



OPEN ACCESS

Without a trace: Why did corona apps fail?

Lucie White ¹, Philippe van Basshuysen ^{1,2}

ABSTRACT

At the beginning of the COVID-19 pandemic, high hopes were put on digital contact tracing, using mobile phone apps to record and immediately notify contacts when a user reports as infected. Such apps can now be downloaded in many countries, but as second waves of COVID-19 are raging, these apps are playing a less important role than anticipated. We argue that this is because most countries have opted for app configurations that cannot provide a means of rapidly informing users of likely infections while avoiding too many false positive reports. Mathematical modelling suggests that differently configured apps have the potential to do this. These require, however, that some pseudonymised data be stored on a central server, which privacy advocates have cautioned against. We contend that their influential arguments are subject to two fallacies. First, they have tended to one-sidedly focus on the risks that centralised data storage entails for privacy, while paying insufficient attention to the fact that inefficient contact tracing involves ethical risks too. Second, while the envisioned system does entail risks of breaches, such risks are also present in decentralised systems, which have been falsely presented as ‘privacy preserving by design’. When these points are understood, it becomes clear that we must rethink our approach to digital contact tracing in our fight against COVID-19.

INTRODUCTION

After extreme reluctance and delays, at the end of October 2020, many European countries have reinstated some level of lockdown measures in order to contain the severe second wave of COVID-19 sweeping across the continent. Cases in other countries, such as the USA and Russia, are also reaching unsurpassed heights.¹ During the initial first wave, there was a lot of talk about digital contact tracing as a potential means of allowing us to emerge from lockdown safely, preventing the need for further lockdown measures. At this stage, with cases spiralling out of control, the hope for this strategy seems to be all but lost. We contend that, on the contrary, this is the crucial moment to reconsider digital contact tracing as a potential means of avoiding a continuing cycle of repeated, economically devastating lockdowns. This, however, will involve radically rethinking the way we approach contact tracing, and what we are prepared to entertain. More specifically, the seemingly long-resolved debate between so called ‘centralised’ and ‘decentralised’ contact tracing apps (in favour of the more ‘privacy preserving’ decentralised option) needs to be reopened. When it comes to COVID-19, the speed of contact tracing efforts is critical for the success of the measure.² In order to sufficiently speed the process, we contend, we need a centralised

system which does not require a confirmed test before reporting as positive. Although this measure has been all but ruled out due to privacy concerns, we argue that we need to look at the overall ethical risks of each option—the decentralised option, due to its inherent delays, has a low chance of success and still includes risks. A centralised option also involves risks, but shows promise in speeding up the process and turning digital contact tracing into a more effective measure.

EFFECTIVE DIGITAL CONTACT TRACING

The consensus in mathematical modelling of automated contact tracing to control the COVID-19 pandemic is that two factors are absolutely crucial; ‘population uptake...and timeliness of intervention’.² In order to produce an effective means of contact tracing, we need to revisit our approach, with these two factors in mind. Concerning uptake, we need to consider whether the use of contact tracing apps, which have not achieved sufficient uptake in many countries thus far to put a sufficient dent in the spread of the virus, should be incentivised,³ or activated by default, while allowing users to opt out,⁴ or perhaps even made mandatory. We wish here, however, to focus on the latter factor: speed.

It is clear that speed is of the essence when it comes to identifying and quarantining suspected cases of COVID-19. It appears that individuals become infectious shortly after they themselves are infected, and that a substantial degree of virus transmission occurs before the onset of symptoms.^{5,6} Any delay in tracing and quarantining infected persons thus means that they are unlikely to be identified before they are well into the window of infectiousness, in many cases without any indication that something is amiss. It is partially for this reason that so much attention has been focused on digital (as opposed to manual) contact tracing methods, which aim to produce a smartphone app that can collect information about the user’s contacts, and, in the event of (likely) infection, alert these contacts instantaneously, replacing ‘a week’s worth of manual contact tracing’.⁷ Various mathematical modelling studies indicate that ‘a contact tracing strategy will only contribute to containment of COVID-19 if it can be organised such that delays in the process from symptom onset to isolation of the index case and their contacts are very short’,⁸ and that ‘delaying contact tracing by even half a day from the onset of symptoms can make the difference between epidemic control and resurgence’.^{9,1}

¹It should also be noted that increasing the speed and thus efficiency of the app might compensate for lower uptake to a certain degree, allowing for outbreak control without requiring an unrealistically high (as long as it is not

For numbered affiliations see end of article.

Correspondence to

Dr Lucie White, Institut für Philosophie, Leibniz Universität Hannover, Hannover 30167, Germany; lucie.white@philos.uni-hannover.de

Received 9 November 2020

Revised 30 November 2020

Accepted 3 December 2020



© Author(s) (or their employer(s)) 2021. Re-use permitted under CC BY-NC. No commercial re-use. See rights and permissions. Published by BMJ.

To cite: White L, van Basshuysen P. *J Med Ethics* Epub ahead of print: [please include Day Month Year]. doi:10.1136/medethics-2020-107061

A key objective in developing an effective digital contact tracing system should, therefore, be increasing the speed of the intervention as much as possible. We suggest that in order to do this, we need two things—an app that stores some information on a centralised server, and that allows for reporting before a confirmed test result.¹¹ In order to see why this is, let's turn to the ways in which data can be stored in contact tracing apps.

Centralised versus decentralised

The most popular digital contact tracing app variants use Bluetooth signals to gauge when two people with the app come into close contact, and for how long. Each person is assigned a frequently-changing series of ID-numbers ('ephemeral identifiers'). When two people come into close proximity, their phones exchange their ephemeral identifiers via Bluetooth. If someone reports on the app that they are positive for COVID-19, anyone who has this person's ephemeral identifiers stored on their phone (and meets the conditions of 'close contact') during the period of infection can be instantaneously alerted.

These apps are often divided into two camps, depending on where information is stored. In a so-called 'decentralised' app, ephemeral identifiers are generated on the user's own smartphone and exchanged directly between users when they come into contact. When an index case registers as positive, his ephemeral identifiers for the period of infection are uploaded to a central server, and broadcast to all other app users. Anyone who has one of the index case's ephemeral identifiers stored will be alerted. In a 'centralised' app, each user is assigned a permanent identifier, which is stored on a central server. The server then creates ephemeral identifiers for each user and sends them to the user's phone. Phones exchange ephemeral identifiers, and when a user registers as positive, the ephemeral identifiers of his contacts¹² are sent to the central server. These are then matched with their permanent identifier, and the corresponding contact is alerted.¹⁰

The centralised app's storage of a permanent identifier for each user provides us with a way to speed up the contact tracing process significantly. The reason for this can be drawn out by looking at a problem for digital contact tracing: How do we make sure that positive reports from users are accurate? There are two options here. The first (which many countries with both decentralised and centralised apps opt for) is to require a positive test before it is possible to report as positive in the app.^{11 12} This is clearly a good way to ensure that reports are accurate, but it leads to delays in reporting. Although testing turnaround times are rapidly reducing in some countries, the wait between experiencing symptoms, accessing a test, and receiving the results could undermine the success of the measure.

The other option is to allow users to self-report that they are positive for COVID-19 as soon as they experience symptoms. This would speed up the process enough to make the success of the measure much more likely, but we are then hit with the risk that the system will become flooded with false positive reports,¹³ either because users are genuinely mistaken, or because of malicious reporting. This would result in many contacts being falsely quarantined, possibly sending us back into conditions approaching a general lockdown. We could require that positive self-reports are followed up with a test within a certain window

made mandatory) proportion of the population to use it.⁴

ⁱⁱThis type of system was originally proposed by Ferretti *et al.*⁷ and Hinch *et al.*⁹

ⁱⁱⁱRather than the user's own identifiers, as in the decentralised app.

of time, allowing for the rapid release of unnecessarily quarantined contacts, but this will only work if every user can quickly and easily access a test, and if users are sufficiently diligent to voluntarily seek a test and submit a follow-up report quickly. It seems, in sum, very difficult to sort the false from the true positives in a system that allows for self-reporting.

The need for centralised data storage

It is here that the permanent identifiers stored by the central server become crucial. These give us a way to distinguish the (likely) false from the (likely) true positives when no follow-up test is forthcoming. In a centralised system, where the permanent identifiers of users can be associated with each other, the server can keep track of whether a positive report is followed by further positive reports from the people with whom the index case has been in contact. Where an initial positive report is not followed by further positive reports, the index case could be identified by the server as a likely false positive, and all his contacts could be rapidly released from quarantine. Such a measure is not possible in a decentralised system, because there is no way to keep track of who has been in contact with whom over time—the server only holds the ephemeral identifiers of infected persons during their period of infection, and each smartphone only holds the ephemeral identifiers of direct contacts. There is no way to get an overview of clusters of infections, and thus see where no cluster results from an initial report, indicating a likely false positive.

There are three things that should be noted about this proposed system. The first is that the process of identifying clusters (or the lack of one) can take place without directly identifying any individual user—this can all proceed on the basis of assigning each user in a centralised app a permanent pseudonymous identifier. This, however, leads to concerns that users will be easier to unmask in a centralised system than in a decentralised one, as we shall see below. The second is that nothing like this kind of system has actually been implemented in real life thus far. The centralised systems initially trialled in the UK and Australia were abandoned after they could not identify contacts with sufficient accuracy^{14 15}—not due to an inherent problem with the system, but as a result of the difficulty of designing a functional app on Android and iPhones without the support of Apple and Google (who only support decentralised app architectures). The centralised system in France has been plagued by low uptake, and requires a positive test before reporting is possible.^{12 16} Thus, the systems developed so far have not taken full advantage of the opportunity offered by digital tools for effective contact tracing. Third, it should be noted that digital contact tracing alone, even as envisioned here, which has the potential to improve on systems currently in operation, will not suffice to curb the pandemic. Rather, it should be embedded in a comprehensive strategy that includes testing and ensuring that individuals do self-quarantine when alerted by the app, which may require removing disincentives related to missing work, among other things.^{iv}

ETHICAL RISKS

So far we have shown that the likelihood of an app being effective can be increased by storing some of the users' data on a central server. However, as advocates of decentralised systems are quick to point out, storing user data on a central server entails risks of breaches. As we noted, no individual user is directly identifiable

^{iv}We thank an anonymous reviewer for pointing this out.

on the central server (they are identified with a pseudonym). But many developers and defenders of a decentralised app are concerned that it would be too easy for governments to use the information provided to identify individuals and expand the system,^{17–19} allowing the government to surveil its citizens, and potentially pass this information on to law enforcement or use it for other purposes.¹⁷ It is also possible that information on a centralised server could be accessed by a hacker.¹⁷

Advocates of decentralised systems not only point out that centralised systems entail risks of breaches; they also contend that these risks are prevented in decentralised systems, which are ‘privacy preserving by design’.¹⁷ This putative advantage of decentralised systems can be compared with a principle from safety engineering, namely that of ‘inherently safe design’. This refers to the practice of eliminating potential hazard, rather than merely containing it.²⁰ For instance, a design using fire-proof materials instead of inflammable ones is inherently safe, whereas using inflammable materials while preventing a major fire through a sprinkler system would constitute a ‘secondary prevention’. Other things being equal, inherently safe design is preferable to secondary prevention because even if secondary prevention is installed, the hazard is still present, and thus the unwanted outcome can still be triggered through some series of events (eg, the sprinkler system might be destroyed through some unanticipated event, and a major fire might occur); whereas this possibility is ruled out if the hazard is entirely removed.²¹

Advocates of decentralised systems can then be understood as arguing that, by minimising the amount of centrally stored data, decentralised apps are inherently safe (‘privacy preserving by design’). Centralised apps, in contrast, are in need of a sprinkler system—regulations preventing the misuse of data and their effective enforcement—but such secondary preventions are prone to error and thus involve ethical risk, namely possible breaches, caused either by government abuse or by a hacker. Because decentralised systems do not involve this risk, they are preferable. As we will show next, this argument cannot be sustained. There are two reasons for this: first, inherent safety is only preferable to secondary prevention if it minimises overall ethical risk, but decentralised systems entail considerable ethical risks as they are less likely to be effective; second, it is not true that decentralised systems are inherently safe. Taken together, these arguments provide reason to think that, if equipped with suitable secondary prevention, centralised systems could minimise overall ethical risk.

Ethical risks are unavoidable

Safety engineering does not require that inherent safety always be realised. Inherent safety may reduce the likelihood that the purpose of a given design be achieved, or it may even entirely preclude the achievement of this purpose, such as when a hazardous element is an essential part of the system, as in the case of nuclear power plants.²⁰ In such cases, secondary prevention is preferable if this allows the achievement of the design purpose while sufficiently minimising the risks. We have argued that a centralised design increases the likelihood of effective contact tracing, and we may thus be in a situation where a centralised design is preferable, if its associated risks of breaches can be sufficiently minimised through secondary preventions.

There is reason to think that they can. In many countries, data pertaining to health are among the most strictly regulated, and in many places, such as some US states, legislation has been introduced specifically preventing information gathered by COVID-19 contact tracing efforts from being shared for non-public health purposes under any circumstances.²² The

Provincial Court of British Columbia in Canada similarly blocked the sharing of information pertaining to an HIV-infected person with law enforcement, accepting that the ‘compelled disclosure of confidential information would undermine the ability (to pursue) effective treatment of HIV and endanger the lives of HIV-positive persons, thereby placing at risk the health interests of the population as a whole’.²³

This is all to say that, just as sprinkler systems can be an effective means of preventing a fire, whether the information gathered by contact tracing can be shared, and with whom, can be successfully regulated. As with any secondary prevention, risks of breaches cannot be eliminated entirely. However, it should be noted that the alternatives are also not risk-free. Decentralised apps, even if they were privacy-preserving (which we argue below is not the case), entail considerable ethical risks as they are less likely to be effective in containing the pandemic, and failing to do so will involve harm (loss of lives, economic damage etc.). Furthermore, alternative means of fighting the pandemic, such as selectively locking down the elderly, involve ethical risks too, as they severely discriminate against the elderly.²⁴ In our fight against the pandemic, there are no risk-free options.

Decentralised systems are not inherently safe

The second fallacy of the above-stated defence of decentralised systems stems from the assumption that these systems are inherently safe (‘privacy preserving by design’). Cryptographers have noted that both centralised and decentralised systems are vulnerable to hacking attacks, but their vulnerabilities differ.^{10 25} A major concern for centralised systems is that a hacker might be able to identify app users through their centrally stored permanent pseudonymous identifiers, as well as the identities of the people they have been in contact with. Serge Vaudenay argues that such an attack would be difficult to achieve, and would probably require a malicious government authority to store additional information as an app user registers in order to make identification possible.¹⁰ Others, such as Richard Baskerville *et al.*, emphasise that highly integrated systems in general have fewer, controllable vulnerability points, and can thus be expected to allow for better auditing. At the same time, increased integration means that they are more severely affected if compromised, as breaches affect all integrated functions.²⁶ In the context of contact tracing, we might thus expect breaches to be more likely in dispersed decentralised systems while being less likely but more severe in more integrated centralised systems.

The kinds of breaches we should be concerned about in decentralised systems are attacks that could expose the identities of infected users. When a user of a decentralised system reports that he is infected with COVID-19, all of his ephemeral identifiers are uploaded to the central server, where they are accessible to everyone. This makes it possible to record the ephemeral identifiers broadcasted by particular users, and then to check them later against the identifiers stored on the server. This would enable the identification of those that have become infected.²⁷ Such attacks, Vaudenay contends, could be conducted by any tech-savvy user, and ‘are undetectable, can be done at a wide scale, and...proposed countermeasures are, at best, able to mitigate attacks in a limited number of scenarios.’ Attacks on centralised systems, on the other hand, can be better identified and mitigated ‘by accounting and auditing’. Vaudenay even suggests that privacy-conscious users would in fact be less likely to report that they are infected in a decentralised system than a centralised one. Vaudenay also notes that information stored in a decentralised manner could be utilised by law enforcement—for example, ‘after a burglary during which a Bluetooth sensor

captured an ephemeral identifier, suspects could have their phones inspected for 2 weeks to find evidence.’ Because one’s own ephemeral identifiers are stored on one’s phone in a decentralised system, access to someone’s phone would yield more information than under a centralised system.

It follows from these arguments that it would be false to see decentralised systems as inherently safe design, or ‘privacy preserving’. Rather, both centralised and decentralised systems entail risks of different kinds of breaches. Our evaluation of these risks should then be sensitive to the question of how much they are worth taking: if a technology has only a small chance of being effective, there is reason to be wary of its associated risks (and cognisant of the risks of opting for an inefficient means of preventing significant harm); while if a technology has a higher chance of preventing significant harm, the level of risk we should be prepared to accept should arguably be higher. As we have argued in the previous section, centralised apps are more likely to be an effective means in our fight against the pandemic. Given the potential risks and benefits of each option, we need to revisit the assumption that decentralised apps are clearly ethically superior.

CONCLUSION

Concerns about privacy have dominated the debate about digital contact tracing. While these concerns are legitimate, this debate has ignored the fact that the failure of a system being effective involves ethical risks too. We have argued that the effectiveness of contact-tracing apps can be enhanced if we embrace a ‘centralised’ app architecture, in which users’ permanent pseudonymous identifiers are stored on a central server. While this involves risks of breaches, these risks can be minimised through ‘secondary prevention measures’. Moreover, risks of breaches are also present in decentralised systems, which have been falsely presented as ‘privacy preserving by design’. Of course, issues concerning implementation of the proposed strategy remain; a critical aspect of this will be to properly address and assuage the privacy concerns that have led Apple and Google, and much of the public at large, to regard centralised architectures with suspicion.^{16 28} But the considerations outlined here do serve to make clear that we must revise some of the most prominent assumptions underlying the debate on digital contact tracing.

Author affiliations

¹Institut für Philosophie, Leibniz Universität Hannover, Hannover, Germany

²Department of Philosophy, Logic and Scientific Method, The London School of Economics and Political Science, London, UK

Acknowledgements Our sincere thanks to an anonymous reviewer for his swift and comprehensive comments and suggestions.

Contributors Both authors contributed to the conception, drafting and revision of the manuscript.

Funding This research was funded by the Volkswagen Foundation within the project ‘Digital Contact Tracing, Privacy, and Discrimination: On the Ethics of Fighting Corona’.

Competing interests None declared.

Patient consent for publication Not required.

Data availability statement There are no data in this work.

Open access This is an open access article distributed in accordance with the Creative Commons Attribution Non Commercial (CC BY-NC 4.0) license, which permits others to distribute, remix, adapt, build upon this work non-commercially, and license their derivative works on different terms, provided the original work is

properly cited, appropriate credit is given, any changes made indicated, and the use is non-commercial. See: <http://creativecommons.org/licenses/by-nc/4.0/>.

ORCID iDs

Lucie White <http://orcid.org/0000-0001-8292-3789>

Philippe van Basshuysen <http://orcid.org/0000-0003-1947-9309>

REFERENCES

- 1 Coronavirus Resource Center, Johns Hopkins University and Medicine (JHU). New cases of COVID-19 in world countries, 2020. Available: <https://coronavirus.jhu.edu/data/new-cases> [Accessed 18 Oct 2020].
- 2 Braithwaite I, Callender T, Bullock M, et al. Automated and partly automated contact tracing: a systematic review to inform the control of COVID-19. *Lancet Digit Health* 2020;2(11):e607–21.
- 3 Loi M. How to fairly incentivise digital contact tracing. *J Med Ethics*.
- 4 Hernández-Orallo E, Calafate CT, Cano J-C, et al. Evaluating the effectiveness of COVID-19 Bluetooth-Based smartphone contact tracing applications. *Appl Sci* 2020;10(20).
- 5 Ganyani T, Kremer C, Chen D, et al. Estimating the generation interval for coronavirus disease (COVID-19) based on symptom onset data, March 2020. *Eurosurveillance* 2020;25(17). doi:10.2807/1560-7917.ES.2020.25.17.2000257
- 6 He X, Lau EHY, Wu P, et al. Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nat Med* 2020;26(5):672–5.
- 7 Ferretti L, Wymant C, Kendall M, et al. Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science* 2020;368(6491):eabb6936–7.
- 8 Kretzschmar ME, Rozhnova G, Bootsma MCJ, et al. Impact of delays on effectiveness of contact tracing strategies for COVID-19: a modelling study. *Lancet Public Health* 2020;5(8):e452–9.
- 9 Hinch R, Probert W, Nurtay A, et al. Effective configurations of a digital contact tracing App: a report to NHSX. Available: <https://045.medsci.ox.ac.uk/files/files/report-effective-app-configurations.pdf> [Accessed 2 Jul 2020].
- 10 Vaudenay S. Centralized or decentralized? The contact tracing dilemma. *IACR Cryptology ePrint* 2020.
- 11 Robert Koch Institute. Infektionsketten digital unterbrechen mit der Corona-Warn-App [German], 2020. Available: https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/WarnApp/Warn_App.html [Accessed 10 Aug 2020].
- 12 Rowe F. Contact tracing apps and values dilemmas: a privacy paradox in a neo-liberal world. *Int J Inf Manage* 2020;55:102178.
- 13 Sweeney Y. Tracking the debate on COVID-19 surveillance tools. *Nat Mach Intell* 2020;2(6):301–4.
- 14 Clarke L. Australia is set to abandon its centralised coronavirus app – will the UK be next? 2020. Available: <https://tech.newstatesman.com/coronavirus/australia-centralised-app-will-uk-be-next> [Accessed 6 Nov 2020].
- 15 Burgess M. Why the NHS Covid-19 contact tracing app failed, 2020. Available: <https://www.wired.co.uk/article/nhs-tracing-app-scrapped-apple-google-uk> [Accessed 6 Nov 2020].
- 16 Rowe F, Ngwenyama O, Richet JL. Contact-tracing apps and alienation in the age of COVID-19. *Eur J Inf Syst* 2020.
- 17 Joint statement on contact tracing, 2020. Available: <https://www.esat.kuleuven.be/cosic/sites/contact-tracing-joint-statement/> [Accessed 6 Nov 2020].
- 18 Lomas N. EU privacy experts push a decentralized approach to COVID-19 contacts tracing, 2020. Available: <https://techcrunch.com/2020/04/06/eu-privacy-experts-push-a-decentralized-approach-to-covid-19-contacts-tracing/> [Accessed 6 Nov 2020].
- 19 Troncoso C, Payer M, Hubaux J, et al. Decentralized privacy-preserving proximity tracing (DP-3T white paper). Available: <https://arxiv.org/abs/2005.12273> [Accessed 7 Aug 2020].
- 20 Möller N, Hansson SO. Principles of engineering safety: risk and uncertainty reduction. *Reliab Eng Syst Saf* 2008;93(6):798–805.
- 21 Hansson SO. Promoting inherent safety. *Process Saf Environ Prot* 2010;88(3):168–72.
- 22 New York State Senate. An act to amend the public health law, in relation to the confidentiality of contact tracing information (Senate bill S8450C), 2020. Available: <https://www.nysenate.gov/legislation/bills/2019/s8450/amendment/c> [Accessed 14 Aug 2020].
- 23 Provincial Court of British Columbia. Det S. Cullingworth, VPD v. BC Centre for Excellence in HIV/AIDS. March 26 2014, Vancouver. (Production order – Confidentiality of medical records). Available: <http://www.aidslaw.ca/site/download/14135/> [Accessed 14 August 2020].
- 24 White L, van Basshuysen P. How to overcome lockdown: selective isolation versus contact tracing. *J Med Ethics* 2020;46(11):724–5.
- 25 Ahmed N, Michelin RA, Xue W, et al. A survey of COVID-19 contact tracing Apps. *IEEE Access* 2020;8:134577–601.
- 26 Baskerville R, Rowe F, Wolff FC. Integration of information systems and Cybersecurity countermeasures: an exposure to risk perspective. *ACM SIGMIS Database: the DATA BASE for Advances in Information Systems* 2018;49(1):33–52.
- 27 Tang Q. Privacy-preserving contact tracing: current solutions and open questions, 2020. Available: <https://arxiv.org/abs/2004.06818> [Accessed 6 Nov 2020].
- 28 Fahey RA, Hino A. COVID-19, digital privacy, and the social limits on data-focused public health responses. *Int J Inf Manage* 2020;55:102181.