HARVARD Kennedy School
BELFER CENTER
for Science and International Affairs

BRIEF
belfercenter.org/STPP

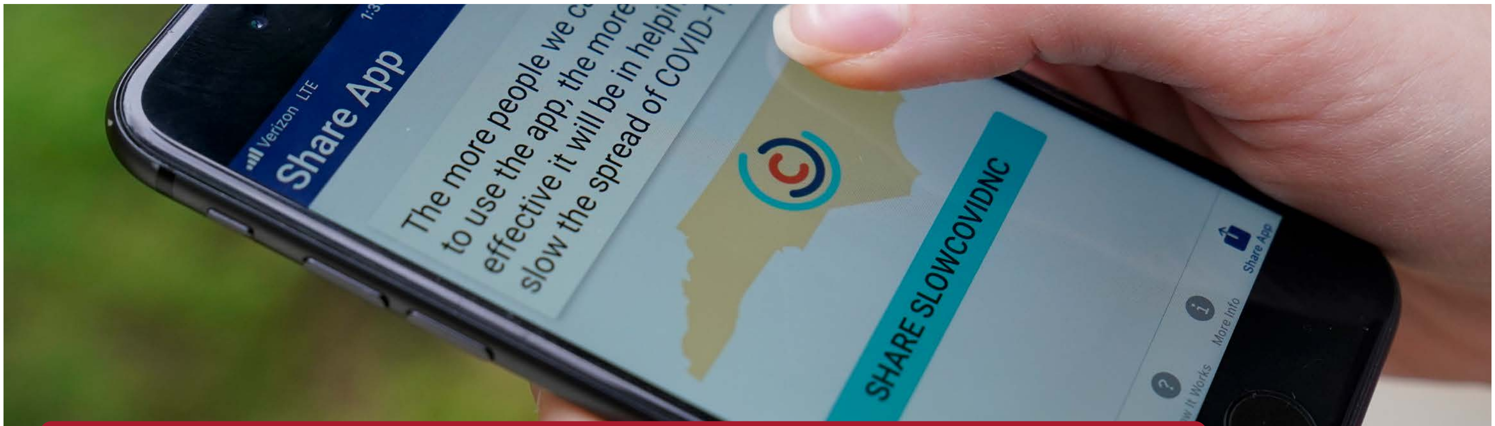

SCIENCE, TECHNOLOGY, AND PUBLIC POLICY PROGRAM | FEBRUARY 2021

Technical Difficulties of Contact Tracing

Amy Robinson and Jim Waldo

In mid-October, thousands of English and Welsh citizens received phantom alerts that they had potentially been exposed to COVID-19. A quick Twitter tour reveals the spiraling fear, frustration, and confusion that ensued. Even though National Health Service (NHS) later updated the app, built using an Exposure Notification System (ENS) developed by Apple and Google, the incident still amplified mass hysteria and confusion.

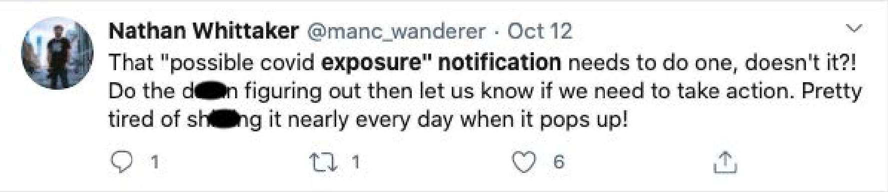

Nathan Whittaker @manc_wanderer · Oct 12
That "possible covid **exposure" notification** needs to do one, doesn't it?! Do the d[redacted]n figuring out then let us know if we need to take action. Pretty tired of sh[redacted]ng it nearly every day when it pops up!

Matt Collins @Mattophobia · 4h
I just got a NHS COVID **Exposure Notification**.

I couldn't read any more details, the app actually didn't even mention the notification, and I don't know where or when the contact is from. The notification also said I didn't need to do anything.

What is the f[redacted]ng point.

The NHS bug demonstrates the real problem of false positives in digital contact tracing. A false positive occurs if the app alerts someone of possible exposure to coronavirus when no such exposure has occurred. A high rate of false positives has two potential problems. First, it could overburden a state's limited testing capacity, as concerned citizens flood the already overwhelmed testing sites. On the other hand, people could become numb to notifications if the app continues to ping them with possible exposure. Then, people who really have been exposed will ignore the warning and not get tested.

While not as panic-inducing, false negatives can be just as deadly. A false negative occurs when a person who was actually exposed to the coronavirus does not receive a notification. If asymptomatic and unaware of a possible infection, she will continue her daily business and further spread the virus. Medical experts have dubbed such oblivious asymptomatic transmission "the Achilles' heel" of the pandemic, especially as social distancing restrictions are relaxed.

Therefore, a digital contact tool must sufficiently minimize false positives and false negatives to ensure it does more good than harm. This is especially true as the number of U.S. states deploying digital contact tracing apps grows. In July, Google announced that 20 states and territories were "exploring" apps based on the Apple | Google ENS, which would represent approximately 45 percent of the U.S. population. New York and New Jersey's recent app rollouts bring the total of state public health authorities currently using the Apple | Google ENS to eleven. In order to understand if the Apple | Google ENS is up for the challenge, we must understand the accuracy of the underlying Bluetooth technology. Long story short, Bluetooth technology simply cannot provide location information that is granular or consistent enough for digital contact tracing apps to reliably function.

GPS Technology is Insufficient

In order to understand why most digital contact tracing apps rely on Bluetooth low energy (BLE), it is important to understand the shortcomings of GPS technology. Early in the pandemic, location tracking based on smartphones' global position system (gps) was proposed as a potential aid to contact tracing. For instance, projects showing the possible spread of the virus by tracking the movements of college students during spring break gave the impression that individual locations could be tracked using the same gps technology.[1]

However, while the gps tracking on phones may have been accurate enough to track students dispersing after spring break, it is not accurate enough to be used for the kinds of measurements needed for precise contact tracing. Contact tracing requires a resolution that can distinguish between co-location within six feet (approximately two meters) for a particular minimal duration of time (currently 15 minutes). But gps has known problems of not being available in some buildings (where contacts are the most likely to cause viral spread). Even when gps is working at its best, the resolution of the gps in a smartphone is only 7-13 meters[2] in an urban environment. Many things can confuse the workings of gps, including time of day, strength of wifi signals, and the time of the year. This accuracy is between 2 and 3 times lower than what is required to do reasonable contact tracing.

Due to gps's shortcomings, most mobile contract tracing has attempted to use the bluetooth networking built in to most phones. Bluetooth is a near-field networking technology, and thus would seem more appropriate for the finer-grained location information that is needed for this kind of application. However, even bluetooth might not provide the granularity required by digital contact tracing to estimate distance.

---

[1] See, for example, https://www.masslive.com/coronavirus/2020/03/coronavirus-tech-company-map-shows-covid-19-impact-of-florida-spring-break-on-rest-of-us.html

[2] See Krista Merry and Pete Bettinger, Smartphone GPS accuracy study in an urban environment, PLOS One, July 18, 2019, accessed at https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0219890.

Bluetooth's Insufficiencies as a Distance Proxy

First and foremost, it is important to underscore that bluetooth does not measure distance - at least not directly. Instead, the Apple | Google ENS uses the simple idea that a signal becomes weaker the further away it is. Therefore, one can use the attenuation, or reduction in signal, to infer distance. A weak received signal strength indicator (RSSI), as approximated by the power measurement of dBm, would mean that you are further away; a strong RSSI indicates you are closer.

Fluctuations and Calibrations

Yet this method is not so simple in practice. RSSI does not provide a clear and consistent measure of distance for a variety of reasons. First, RSSI naturally fluctuates by as much as 5dBm even in very controlled settings. Second, RSSI fluctuates by device. RSSI is itself a bit of a strange measurement because it depends on both the strength of the sending device as well as the sensitivity of the receiving device. Therefore, a weak RSSI could be caused by the phone emitting a weak signal, or the other phone not having a very sensitive receiver. In tuning their recently released TraceTogether app, Singapore researchers found that different devices varied by as much as 20 dBm in highly controlled settings due to different bluetooth hardware, antenna layout, and even the operating system (OS) configurations such as battery saving features. A difference of 20 dBm roughly translates to a factor of 10 in terms of distance measured. In practice, that could be the difference between measuring 1 meter and 10 meters.

The good news is that the Apple | Google ENS can tune the algorithm based on these known variations of supported devices. Apple and Google explained that they can even extend calibrations to currently unsupported devices based on averaging associated devices.

The most recent calibration file lists almost 12,000 devices. However, Apple and Google provide a confidence level of either low, medium, or high for each calibration. Of the 11,809 listed devices, only 306 devices had high confidence levels from measurements taken directly from that specific model. The calibration numbers for the other 97 percent of devices were listed as low confidence. Apple and Google themselves note that this is "a very coarse method of calibration" to just serve as a "stopgap" until data from more devices is available.

Real World Scenarios

Third, and most importantly, RSSI fluctuates significantly depending on real world situations. After all, Bluetooth LE signals, which are a type of radio wave, can absorb into or reflect off of various surroundings. For instance, a person's body can absorb Bluetooth radio signals, making the signal look much weaker and therefore that the devices are much further away than they actually are. Therefore, a person simply rotating their body can alter the signal strength by as much as 20 dB. Changing the position of a phone in a purse just one meter away from another phone can also alter the signal strength by 10-20 dB. Even the presence of carpet and furniture could make the phones seem further apart than they actually are. Signal absorption can therefore increase the level of false negatives since the ENS would not register that a person came in contact with an infected person.

On the other hand, real world fluctuations can also increase the chance of false positives. Especially in indoor settings, signals can reflect off of metal and other reflective surfaces, making them appear much stronger and therefore much closer. Signals could even reflect off of wet pavement. In replicating a metro car, researchers found that the RSSI actually *increased* as the phone moved from 2 meters to 4 meters away. Therefore, the ENS would register that the people were getting closer when in reality they were moving further apart.

Bluetooth signals can also propagate easily through certain walls. While signals have a difficult time passing through blockwork or cement walls, they can easily pass through stud walls. Stud walls are generally used to separate rooms whereas blockwork is generally used to separate adjacent houses and apartments. Therefore, the risk of a false positive may be low for neighbors, but it could be substantial for people in adjacent rooms, such as in an office space. RSSI simply cannot capture real-world conditions, such as a dividing wall, that would drastically change the risk of infection.

After taking into account all of these fluctuations and real-world scenarios, it is apparent that RSSI is a very tricky and inconsistent measurement. In trying to determine if two people came within 1.5 meters of one another, one study ultimately concluded that RSSI yields an error rate of around 50%, even when the two people were next to each other for ten minutes. Apple and Google even acknowledge in their documentation that "Attenuation is a very noisy proxy of distance." So if the ENS can't tell you if you were within 1 meter or 10 meters of an infected person, then what can it even really tell you?

Threshold Settings

On top of the difficult task of approximating distance from Bluetooth, the Public Health Authorities (PHAs) are also faced with the actual configuration of the app. The Apple | Google ENS takes into account several different factors, such as distance and time, in calculating someone's Exposure Risk Value (ERV) as measured by Meaningful Exposure Minutes (MEMs). The ENS then determines what level of MEMs triggers a notification.

While earlier ENS versions provided the framework for determining what exposure triggers a notification, iOS 13.7 enables state PHAs to fiddle with the knobs themselves. The PHA can customize the weight attributed to levels of infectiousness (exposure to a sick individual the day after he experienced symptoms is more significant than exposure to him 14 days later). The PHA can also customize the weight associated with different types of reported diagnosis (a confirmed positive test from a doctor is more significant than a self-diagnosis).

But most importantly, the PHA can define distances (immediate, near, medium, and other) according to various Bluetooth attenuations. For instance, Apple's example defines "immediate" as less than 40 dB, "near" as between 40 dB and 53 dB, "medium" as between 53 dB and 60 dB, and "other" (i.e. far away) as greater than 60 dB.

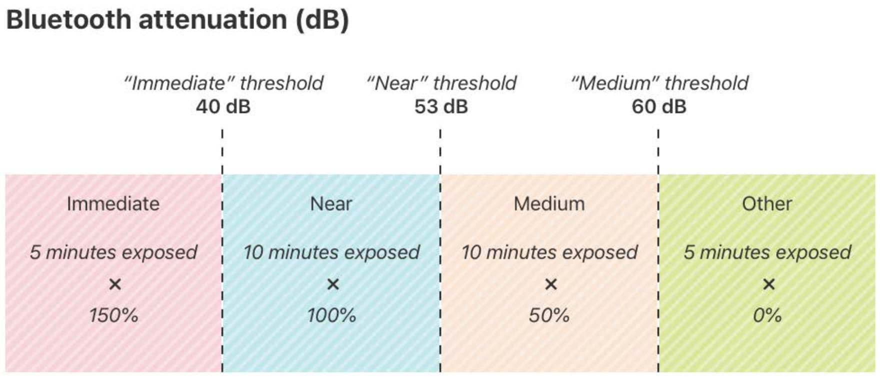

While state PHAs can customize the threshold Bluetooth attenuations, Apple's example reveals the impossible level of granularity necessary for calculating someone's exposure risk. Apple uses intervals of just 7 dB to set critical distance categories. Only 20 dB separates a serious "immediate" exposure from the "other" exposure that is supposedly far enough away so as not to be factored into the calculation. This granularity is next to impossible given that just the positioning of a phone could result in a change of 20 dB.

Enabling state PHAs to tune and customize the ENS parameters gives them some power in determining the level of false positives and false negatives. After all, a PHA deciding that they only cared about interactions lasting longer than 60 minutes with an infected person would generate far fewer notifications than if they set it at 5 minutes. Some epidemiologists believe that this type of conservative tuning can still be useful even in the face of imperfect measurements. In regards to the Swiss app (set to notify someone if she's been within 2 meters of an infected person for at least 15 minutes), epidemiologist Marcel Salanthé commented "if somebody gets an exposure notification, we will feel damn sure it's actually been a contact."

Yet the fluctuations of Bluetooth's precision as a distance proxy in real world scenarios simply does not support such a confident statement.

Opt-In

Accurate location data is far from a contact tracing app's only hurdle. Since the Apple | Google ENS is an opt-in app, people have to actually take the step of enabling the app on their phone. Digital contact tracing is only effective if a significant percentage of people actually opt-in to use it. The threshold level of adoption in order for digital contact tracing to be effective is a bit up in the air. The World Health Organization (WHO) noted that 60 to 75 percent of a country's population needs to opt-in in order for digital contact tracing to be effective. However, the authors of the original Oxford study cited by WHO later clarified that their research was misinterpreted and that even an opt-in rate as low as 15 percent could help reduce new COVID-19 cases if paired with traditional contact tracing.

Yet even 15 percent may be a difficult number to achieve in the U.S. First, less than 75 percent of Americans own a smartphone and are therefore automatically excluded. Second, the contact tracing app will have to overcome strong public distrust. Even though Apple and Google have repeatedly stated that they will not collect personal information, people are already suspicious of data collection from Apple, Google, or the government. People have even been unwilling to work with contact tracing callers, as shown by New York contact tracers struggling to complete an interview.

We already have preliminary data on adoption rates in some of the states that have rolled out the Apple | Google ENS. A month after its launch date, the New York Covid Alert app passed 600,000 downloads. While an impressive number, that figure only represents 3 percent of New York's 19.45 million population. Pennsylvania reported that 330,000 people had downloaded the app. In addition, 107 people had uploaded positive test results, which triggered exposure notification for 22 people. Three of these 22 people then contacted the Department of Health for advice. Accordingly, there can be a significant drop-off of people receiving a notification and acting upon it.

It's easy to say that any effort to improve contact tracing would provide a social benefit. After all, three potentially asymptomatic people reaching out to their public health authority is better than zero, right? However, that is not necessarily true. What happens when thousands of people receive a ghost notification on their phones and flood clinics already straining for resources, as we saw in the UK? We simply cannot afford to sow more public confusion and government distrust. Unfortunately, a digital notification system that presents a high number of false positives or false negatives threatens to do just that. As Bluetooth technology currently stands, digital contact tracing tools risk undermining rather than enhancing current contact tracing efforts. Without more precise location data, it is best that we stick to the traditional tools of contact tracing for now.[3]

---

[3] A shorter version of this work appeared December 17, 2020 at https://www.lawfareblog.com/technical-difficulties-contact-tracing

Science, Technology, and Public Poilicy
Belfer Center for Science and International Affairs
Harvard Kennedy School
79 JFK Street
Cambridge, MA 02138

belfercenter.org/STPP