—as might be expected—to the adoption and use of these tools.

Our role was to ensure that this trust in the institutions that keep us safe was not squandered by building adequate privacy and security safeguards into TraceTogether.

While trust in institutions can help promote use of the digital contact tracing tools, a lack of trust in the technology can dampen adoption and use of the app—on the part of health professionals as well as lay users.

Do no harm (to public health). As one of many tools in the toolkit of infectious disease practitioners, digital contact tracing should be a supplement to human-led contact tracing (especially at lower levels of population adoption). It should not be positioned as a replacement for the need to undertake testing, tracing, and isolation of infected individuals, with humans in the loop [1]. This means ensuring that product design choices are fit for purpose, from a public health perspective.

To that end, we worked closely with public health officials right from week one. To avoid generating unnecessary noise and consuming already stretched public health resources, we initially biased the system away from false positives, based on feedback from our key stakeholders.

This meant eschewing technological triumphalism in promoting the use of TraceTogether and taking a broader perspective of its role. Contact tracing processes in Singapore had been built up following our experience with SARS in 2003 and were not fundamentally broken. TraceTogether was to be a supplement to traditional contact tracing, not supplant it.

By erring on the side of limiting the yield of close contacts from TraceTogether, we avoided generating large numbers of false positives and creating a "pingdemic." From a technological perspective, this meant scaling back ambitions and curtailing a natural impatience to demonstrate the potential of the technology—especially when such a parochial approach would have risked swamping public health systems. In short, we prioritized the ends (pandemic response operations) over the means (proving the value of the technology in the short term).

We were also able to run *ex post* analyses of correctly and incorrectly identified close contacts. Over time, this built up confidence on the part of public health officials that the system was working as intended. With this trust, we were able to subsequently expand TraceTogether's use beyond identification of close contacts, to provide tiered alerts to users, depending on their risk level. Without this closed loop feedback, we would have risked inundating public health officials with the very tool that was supposed to support their lifesaving work and breached a principle we held dear—"first, do no harm".

Privacy. Where lay users are concerned, a potential dampener on public adoption of the app is users' privacy concerns. People do not decide to use a bad product simply because it is

Figure 3. Sketch illustrating the differences in the relationship between signal power and distance, after factoring in multipath effects (from reflections off ground surfaces, walls, and ceilings).

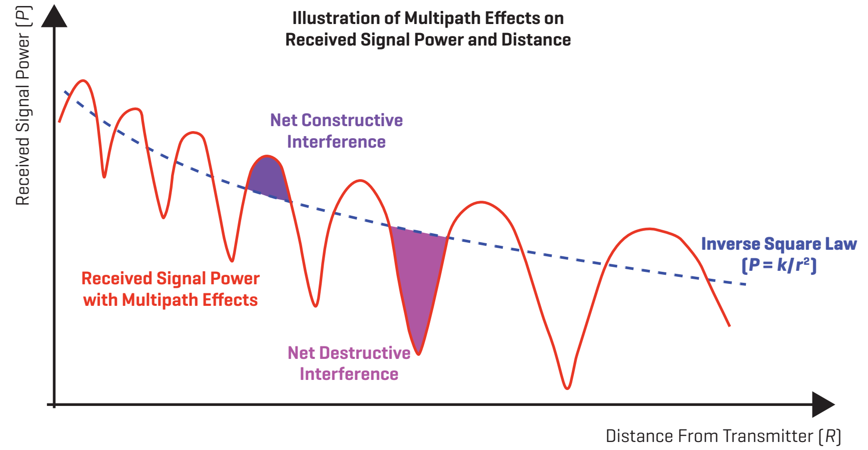

Image courtesy of Jason Bay

privacy preserving; however, they can be deterred from using an otherwise good product if it is callous with privacy sensibilities.

Much has been written about privacy and security considerations. Digital contact tracing systems in countries such as the United Kingdom and Germany have been redeveloped from scratch to reflect the privacy sensibilities and socio-political contexts of the communities they serve. It is important to be sensitive to the nuances in the social milieu around privacy in various jurisdictions.

Recognizing this, TraceTogether was built incorporating input from privacy and data security professionals. Without losing sight of the *raison d'etre* for contact tracing, we designed the system to decentralize the recording and storage of proximity logs on users' own devices (or on a non-connected hardware device called the TraceTogether Token), so that data is processed and information on a user's close contacts is made available to contact tracers only when the user is diagnosed with COVID-19 and consents to upload their data. No other proximity data is available to the public health authorities.

However, we chose not to deliberately withhold critical information that would allow public health personnel to piece together transmission chains. Doing so would prevent TraceTogether from functioning as a complement to contact tracing and epidemiological research and limit the system to notifying users of potential exposures with no feedback loop to public health. That would risk placing the cart before the horse, and mistaking privacy as the ends rather than the means.

Figure 4. Government Digital Services team members on March 21, 2020, a day after TraceTogether launched.

Image courtesy of Team TraceTogether

THE SINGAPORE EXPERIENCE

Use of TraceTogether was effectively obligatory at most public places in Singapore by the end of 2021, and adoption was virtually pervasive [2]. However, prior to TraceTogether being a requirement for entry to public spaces, by December 2020 voluntary adoption had already risen past 70% in response to exhortations by the government to participate in the system as Singapore prepared to relax restrictions on gathering and daily activities. (Note: Even at 70%, only about half of all patients' contacts would have been captured.) This is probably a reflection of the high levels of social capital and trust in institutions—governmental and non-governmental—that exist in Singapore; the annual Edelman Trust Barometer for 2021 ranked Singapore in the top five globally.

Against this backdrop, our role was to ensure that this trust in the institutions that keep us safe was not squandered by building adequate privacy and security safeguards into TraceTogether. These include ensuring that "chirps" do not contain persistent information that can reliably identify an individual, by rotating these identifiers several times an hour, and decentralizing the storage of data on users' own devices, so that the only proximity data that is accessible to public health institutions is that of users who consent to upload their data upon a positive COVID-19 diagnosis.

Although not everyone who became diagnosed with COVID-19 agreed to upload their proximity logs so that their close contacts could be alerted, enough did so such that by early 2021 about half of all close contacts (the exact proportion rises and falls due to changing circumstances) that would have been identified through human-led contact tracing were also identified by TraceTogether. Critically, these at-risk individuals were notified within minutes of an upload—potentially days faster than a wholly manual process. This has allowed the system to contribute significantly to Singapore's pandemic response, resulting in a shorter time elapsed before isolation of at-risk contacts, truncation of transmission chains, and low death rates. Singapore's excess mortality was negative in both 2020 and 2021, and overall death rates in both 2020 and 2021 were lower than in 2016–2018 [3].

When concerns arose as to the use of TraceTogether data for criminal investigations, the government responded decisively to pass legislation to circumscribe the circumstances under which TraceTogether data could be used. These guarantees are *sui generis*, going beyond confidentiality assurances for information about financial transactions, banking information, social media and instant messaging content, and personal correspondence—all of which are subject to properly served warrants.

More background on the Singapore TraceTogether experience has been captured in case studies written by

Harvard Business School [4] and the Singapore Management University [5].

WHERE DO WE GO FROM HERE?

The COVID-19 pandemic has now been raging for over two years. When will it end?

I know Team TraceTogether and I continue to hold our breath, waiting for the public health authorities to stand down contact tracing and exposure notification operations. After all, when we launched TraceTogether, we had committed to disabling the system when contact tracing ceases.

Turning off the systems will be a watershed moment. It will signal the end of this pandemic chapter in our shared experience.

But the work should not stop there. When my team and I adapted this technology for a public health purpose, we did not—in our wildest dreams—believe it would see widespread adoption and effectiveness during this pandemic. After all, our closest experience had been SARS in 2003, and that had only lasted a few months. As a result, our ambition was merely that we demonstrate a proof of concept, collect some data, exercise the process of integrating our data into public health systems, and after the outbreak ended in a few months—or so we thought—mothball the technology and concept in preparation for future disease outbreaks. We were not interested in creating a flash in a pan, without any lasting impact.

Stewarding the potential of a novel technology. In short, we wanted to steward and safeguard the future potential of the technology. Arthur C. Clarke famously said, "Any sufficiently advanced technology is indistinguishable from magic." Now, there is nothing magical about digital contact tracing technologies, but it is natural to be wary of what we are unfamiliar with. The relatively low adoption rates (and low effectiveness) of many such apps underscores the importance of trust, not only in the underlying technology, but also in the entity deploying the technology.

If our societies lose trust in the technology because of how it was deployed for COVID-19, it will ultimately be to all our detriment in the long run. In some sense, we are all ultimately stewards. Our values, our design, and our implementation choices, along with those of other digital contact tracing teams and stakeholders around the world, will shape the future of this technology and its potential to serve the public good for future infectious disease management.

Stewarding future technological options. In terms of the underlying technology, the pioneering contact tracing apps essentially improvised on existing functionality within mobile handsets before Apple and Google (and Huawei through its Contact Shield APIs) announced proprietary OS-level features to support contact tracing functionality.

Future contact tracing apps and systems should not, however, be subject to the same constraints. The Bluetooth standards are evolving to incorporate high accuracy distance measurement (HADM), which would alleviate some of the existing issues that plague the use of Bluetooth. Alternatively, future systems might not even be based on narrowband Bluetooth, with its inherently error-prone distance estimation capability. Instead, ultra-wideband (UWB) radio—which is inherently capable of sub-10 cm accuracy—could be used to provide more accurate distance measurement. Even more speculatively, some developers have proposed using ultrasound to measure distance. However, while audio signals are promising and their use is available on a wider range of devices than Bluetooth, the processing of audio from a user's environment would certainly raise a distinct set of privacy concerns.

Giving public health a greater say. Ultimately, the design and deployment of such systems have to be informed by public health expertise and the evolving body of knowledge of each novel pathogen. Historically, our experience with SARS led to the use of extensive thermal screening for febrile, symptomatic patients. However, the prevalence of SARS-Cov-2 asymptomatic and presymptomatic transmission meant that public health efforts had to rely more heavily on contact tracing. Even so, the parameters used to define and identify risky exposures have had to evolve.

When droplets were believed to be the primary mode of transmission, digital contact tracing tools that flagged prolonged proximity seemed ideal; when the evidence in favor of aerosol transmission began to accumulate, being in the same space (and breathing the same circulating air) as an infected individual became increasingly important as a proxy for exposure risk, leading to various countries deploying QR code-based check-ins. In some cases, the rationale was simply that an adoption-sensitive system did not generate enough contacts or provide information to epidemiologists hoping to understand how the virus spreads in different types of venues and environments, and when people participate in different activities.

This is an area that developers of contact tracing technology need to do better at, for the next pandemic. Having technologists primarily guiding the development of a contact tracing solution, and determining how it can be deployed, creates a bias toward "tech push." This came to a head in April 2021 when an update to the U.K.'s National Health Service app to allow consenting users to upload a historical log of venues they had visited, was blocked by Apple and Google for breach of their policies surrounding the use of the proprietary Exposure Notification protocol.

A better approach would be for public health authorities to be given a greater say in the design and implementation of these apps, and in a context-sensitive manner. This means shifting away from a one-size-fits-all approach promulgated by technology companies, by putting the choice—and responsibility—back into the hands of health authorities backed

We prioritized the ends (pandemic response operations) over the means (proving the value of the technology in the short term).

by legitimate mandates, and into the hands of users who consent to participate in such a system. This also ensures the systems are ultimately useful—not only to their users, but also to the health authorities we trust to shape public health policy and to take charge of non-pharmaceutical interventions. Ultimately a contact tracing app is but one of many tools that public health authorities can and have used to address the COVID-19 pandemic.

Adopting an open, inclusive, multistakeholder model. With so many issues to grapple with, how can we make progress on areas of consensus, while respecting the diverse social, political, and economic contexts that digital contact tracing technology will be deployed in?

Well, we can take a leaf from the internet's governance model and the concept of a multistakeholder process, with open participation from all relevant stakeholder constituencies: academia, epidemiologists, technologists, public health professionals and authorities, governments, international organizations like the World Health Organization (WHO), app development teams from around the world, and professional and standards bodies such as the ACM, IEEE, and Bluetooth SIG. The goal of such a process should be to steward a process to evolve a consensus around principles that should guide the actions of participants across domains and jurisdictions in respect of the design of these systems and technologies. In turn, this could guide the development of open and transparent standards, as well as implementations.

A set of shared standards and implementation principles will also stand a better chance of allowing true interoperability across geographies. Shortly after TraceTogether launched, my team and I wrote the BlueTrace manifesto, in which we stated, "COVID-19 and other novel viruses do not respect national boundaries. Neither should humanity's response… We believe that TraceTogether and its sister implementations should be interoperable, and that's what we're building towards." This work remains unfinished.

Turning off the systems will be a watershed moment. It will signal the end of this pandemic chapter in our shared experience. But the work should not stop there.

One promising platform for advancing these discussions is the Linux Foundation Public Health's Project Herald, which currently hosts an open-source repository derived from the original TraceTogether code base, which my team contributed. Herald is currently used by several other governments that prefer not to be locked into a proprietary system and want to maintain the option of a complementary hardware device—especially for persons who do not have access to Bluetooth-capable smartphones. Unlike a closed ecosystem, Herald envisions a modular architecture that is not tied to specific technologies, or hardware. Adopters are also free to integrate the Herald system into their own public health systems.

Technology aside, we also need to shape norms around the governance of the technology. Just as there are differing perspectives on the role that social media companies play in shaping the online information landscape and regulating freedom of expression, there is a lack of consensus on the appropriate role that Apple and Google should play as key arbiters of what is possible with the available technology, in the context of public health—an otherwise Westphalian sovereign imperative.

Much of how digital contact tracing technology was governed and defined during the COVID-19 pandemic was an inevitable consequence of a lack of framework and norms for how such systems should be developed and used, and the ethical choices that were *de facto* encoded in the design of the systems by the various technologists that worked on them— including for TraceTogether. From first principles, it is not at all consistent that digital contact tracing data, when used for public health purposes, should not be subject to the same ethical and governance frameworks that are in place for other types of medical data.

CONCLUSION

What will the future of digital contact tracing technology look like? How will it be used in future disease outbreaks? Will our stewardship bear fruit for a future pandemic? Perhaps, perhaps not. Personally, I hope that we never have to resurrect such technology in future. But we should be prepared, regardless. And that work starts now.

Note

The views, thoughts and opinions expressed here belong solely to the author and do not necessarily reflect those of the Singapore Government, GovTech, or any other group or individual. Please contact Jason at jasonbay [plus] tracetogether [at] alumni [dot] stanford [dot] edu if you would like to contribute to discussions in the digital contract tracing space.

References

[1] Bay, J. Automated contact tracing is not a coronavirus panacea. April 10, 2020; https://blog.gds-gov.tech/automated-contact-tracing-is-not-a-coronavirus-panacea-57fb3ce61d98

[2] Baharudin, H. Nearly all in Singapore are on Covid-19 contact tracing programme TraceTogether. *The Straits Times*. January 20, 2022; https://str.sg/wA4U

[3] Teo, Y. Y. and Lee, V. What's behind the recent rise in Covid-19 deaths in S'pore? Here are some facts you may not know. *The Straits Times*. November 6, 2021; https://str.sg/3pye

[4] Weiss, M. B. and Mehta, S. *TraceTogether*. Harvard Business School Case 820-111, June 2020.

[5] Wong, Y. N., Miller, S., Cheah, S. M., and Ang, J. Digital product management under extreme uncertainty: The Singapore TraceTogether story for COVID-19 contact tracing. Singapore Management University Case SMU-21-0045. February 2022.

Biography

An engineer-turned-policymaker, Jason Bay bridged the worlds of technology and policy as senior director of government digital services at the Government Technology Agency (GovTech) in Singapore. In addition to managing more than 400 software developers and 50 digital transformation initiatives, when COVID-19 struck, Bay masterminded the development and launch of TraceTogether, the world's first national contact tracing app, and oversaw its integration into Singapore's pandemic response. He has consulted with numerous national and international digital contact tracing and exposure notification teams and is a regular speaker on the subject. Besides technology, Bay's policy experiences span the domains of defense, international relations, law, telecommunications, and futures thinking. He graduated from Stanford University with bachelor's and master's degrees in electrical engineering, as well as a master's degree in management.