Salvation by Gibberish

Henry Price and Hugh Lawson-Tancred
Birkbeck Institute for Data Analytics

In our previous paper, we considered 8 questions proposed by the Computational Privacy Group (CPG) at ICL to assess the privacy-acceptability of possible automated digital contact tracing apps [1].

At this stage of the pandemic, a large number of such apps have already appeared, are currently being launched or are in the design phase. It is widely hope that they will be able to ensure a smooth exit from the current lockdown and also contain any future reemergence of an upward trend in infection.

There is an extremely useful review of the possible objections to such apps from both the privacy and the practicality perspectives from the Ada Lovelace Institute [2], and a benchmark model of how such apps might work in practice from the Big Data Institute at Oxford University [3].

Linklaters [4] found on 16 April that 28 countries had already launched some form of contact tracing app, while a further 11 were engaged in the development process. The leading continent amongst those countries that have already launched was Asia with 13 apps.

In addition to national governmental efforts (and the EU-wide initiative PPPE-PT [5]), there have also been independent academic and corporate studies, of which by far the most important is the unprecedented Google-Apple collaboration [6].

In all cases, however, there are also significant concerns about the effectiveness and ethical acceptability of such apps. These objections can be categorised as concerning a) infrastructure/rollout issues, b) uptake issues and c) privacy concerns. A good example of the anxiety provoked in particular by the last of the three is the letter sent by a large number of French academics voicing concern about the French app StopCovid [7]. A similar, but international collective letter has been sent by academics from 26 countries emphasising to what extent public trust is at stake in the event of opting for the wrong design of any widely introduced contact tracing app [8].

Proposals for automated contact tracing are also subject to criticism from advocates of conventional manual contact tracing such as Prof Allyson Pollock of the Population Health Sciences Institute Newcastle University [9]. It is also noteworthy in this context that arguably the most high-tech city in the world, San Francisco, is relying on conventional manual contact tracing [10].

There are, however, a number of material objections to the manual approach. In the first place, arranging and conducting interviews will inevitably take time, but secondly and more importantly human beings are notoriously unreliable as sources of information about their

own, even recent, past. Above all, users of public transport cannot be expected reliably to recall any details about those with whom they shared possibly constricted space.

By contrast, the types of automated contact tracing currently being proposed are much more likely not simply to forget contacts. The problem that they face, rather, is that they may not detect contacts in the first place because of technical obstacles such as wind or other atmospheric disturbances, intervening obstacles, weak signals, battery failures and so on. Decisively, however, automated systems can maintain real-time contract establishment 24/7. They also do not rely at all on the state of alertness of the app user. More generally, it seems reasonable to ask why, given the huge impact made in the commercial and political sectors by automated information gathering technologies, we should not seek to adopt the same advantages in the health sphere.

A further decisive consideration in favour of automated as against manual contact tracing is that contact should not be confined exclusively to human proximity. It is also possible to come into contact with an immaterial object. Obviously, such objects fall entirely outside the scope of manual contact tracing, but it seems extremely reasonable to expect, using IOT technology, that it will be possible for RFID sensors or similar devices to record contacts with both humans and other material objects, thereby vastly increasing the scope of contact tracing to cover possible forms of contract which are extremely relevant for a condition, such as Covid-19, which can be transmitted through droplets on surfaces as well as transfer directly from person to person.

Concentrating, therefore, on a few representative examples of the plethora of automated contact tracing apps currently or imminently available, we explore in this paper the extent to which they comply with the ethical requirements emerging from our previous paper in response to the ICL group.

On a broad overview, we find that all existing apps can be divided in two ways. The first division is between apps that collect entire trajectories/social graphs of the users and apps that gather only the bare anonymous contact data. The second is between those apps that store the data collected on a central server curated by authority and those were the distribution of the relevant data is handled with the greatest possible decentralisation.

This naturally leads to two privacy-acceptability criteria:

1) Does the app collect only bare contact data (minimalist criterion)?
2) Is the data collected stored on a decentralised basis (decentralised criterion)?

In our assessment, if a model complies with both criteria, then it meets all the privacy concerns which have been voiced about automated contact tracing while remaining capable of sustaining the crucial benefits which have been held out for this technology.

We look first at two apps which seem to embrace both the full trajectory approach (and therefore fall foul of the first, minimalist criterion) and the centralised storage of data (thereby failing to meet decentralised criterion). These are the NHSX proposal in the UK and the PPPE-PT proposal in the EU.

The NHSX app has yet to be fully rolled out but according to its website [11]it will certainly involve the centralised approach. The user, on becoming ill, optionally notifies the NHS,

which is then responsible for notifying all those with whom the user has come "into significant contact over the previous few days". This clearly implies disclosure of the infected user's social graph to the central authority (in breach of both criteria).

The European approach, PPPE-PT, appears to be undecided as between the centralised and decentralised approach. According to its website [5], "PPPE-PT currently considers two privacy preserving approaches: "centralised" and "decentralised"". The PPPE-PT proposal is based on creating an "encrypted proximity history" which is not accessible even to the user of the phone on which the app is installed, comprises no geolocation, personal information or other data and will automatically expire when it becomes "epidemiologically unimportant". The PPPE-PT system therefore appears to comply with the minimalist criterion.

However, it is not clearly committed to decentralisation of the storage and usage of the proximity history. If the user of the app reports sick, she is sent a TAN code with which she then notifies the central authority. This (pan-European) authority is then in possession of such a user's full proximity history. The purpose of the TAN code is to prevent the malicious injection of "incorrect infection information" into the PPPE-PT system. In other words, the function of the code is to protect the data integrity of the central authority, not the data privacy of the infected user. The existence of the code alone is not sufficient to ensure satisfaction of the decentralised criterion.

Both these embryonic state-sponsored systems therefore still rely on the centralised storage model. This is also the case with the most fully developed contact tracing app, namely the TraceTogether app based on the BlueTrace protocol and used in Singapore. According to its portal [12], this app assigns a random anonymised user ID to all subscribers to the app. This user ID is securely stored and not disclosed "to the public". However, it remains available for identification by the central authority at all times, irrespective of whether or not the user is infected. This system, like other systems deployed in Asia, appears to breach both criteria: the information provided can lead to disclosure of wider personal information about the subject and that information is also accumulated in storage.

The most forceful and coherent objection to such centralised storage approaches has come from the pan-European academic consortium operating under the banner of Decentralised Privacy-Preserving Proximity Tracing (DP3T) [13]. The DP3T group would regard the British, European and Singapore models as falling under what they call either the data grab or the merely "anonymised" data approach. In design terms, they argue that the centralised concentration of data on the above three models is not adequately protected against the possibility of extension or use/repurposing and therefore could effectively morph into the worst form of model, a "data grab" model.

The central thrust of the initiative is to advocate a fully distributed structure, on which no sensitive data is maintained on a central server. A central repository cannot be dispensed with, but it need only contain nonsensitive data. On the DP3T approach, nonsensitive data takes the form of "ephemeral indicators". These indicators contain no information whatsoever about their source or about any location or temporal context. Such indicators can be safely stored on a central silo, as they contain absolutely no information which could be abused either by the authority itself or by an adversary who has successfully hacked it. Rather, the mobile phones of other app users will automatically access the central repository of ephemeral indicators and determine for themselves whether or not they have been in contact with another phone which has now become infected (or possibly merely at risk).

We entirely subscribe to the DP3T preference for such decentralised storage of ephemeral indicators, and we consider that decentralised ephemeral indicators are able to achieve the objectives of contact warning without any disclosure to the central authority or other users of potentially privacy-compromising information. In design terms, this is therefore the gold standard. It is also clear that the privacy-protecting advantages of this approach have been recognised by Google and Apple, who are proposing to build it into their future operating systems [14].

This approach is also adopted by the MIT Private Kit: Safe Paths solution [15], which explicitly frames the debate in generational terms. Centralised storage approaches represent a first-generation of large-scale biosurveillance technology, whereas fully decentralised systems represent effectively a second-generation. The key concept behind Private Kit: Safe Paths is that of a "pull" approach. Individual users pull from the central repository anonymised data which will enable them to determine whether or not they have had an infected encounter. The regularity of the pulling can either be predetermined by the app itself (without the human user being involved at all) or made optional for the human user to operate at will. Although Private Kit: Safe Paths is entirely clear on the need for decentralised storage and individual retrieval of data on the pull model (thereby satisfying our second criterion), it is strangely uncertain about the need for such data itself to take the form of pure contact data. The system envisages two "iterations". On the first iteration, it is the full trajectory/social graph that is effectively disclosed at the choice of the user, and it is only on the second iteration that the user only makes available pure contact data to the central repository. We should make it clear that we consider that only the second iteration meets our minimalist criterion (although both meet the decentralised criterion).

Our concern is about the extent to which, even given a gold standard system, it is possible to minimise the leakage of information in real-world applications.

We envisage the following practical deployment scenario (which, on our understanding, complies with best practice as advocated by DP3T, Google/Apple and (the second iteration of) Private Kit: Safe Paths):

The app assigns time-limited unattributable tokens to its users. We call these tokens unique gibberish units (UGUs). They are unique because they are never reused and they are gibberish in the sense that it is impossible to infer any information at all from their content. Each mobile phone on which the app is installed keeps a record of its own UGUs, linking them to that phone. This information is not disclosed at any stage to any other party. The apps constantly both emit and receive UGUs to and from phone users they come into contact with. The record (history) of such contacts is also, of course, maintained, but only on the relevant phone. Two users, Bob and Alice, both download the app. Both Bob and Alice both emit and receive UGUs.

Bob now becomes symptomatic and reports himself as sick. His UGUs are now uploaded to a central repository. It is not possible for any curator of the repository to derive any information about Bob from the UGUs that he has uploaded. Alice's phone regularly checks the repository and pulls from it the latest uploaded UGUs. If one of the UGUs pulled is included in the contact history kept exclusively on Alice's phone, Alice becomes aware that she is now at risk and takes appropriate steps (possibly set out for her by the app itself on the basis of the latest instructions from the authorities). At no stage does Alice disclose any information

whatsoever (not even UGUs) to the central repository. It is only if she herself becomes sick that she will submit her own UGUs to the repository.

On this system, Bob will be required to disclose his identity to the health authorities whom he contacts on becoming sick. (He will be incentivised to do so by, for example, promises of furlough if he needs to self-isolate and the availability of appropriate medical support at all stages - the incentivisation issue is not specifically part of the privacy design problem.) Bob therefore loses his complete anonymity and becomes a medical statistic. This does not change the general problem of the confidentiality of medical statistics, which preexists and is independent of the specific privacy issues of contact tracing apps.

The anonymity of Alice is slightly more problematic. Alice is required to take certain measures, including, presumably, self-isolation and, if available, testing. In the latter case, it will be possible for her to receive on her phone a notice from the authority that she has submitted to testing. If her mobile phone is subsequently inspected, then she will be able to display both that she has had an infected contact and that she has been tested. This will in practice amount to an infection passport in the Chinese style. Any authority inspecting her phone will only be able to see that that particular phone has been exposed to potential infection and also been certified by an authority as being owned by a person who has received testing. If Alice tests positive, she is in the same position as Bob. If she tests negative, she goes back to being in the position she was before pulling the UGUs recording the infected encounter.

If, on the other hand, the recommendation is for Alice to self-isolate, it is more problematic how she will be able to demonstrate that she has done this without disclosing her identity. If the authority wishes to enforce self-isolation by monitoring people in the position of Alice, then it would appear that there is no obvious technical solution at present to reconciling such monitoring with the complete anonymity of Alice. In these circumstances, the only two possible solutions are either to rely on the civic responsibility of Alice herself (together with her desire not to infect those close to her) or to accept the need to compromise the anonymity of at-risk users in the same way as with positively infected users. (Obviously Alice can herself also be incentivised to comply with the self-isolation regime, but this does not circumvent the problem of compromising anonymity, albeit now with consent.)

The conclusion of this simple model is that it is possible to achieve effective contact tracing with the level of penetration needed to block or considerably slow the spread of the pandemic and to do so with either no or only very marginal increases in the necessary extent of the disclosure of privacy-sensitive personal information. If users remain uninfected then no data whatsoever is disclosed to any central authority or to any other user which could possibly be used to identify them. Once a user, in our case Bob, becomes infected, then his anonymity is necessarily breached, but no information is needed about his trajectory or his social graph. The position of Alice, who has been the subject of an infected encounter, is that her identity may have to be disclosed to a central authority for monitoring purposes, but that is the maximum extent of any disclosure that will be required of her. In no circumstances will she be required to disclose her trajectory, social graph or any other information beyond her infection (and possibly testing) status.

We consider that, given the present state of technology, this represents the optional reconciliation between the objectives of public safety and privacy protection.

References

[1] "Evaluating COVID-19 contact tracing apps? Here are 8 privacy questions we think you should ask." https://cpg.doc.ic.ac.uk/blog/evaluating-contact-tracing-apps-here-are-8-privacy-questions-we-think-you-should-ask/

[2] "Exit through the App Store?" https://www.adalovelaceinstitute.org/wp-content/uploads/2020/04/Ada-Lovelace-Institute-Rapid-Evidence-Review-Exit-through-the-App-Store-April-2020-2.pdf

[3] "Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing" https://science.sciencemag.org/content/sci/early/2020/03/30/science.abb6936.full.pdf

[4] "28 countries race to launch official Covid-19 tracking apps to reduce the spread of the virus | Deals | About Us | Linklaters." https://www.linklaters.com/en/about-us/news-and-deals/deals/2020/april/28-countries-race-to-launch-official-covid-19-tracking-apps-to-reduce-the-spread-of-the-virus

[5] "Pan- European Privacy-Preserving Proximity Tracing" https://www.pepp-pt.org/

[6] "Exposure Notification Cryptography Specification" https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ExposureNotification-CryptographySpecificationv1.2.pdf

[7] "Mise en garde contre les applications de traçage." https://attention-stopcovid.fr/

[8] "Contact-tracing apps could ‘catastrophically’ hamper trust, academics warn." https://www.digitalhealth.net/2020/04/contact-tracing-apps-could-catastrophically-hamper-trust-academics-warn/

[9] A. M. Pollock, P. Roderick, K. Cheng, and B. Pankhania, "Covid-19: why is the UK government ignoring WHO’s advice?," doi: 10.1136/bmj.m1284.

[10] "‘Contact Tracing’ Efforts Growing on COVID-19 Battle – NBC Bay Area." https://www.nbcbayarea.com/news/contact-tracing-efforts-growing-on-covid-19-battle/2277576/

[11] "Digital contact tracing: protecting the NHS and saving lives - NHSX." https://www.nhsx.nhs.uk/blogs/digital-contact-tracing-protecting-nhs-and-saving-lives/

[12] J. Bay *et al.*, "BlueTrace: A privacy-preserving protocol for community-driven contact tracing across borders."

[13] "GitHub - DP-3T/documents: Decentralized Privacy-Preserving Proximity Tracing -- Documents." https://github.com/DP-3T/documents

[14] "Apple and Google partner on COVID-19 contact tracing technology - Apple (UK)" https://www.apple.com/uk/newsroom/2020/04/apple-and-google-partner-on-covid-19-contact-tracing-technology/

[15] R. Raskar *et al.*, "Apps Gone Rogue: Maintaining Personal Privacy in an Epidemic."