The NHS COVID-19 App, Separating Fact from Fiction.

An essay by Steve Clark

Version 1.2 dated 11th May 2020

(v1.0 first published 9th May 2020)
For overview of changes between versions see the final page.

Executive Summary

Being a technology company based on the Isle of Wight we've been in the very fortunate position of having early access to the new NHS COVID-19 Track and Trace App. We've also seen a lot of the discussions and concerns of our fellow Islanders as they decide whether to download it and their early user experience and what they are saying on social media.

This essay is for anyone genuinely interested in how the App works and to help make sense of what they are reading about it and the alternatives in the media and social media and is aimed at everyone, not just technical readers.

It "gets under the hood" of the App and the technical infrastructure and cuts through the myths doing the rounds on social media to separate the Facts from Fiction and tries to put to bed those issues for which there is a technical answer, and provide a sensible context to where there isn't one answer.

- We explain the Bluetooth challenges and the difference between the "centralised" and "decentralised" approaches that have been widely talked about and look at what Apple and Google are doing.
- We explain why we have no worries about participating in the trial ourselves (nor would have worries about continuing to use the App after a national rollout).
- We look at the privacy issues both from the perspective of a private individual on the trial and how the privacy issues change when looked at from the perspective of the ICO and Parliamentary Joint Committee on Human Rights with respect to safeguarding population scale data as well as views from other commentators.
- We explore some of the other worries people are having about the App on social media and recommend looking at source documents not social media for information to make decisions.

Our opinion is that a digital App operating within the centralised model undoubtedly brings the better potential for getting on top of the virus especially as part of an overall track and trace programme with notable benefits over the decentralised model including hotspot tracking and identifying "Typhoid Marys" (asymptomatic carriers), though unfortunately carries greater privacy concerns.

The App is being developed in a time of crisis where time is of the essence and whilst this certainly doesn't give carte blanche to ride roughshod over

privacy and human rights legislation it is important to appreciate the accelerated time scales involved in both the App development and privacy compliance. Risk shouldn't be seen as one dimensional and it's important to look at both the likelihood and impact of risks in the context of a single individual, the Isle of Wight Trial, and a national rollout and to remember many risks apply to both the centralised and decentralised approach. We should contrast those risks with the benefits, whilst at the same time implementing the necessary measures to mitigate them over the lifetime of the programme.

We personally have already installed the App and have no concern in using it both during the trial and during a nationwide rollout as we believe the benefits outweigh the risks, however we recognise that some may assess the individual privacy risks differently from us and choose not to use the current App on trial in the Isle of Wight, in which case there is merit in the UK also pursuing a decentralised version built on the Apple/Google platform.

Introduction

Being a technology company based on the Isle of Wight we've been in the very fortunate position of having early access to the new NHS COVID-19 Track and Trace App. We've also seen a lot of the discussions and concerns of our fellow Islanders as they decide whether to download it and their early user experience and what they are saying on social media.

With so much attention on it we thought we'd do our best to "get under the hood" to explain how it works, look at some of the inevitable controversy that is doing the rounds in the media and on social media and see what we need to worry about and what we don't. We'll also take a look at the UK approach which is the "centralised" model that you may have heard about and the other approach which is the "decentralised" model.

This essay is for anyone genuinely interested in how the App works and to help make sense of what they are reading about it and the alternatives in the media and social media. It is aimed at everyone, not just technical readers.

We're attempting for this paper to be as accurate as possible – it's written by an objective "techie" and not trying to push any political agenda, but inevitably some of things doing the rounds that people are worried about do have a political overtone so we will even look at the "what if I am giving all my data to one of Dominic Cummings' mates" that features regularly on social media. Where applicable we explicitly say something is our comment, speculation or opinion.

To make this piece as readable as possible we're going to try to stay away from too much technical language, but still be accurate, though later on we will look at a commonly used risk analysis approach. At various stages we've also put in some FAQs to answer some of the questions we've seen commonly on social media.

This is an evolving piece and with everything moving so fast during the crisis there may be changes so if you do find any errors that can be backed up with an authenticated source please ping us the reference via the contact page on our web site at www.clark.associates. At the end of the blog are links to the references we have used.

At the moment this article looks at the App and technical approach, not the overall Track and Trace process around it.

Given things are moving so quickly some of what is discussed here could rapidly be history the datestamp for core information in this essay is 9th May 2020.

How the NHS COVID-19 App works

The NHS COVID-19 App is designed to help the NHS track and monitor the spread of the Coronavirus and help us get early warning of being exposed to the virus and for those with symptoms to get rapid access to testing.

When you download the App it installs and prompts you to enter the first half of your postcode (not all your postcode). It will also ask for permission to use Bluetooth and permission to send you push notifications. On Android phones it also will ask you to enable location services. This isn't to track you using other means, it is because the Android operating system requires location services to be enabled in order for Bluetooth to also be enabled.

It presently (the Isle of Wight Trial) should work on iPhones with IOS 11 and higher and Android version 8 and higher.

It works by detecting contacts using Bluetooth. This means it can log that a contact happened with someone else but not where that contact takes place. If you don't believe it only uses Bluetooth we have covered that off towards the end of this article.

We probably all have lots of Bluetooth devices in our homes ranging from phones to TV's to audio speakers to toys and gadgets. Bluetooth devices periodically send out a "beacon" to tell other Bluetooth devices they are near and to advertise the kind of "services" they provide, which typically includes their primary service (e.g "I'm a phone") and also subsidiary services (e.g "this is my signal", "this is how much battery life I've got"). Normally there is a pairing process that most of us will be familiar with for example to pair a phone to our car– once paired the phone lets the car use the services to make calls and display the battery life and signal strength for example.

If you are interested you can see this data (though it's usually in a techie form that is difficult to make sense of) by using an App like nRF Connect which is available for iPhone and Android.

Some devices also provide data within the beacon (e.g. a thermometer might broadcast the temperature) or devices can be programmed to be already aware of the services provided by similar devices and this means they can provide some functions without being paired. Obviously the App would be very awkward to use if you had to stop and pair your phone with everyone you encountered so this latter method is how the App works – it broadcasts an encrypted unique identifier which can be received by someone with a phone that is also running the App and listens for the same from others.

FAQ: Isn't Bluetooth unsecure?

Like any wireless signal Bluetooth can be eavesdropped. Bluetooth got its "bad rep" for security mainly when it was discovered that many Bluetooth products used easy to guess pairing keys like 1234, 1111 or 0000 and that the services were provided "in the clear" (i.e. unencrypted). In the scary scenarios published by the media someone nasty could connect to a child's toy with a speaker in it (if they were sitting within Bluetooth range with a bunch of electronics in which case a parent just maybe might notice!).

Now the NHS App doesn't even use pairing so of course if you know what you are doing you could eavesdrop the wireless signal and "steal" the unique ID. However the ID is encrypted so you would be stealing something encrypted and of little use to you. You'd also need to be within the range of the Bluetooth signal to do that and then about the only thing you could do with it is fake that you were the mobile phone that sent it – which to be honest isn't really going to be high up anyone's list of useful things to do. Even so, the App developers have got that covered – the encrypted ID is changed periodically so that even if you did steal it, by the time you had decrypted it (even if you could) you would have little opportunity to use it. Having said that, some commentators are making more of this potential vulnerability which we discuss later on.

FAQ: The NCSC at GCHQ helped develop the security so doesn't that mean the government can snoop on you?

OK, so let's assume they've put a backdoor into the encryption system and can get at the data and decrypt it? Let's park this question for a moment and call them in this sense a "malefactor", someone that wants to do bad things with our data, and we'll cover it off when we look later on at "What if a malefactor can get access to the data". Short answer though is that the stuff they can get from this App they can most likely already get from other means – we've all seen in the dramas and action entertainment programmes like "Hunted" how they can work out which phones have been near each other, and of course if you are on social media you are probably already giving away far more interesting things than just an encrypted unique ID.

So the App sits on your phone and logs any Bluetooth contacts from other phones that are also running the App…

FAQ: hang on a minute before you go on, Lots of people have been saying that Bluetooth doesn't work properly to do this especially on iPhone?

Ahead of the App's introduction many were speculating that the Bluetooth approach wouldn't actually work well on iPhone, with the common theme being that an App can't use Bluetooth in the background. At face value if you look around there is quite a lot to support that, however like many technical things it is only part of the story. Firstly you can wake Apps up periodically – you may see this referred to with reference to terms like "keep-alives". Use of "keep-alives" is very common in the technical world and a typical use is allowing things that consume power to "sleep" but they can be woken up to do a job and then go back to sleep again. Secondly there are some aspects of Bluetooth service request and response that can be done in the background, so for example an iPhone App can "listen" in background mode to the Bluetooth beacon for another device and wake up to deal with it if it "hears" one.

So, many commentators have been quite surprised that following the start of the trial the iPhone App appears to work fine but it also leads to the situation some investigators are reporting that if two iPhones are left on standby for some time they don't wake up to record a contact (with some reports adding "unless they are also in contact with an Android phone"). It looks like there may be a scenario where this is indeed the case, where the iPhone App has gone sufficiently dormant that it isn't allowed to continue to be woken up by "keep-alives" and so if two phones in that state encounter each other neither of the two phones try to wake up the other, however if an Android phone (or no doubt a woken up iPhone) is introduced to the mix then the dormant phones that have been allowed to "listen" but not "speak" hear the message and the contact is recorded. Some technologists have postulated the presence of *any* non-dormant Bluetooth device (not just a phone) is enough to keep the App alive and fully functional.

We've seen several reports on how long the mutual dormancy might take, however given how many people use their phones really frequently and given it requires both to be iPhones and both to have become dormant it's probably not likely to happen that often in a scenario that matters. The developers have reportedly said as much. No doubt the Isle of Wight trial will create more data with which to assess this.

FAQ: What about battery usage?

Some social media users have expressed concern about battery usage as you do need to have Bluetooth enabled and the App open (even if it is in the background). In our experience our iPhones are reporting a usage of approximately 1% battery in a 24 hour period (though we are users that would have normally had Bluetooth enabled in any case, so users that normally have Bluetooth disabled may see a greater hit on battery life – we haven't seen that quantified at this stage).

FAQ: But what about time and distance? It says the App measures distance and time how does it do that?

Well time is easy. The phone has a clock. Distance is a bit "rough and ready". The phone can see the signal strength it is receiving from other devices sending out beacon signals. Walls and clothes can affect this but we are very confident that it's possible to approximate distances that are "close", "social distancing kind of distance" apart and "further away". You can see this yourself on the nRF Connect tool mentioned earlier. It's also been suggested that the App developers have taken account of different device characteristics as the App knows what phone it is running on.

OK, back to the explanation of how it works…

The App sits on your phone and logs any Bluetooth contacts from other phones that are also running the App. Data from contacts that are older than 28 days get discarded.

The App's main screen has a constantly available question asking "how are you feeling today". If you feel ill you follow the instructions to enter your symptoms and if your symptoms make it likely you have got COVID-19, at the request of the App, you can choose to upload the contact data that your phone has recorded.

You now enter the wider track and trace programme and can be given access to a test. If this test comes back positive you can tell the App and your encrypted ID is now updated on the server as positive for the virus.

The AI on the central server is able to look at the circle of contacts that you have been in touch with and for how long and your approximate distance apart. It can then assess the risk and push notifications can be sent to those who have been in the vicinity of someone who has subsequently had symptoms and especially who may also have subsequently tested positive and will offer further advice in the context of the assessed risks and the wider track and trace programme.

The Centralised Model

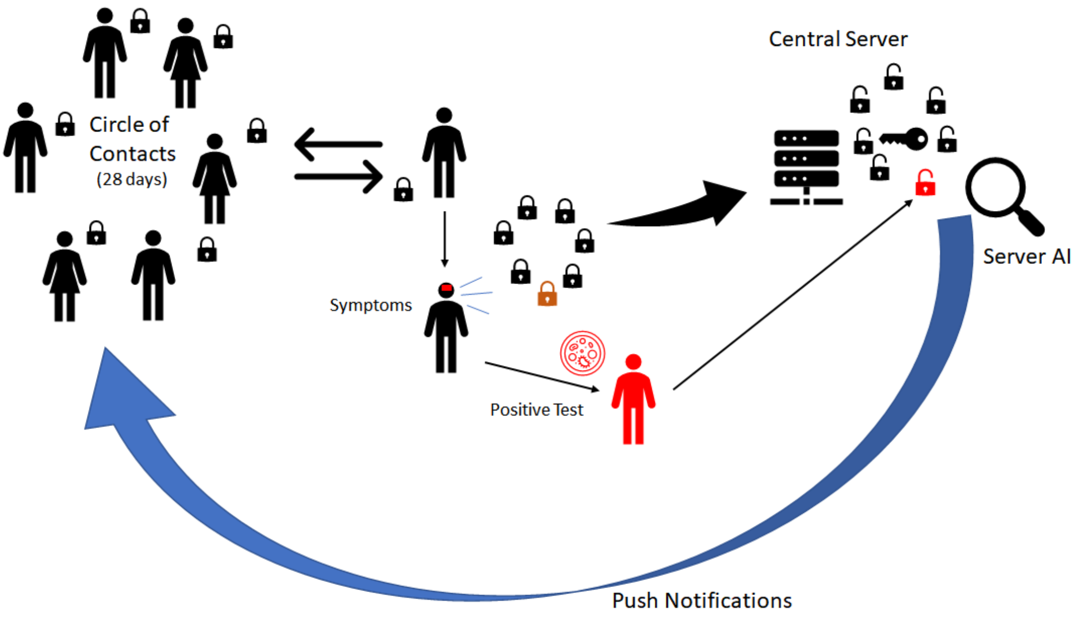

A common early social media comment from the trial is that if we are under lockdown it's of no use as we aren't interacting; people who are out and about are saying things like "it didn't ping and warn me until later at home". That "it told me I walked past someone infected in the supermarket – why were they out?"

OK let's explain all of those. Firstly during the trial, there is much useful information to be had in order to scale the technical infrastructure, so even for people who see little point in having the App, having so may still be useful to the developers (our speculation). It also has the scope to allow valuable statistical data to be derived including different "R" factors depending on different levels of contact, which helps inform the speed and scope of relaxation of lockdown (our speculation).

Secondly the App's primary purpose is not to alert you to someone who is positive when you are in contact with them (at the time of writing we don't know if there is a special message or scenario for that), not least because someone positive with the virus should be self-isolating. So if the App alerts you when you are at home or in the supermarket it is because the system happened to have chosen that moment to push out the information to you about a previous contact who now knows they are positive for the virus. It can't and it won't tell you *who* the infected contact was as the phone and indeed the entire system only knows about encrypted IDs, not actual identities.

FAQ. So far everything seems encrypted and the central server only knows my phone's secret ID that even I don't know and that of people I have been in contact with, so why are lots of people on social media saying 170 scientists and the Information Commissioner's Office and the Parliamentary Joint Committee on Human Rights are giving the impression that it isn't safe and against human rights? And why are people saying we should be doing the Apple/Google "decentralised" model.

OK – big questions. Will answer in several parts. Firstly, the technical bit here and address the functions of those bodies later on and what they are saying and their concerns.

The UK App that is in trial is using the centralised model but it's probably worth discussing how the so called "decentralised" model works first as that helps answer a lot of questions.

A large part of the way the "decentralised" model works is using a "peer to peer" system. Many of us will have come across that term in the context of computer games where you play against other people or even the old music sharing services for those of us that remember that far back. The idea is that there is no central repository for the data and therefore no central body has your data and therefore there is no scope for anyone abusing that data.

Unfortunately a fully peer to peer system can't work in the context of a virus track and trace application, because over time your circle of contacts changes and although a lot of your contact circle as the lockdown eases will be your friends and family, many won't be, so there is no guarantee that you would meet someone that later tests positive a second time in order to know if it was up to the App to detect it by sharing data with another phone. For this reason even the so called "decentralised" model has a "central server". The major difference is that if you have symptoms and are tested positive it is only your positive encrypted ID that ends up in the central server (whereas in the centralised model, the non-positive contacts end up there too). The way your App knows if you have had a positive contact is not from a push notification generated by the server but instead it has to look at the positives in the central server and then compare them locally with the contacts it has recorded to see if you have had contact with them – it is where the "intelligence" of the decision making is done that is centralised or decentralised.

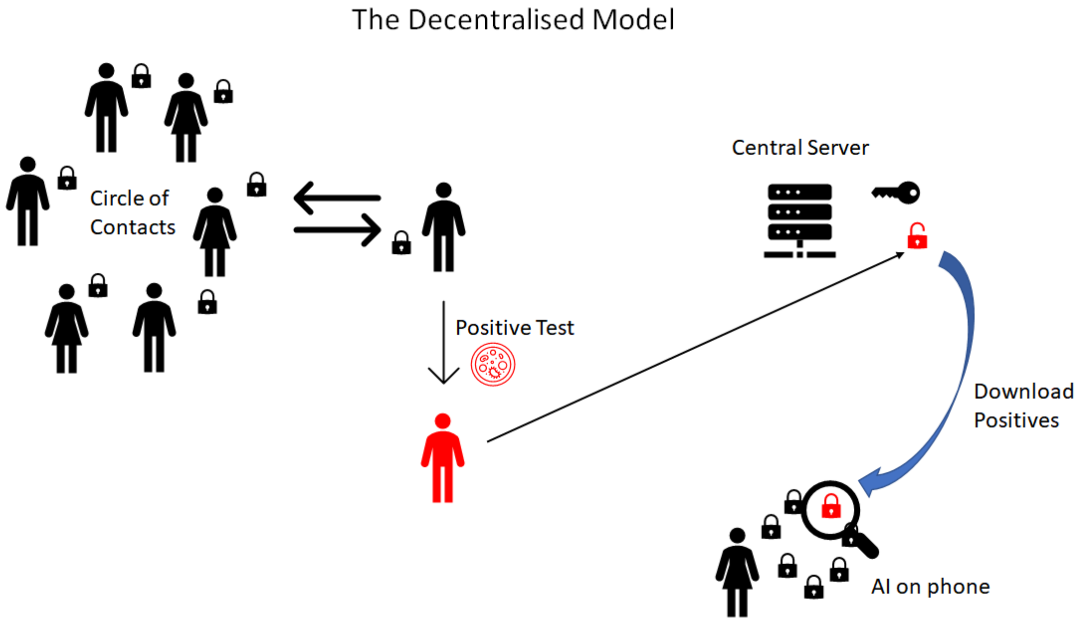

So this is the major reason why people argue that the decentralised model is more secure than the centralised model – in the decentralised model if the centrally held data is compromised, someone gets a bunch of unrelated encrypted IDs and even if they are able to decrypt and relate the ID to a real person by correlating it with other data all they get is a list of people who have tested positive; in the centralised model if the centrally held data is compromised, someone gets contact circles, which means that if by correlation with other data that malefactor is able to relate the ID's to actual people they could, theoretically, if further measures aren't taken, get a list of people's intersecting contact circles. Just to briefly touch on the ICO/Parliamentary Joint Committee for Human Rights concerns, if you apply this logic at the scale of the entire population there is a theoretical risk that a malefactor could work out how everyone in the population relates, in terms of the recorded contacts at any rate, to everyone else, so their concern is to make sure measures are in place to prevent this.

FAQ: "Almost all the other countries" are doing the decentralised model, why isn't the UK?

That’s not actually the case. Many countries are looking at the centralised model because it offers a number of benefits over the decentralised model, however the decentralised approach is also endorsed by Apple and Google. Some countries have been reported to be changing from the centralised approach to the Apple/Google approach because they were unable to solve the Bluetooth problem (source BBC News); the problem apparently “solved” by the developers of the UK App as described earlier.

Apple and Google have proposed a decentralised model where they make an interim API and ultimately operating system changes to better enable contact tracing (presumably in Apple’s case allowing Bluetooth to function without the restrictions discussed earlier). An API is an easy to use Interface to enable Application Programmers to develop software using features of the operating system or existing codebase. At the time of writing an early version of their APIs were released to developers in late April, and the first proper release is anticipated in mid May.

They are intending that in a second phase the operating system itself will carry out the basic functionality of contact detection and alerting of a prior contact becoming positive even in the absence of an App. They are looking to “state actors”, most likely the health or authority of the individual countries, or perhaps individual states in larger countries, to develop Apps on the API and in due course operating system functionality that they will provide.

Comment: Apple and Google have no real choice than to endorse and advocate the decentralised model over the centralised model. Can you imagine the outcry if they tried to build a centralised model thereby potentially giving them access to the entire contact circles of the smartphone owning population of the World! We mention this because some people on social media are using Apple and Google’s choice of the decentralised model to argue that the choice of the model by such technology heavyweights makes the decentralised model the “right choice”.

Looking further at the strengths and weaknesses of each model, firstly the decentralised model is much easier to spoof and already on social media some people are saying, what if I just lie about my symptoms and say I am positive. In the raw decentralised model being built by Apple & Google they implicitly acknowledge this is possible, and they are anticipating that state-built Apps and surrounding track and trace programme would take responsibility for authenticating (maybe with a special test result code that the user has to enter into the App) that positives are indeed people who have tested positive, in a similar way that would be the case in a centralised model.

FAQ: so what does the centralised approach do that the decentralised approach doesn’t?

The centralised approach offers several benefits that the decentralised approach doesn’t. Because it knows the first part of the postcode (in the UK) and because it can see intersecting contact circles of not just positives but also the negatives it can detect virus hotspots that merit further healthcare investigation, and very importantly “Typhoid Marys” i.e. asymptomatic carriers.

These are really valuable tools for the NHS to better understand and fight the virus and such is their importance it's almost ironic that their mention takes barely a couple of paragraphs in this essay.

Detecting "Typhoid Marys" (asymptomatic carriers)
[Greatly simplified to explain the concept]

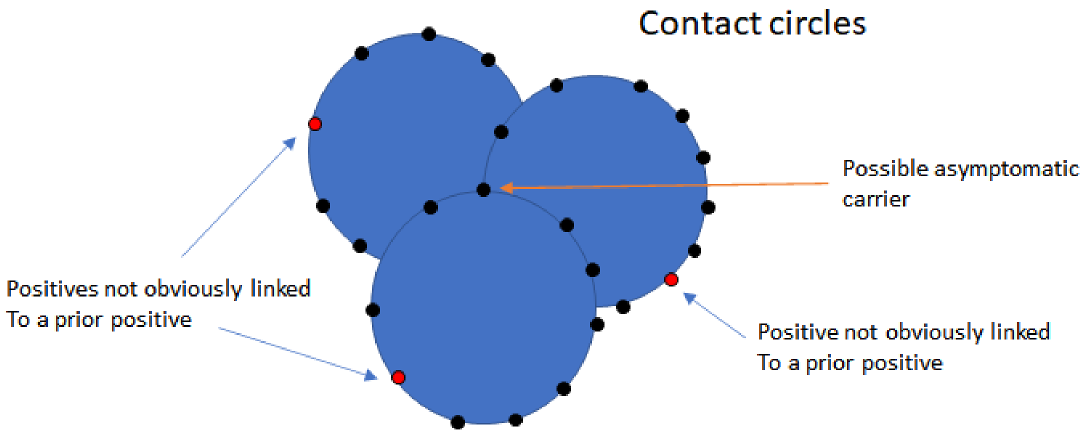
Contact circles
Possible asymptomatic carrier
Positives not obviously linked To a prior positive
Positive not obviously linked To a prior positive

Detection of asymptomatic carriers is very important in combatting the virus especially as a lockdown is eased because they don't get ill or have symptoms themselves and so can inadvertently spread the virus to a large number of other people creating a big "R" number (measure of how many people get infected by one person). The way for example that you would suspect and work out which unique ID an asymptomatic carrier is if you see positives starting to appear in interlinked contact circles with no apparent link to a previous positive. By looking to see if there is a common unique ID that has links with those circles you may have found the asymptomatic carrier. An asymptomatic carrier wouldn't even be in the database in the decentralised model.

Also, using the centralised system, knowing about symptomatic contacts (not just positive ones) the central server can use AI to decide whether to push warnings to contacts even before the person with the symptoms has had the results of the test.

The other benefit right now is that the NHS App is already in trial on the Isle of Wight. The Apple/Google solution isn't ready yet and at time when "every day counts" this is material. Their second phase will also require users to upgrade their operating system.

FAQ: Why are we bothering to go ahead with the trial if we are going to do a U turn and use the decentralised model?

This question has been going around since reports surfaced that money had been allocated to investigate the Apple/Google model. As a technologist, my position would be really straightforward which is that if I've built an App and the operating system is about to be changed in a way that might affect and even positively benefit my App then I'm going to look at it, and although at the time of writing the understanding is Apple/Google are only endorsing the use of the changes in the context of a decentralised model, we are in a fast moving crisis and who knows if Apple won't let the developers of the UK App use the new Bluetooth features to overcome the technical issue discussed earlier? (France also is reported as wishing to have a centralised App and is reported to be specifically lobbying Apple/Google for this. Source BBC News). Also, there is no technical reason not to have two Apps, a decentralised one for those where the App Develop/NHS/Government would never win some users over due to their fear of the privacy risks and a centralised one for those who see the risk reward as worth the increased privacy risk.

Stop Press: As I was checking references for this essay I noticed that it has been announced within the last 24 hours that NHXS are already developing a second version based on the Apple/Google code.

Privacy, Risk and Reward

Although we've kept this essay as non-technical as possible so far, it's worth looking at how risk analysis is done in many businesses especially technology businesses.

The reason we are doing this is because the media loves to portray risk and choice as binary, black and white, but it isn't, and this means that lots of reposts of inaccurate views end up in social media. Even traditionally reputable media outlets increasingly use "click-bait-esque" headlines that actually contradict the body-content of the article they contain.

So it's very common that the media will create a headline that gives the impression that a disaster is waiting just around the corner if we do something, when actually the true likelihood of it happening and impact if it does are very different.

Over the years I've been involved with looking at risk in businesses most often using technology to deliver a service to the public, be it gaming, e-commerce and even e-gambling. In these organisations you have to look at everything, from hardware and software failures, to security breaches and of course privacy.

A well-established technique is to imagine every risk and to quantify it in two dimensions – the likelihood of it happening and the impact if it does happen and draw it on a matrix like the typical one shown below:

		Impact				
		Negligible	**Minor**	**Moderate**	**Significant**	**Severe**
Likelihood	**Very Likely**	Low-Medium	Medium	Medium-High	High	High
	Likely	Low	Low-Medium	Medium	Medium-High	High
	Possible	Low	Low-Medium	Medium	Medium-High	Medium-High
	Unlikely	Low	Low-Medium	Low-Medium	Medium	Medium-High
	Rare	Low	Low	Low-Medium	Medium	Medium

What typically happens in business is that it crafts its business processes around these matrices with the idea of making sure that you contain as much as possible to the green areas by putting working practices and mitigation strategies to move items from the top right as far towards the bottom left as possible.

Alongside this and to help make the judgements as to how much time and money to invest in mitigation strategies, it’s sometimes useful to complement the Impact and Likelihood matrix with a Likelihood and Reward matrix:

		Reward				
		Negligible	**Poor**	**Moderate**	**Good**	**Excellent**
Likelihood	**Very Likely**	Low-Medium	Medium	Medium-High	High	High
	Likely	Low	Low-Medium	Medium	Medium-High	High
	Possible	Low	Low-Medium	Medium	Medium-High	Medium-High
	Unlikely	Low	Low-Medium	Low-Medium	Medium	Medium-High
	Rare	Low	Low	Low-Medium	Medium	Medium

Like so many things though, everyone's perceptions of how to assess something and where to place it in the matrices can vary and even in business meetings where everyone around the table is trying to achieve the same objective it can be difficult. Imagine how difficult things become when dealing with groups that aren't on the same page… Let's just try it with the NHS COVID App.

First here's my personal view and we'll use the risk "a data breach enables a malefactor to steal the data and decrypt the identifiers, correlate them against other data such they can identify by contact circle and exploit that knowledge against us".

		Impact				
		Negligible	**Minor**	**Moderate**	**Significant**	**Severe**
Likelihood	**Very Likely**	🟢	🟨	🟧	🟥	🟥
	Likely	🟩	🟢	🟨	🟧	🟥
	Possible	🟩	🟢	🟨	🟧	🟧
	Unlikely	🟩	🟢	🟢	🟨	🟧
	Rare	🟩 X	🟩	🟢	🟨	🟨

.. and my Reward/Likelihood that complements it which is "I get early warning of contact with a positive and pleasure in the fact that my data may help others."

		Reward				
		Negligible	**Poor**	**Moderate**	**Good**	**Excellent**
Likelihood	**Very Likely**	🟧	🟨	🟢	🟩	🟩
	Likely	🟥	🟧	🟨	🟢 X	🟩
	Possible	🟥	🟧	🟨	🟢	🟢
	Unlikely	🟥	🟧	🟧	🟨	🟢
	Rare	🟥	🟥	🟧	🟨	🟨

So personally I have no issue whatsoever with using the App – I believe that even if my data was decoded and used by a malefactor actually the practical consequences of them using it against me are very low (I'd think totally differently if this was credit card information for example!) , and conversely I think it is very likely participation will either help me directly and in any case my participation may help others and I feel good about that.

But what would your matrices look like? You might use different information to form a judgement and take a different view. Do your matrices look different for the centralised or decentralised model?

A Little Bit of Politics Perhaps – or Genuine Concern?

In this section we're going to look at some of the more "political" aspects of what has been going on in social media as a lot of posts from people that don't want to use the App are arguing why is the government doing one thing when lots of people are advising something else. In addition some commentators are raising genuine concerns about privacy, though sometimes their style of presentation suggests they possibility have an agenda other than genuine concern – you decide!

As the trial has started there has been a lot of social media variously combining the view from the Information Commissioner's Office, the view from the Parliamentary Joint Committee on Human Rights and the open letter signed by 170+ academics many with privacy and information security within their purview.

Cutting to the chase and why we've put this under the heading above, is that one focus of the online propagation of the output from these parties has been made clearly in an effort to affect people's decision whether or not to download and install the App. As a result, a recurring example on social media over the last few days is the views of these parties being used to say that the "App must not be rolled out [full stop]". An instance of where something as simple as leaving out half a sentence can profoundly alter its meaning. As a result this means many people are commenting well in that case they won't bother to download the App at all even in the trial.

However if you return to the source documents that is not what was actually said. For example in the case of the minutes of the Parliamentary Joint Committee on Human Rights they make a number of recommendations of what should happen ahead of the national rollout and what they actually say is "Such an App must not be rolled out nationally unless strong safeguards and protections are in place." Aside from sharing some of the Joint Committee's concerns for sufficient scrutiny the dominant concerns from the 170+ academics is <u>post-crisis</u> "mission creep" of the App and data collection persisting and being used for other purposes, and that in the overall context of the programme nothing is done to make it easier for a malefactor to correlate data such that the anonymous IDs can be de-anonymised.

I can understand their position, let's imagine we are sitting on a committee responsible for safeguarding privacy of data at the scale of the UK population.

If I was charged with that responsibility I can see how we might move the X from the bottom left towards the upper right, especially when you consider that the presence of population wide data starts to make it more attractive to those who actually have the resources to execute and exploit a hack and decrypt and correlate the data; for example the security services of foreign powers. Another thing you most definitely would look at is misuse of the data by people that have legitimate access to it (there are well established techniques that can be employed to mitigate internal attacks).

So in this case my X goes in the "Possible" row now – however I'm not quite sure how significant the impact would be. Damning maybe if it were to happen but that isn't one of our categories. I'm putting it in "Significant", however to be honest it depends on what the malefactor manages to do with it – it may be nothing in which case we should use the "Negligible" column or it may be that they use it to try to get circles of contacts to influence each other which puts it in Significant; but maybe the malefactor can already do that by hacking Facebook so having this data isn't useful at all? Maybe intelligence led information has just come in and we now think a state sponsored hacker has their eyes on this – the likelihood has gone up.

So you see how difficult quantifying some risks and impact can be – so let's err on the side of caution and go with "Likely" and "Significant", but arguments could be made for the X to be almost anywhere in the top right quadrant of the grid.

Risk "a data breach enables a malefactor to identify the contact circles of the entire population using the App for purposes unknown":

		Impact				
Likelihood		**Negligible**	**Minor**	**Moderate**	**Significant**	**Severe**
	Very Likely	🟩	🟨	🟧	🟥	🟥
	Likely	🟩	🟩	🟨	🟧 X	🟥
	Possible	🟩	🟩	🟨	🟧	🟧
	Unlikely	🟩	🟩	🟩	🟨	🟧
	Rare	🟩	🟩	🟩	🟨	🟨

To do a full job of course we'd have to analyse not just the example above each risk that has been identified, for example the risk posed by post-crisis "mission creep", which was one of the main concerns of the 170+ academics.

Just as a business would look now at ways of mitigating the risks, the Joint Committee recommendations (Efficacy and proportionality, Primary legislation, Oversight, Child safeguarding, Efficacy review and transparency) are similarly intended to move the X leftwards and downwards in the matrix ahead of the nationwide rollout.

As well as the privacy issues, the Committee also addresses human rights, and it is here it adjudges "the lockdown measures are a significant interference with the right to family and private life, (Article 8 of ECHR)" It also says that digital contact tracing would be able to ease the lockdown but that it in "itself will interfere with the right to private life (Article 8 ECHR)".

Opinion: One hopes that a common sense of the greater good will prevail, however we have concern that the Joint Committee's view that digital contact tracing might "interfere with the right to private life (Article 8 ECHR)" coupled with the legal opinion obtained from Matrix Chambers on behalf of the Open Society Foundation might tempt human rights lawyers to derail the App.

We haven't commented separately on the ICO statements, however the ICO did provide expert testimony to the Committee which we have dealt with.

In respect of the open letter from 170+ academics, we hope that this essay demonstrates that whilst we have adjudged the App at little personal risk, we have shown how an academic/techie honestly answering the question would never say there are no risks and could indeed speculate on further risks which would also not have zero likelihood of happening, which when translated one dimensionally through the media can create headlines that would likely engender the impression that the App is insecure.

One other comment which may give food for thought is that prior to GDPR it is arguable that an encrypted unique ID was not "Personal Data" under the old data protection act, whereas it is under GDPR and because it could if compromised and correlated with other data identify an individual, is subject to data privacy rules. Interestingly this means it can be argued this App uploads other peoples' "personal data" as well as your own – we haven't seen wide coverage of that but again it is very likely that we may see in coming days this exploited by those wishing to scaremonger – "Your friends and even people you don't know have got your personal data and can upload it to the central server". Comment: We don't care that someone else is uploading our "Personal data" in the form of an encrypted anonymous ID to the central server in the interests of tracking the virus. Dealing with this may be as simple as asking consent in terms of use for the App. On our to-do list is to further evaluate this as the trial and roll out proceeds.

To exhaustively explore all of the privacy ramifications one could almost write a thesis but hopefully the "takeaway" from this is to always look at source documents not click-bait-esque headlines even in the reputable media and most definitely not social media to make judgements!

Finally in this section we need to revisit the possible exploits of the data including the ability to identify a unique user at the time the contact is recorded. We're putting it in this section because some of the examples you may think are quite fanciful (though nevertheless theoretically possible and as per our discussion on risk, it makes sense to imagine as many risks as possible so they can be analysed and mitigated).

One example given is that a malefactor is in an environment where they might be able to identify you through other means and deliberately trigger a Bluetooth contact, and record your unique ID using scanning equipment. As per our earlier description the encryption is rotated so it is unclear in practice if and for how long material damage could be done to you as a result however one commentator has given the risks associated with this and similar exploits as potentially causing the user panic, social stigmatisation, adverse health outcomes or could exacerbate commercial or other forms of tracking.

Let's just think about that for a moment. Someone in a public environment where they can also obtain my real world identity (maybe I hand over a debit card to pay for something or have some other ID they can exploit like an RFID card) steals my encrypted ID.

For this to be a risk related to the App it has to be something they can now do by knowing our phone's unique ID (presumably that they have also managed to decrypt) that they can't already do by knowing our real life identity and our juxtaposition with the place they have done the intercept. If it is about correlating data to the central server, just what proportion of the central server data of millions of people are they going to actually manage to correlate in that way? What is the damage to us individually if they do?

We don't want to appear unduly facetious (actually we do), but here's our grid for some of the more fanciful risks and impacts we've seen:

		Impact					
Likelihood		**Babel Fish***	**Negligible**	**Minor**	**Moderate**	**Significant**	**Severe**
	Very Likely	🟢	🟢	🟨	🟧	🟥	🟥
	Likely	🟩	🟩	🟢	🟨	🟧	🟥
	Possible	🟩	🟩	🟢	🟨	🟧	🟧
	Unlikely	🟩	🟩	🟢	🟢	🟨	🟧
	Rare	🟩	🟩	🟩	🟢	🟨	🟨
	Infinitely Improbable*	🟩 X	🟩	🟩	🟩	🟢	🟨

*With apologies to the late Douglas Adams and HHGTTG fans who might suggest that infinite improbability might result in the impact of a sperm whale and bowl of petunias on Magrathea (that wouldn't fit in the box)

Talking seriously again, it's also worth mentioning that risks associated with the App raised around intercepting the Bluetooth exchange of the encrypted IDs apply to both the centralised and decentralised models and those citing such arguments are actually making an argument against digital tracing, not just against the centralised model.

Other examples cited require deliberate abuse of power by the developers or organisations associated with them. If anything this is the most likely route to a compromise of the data, and in this regard the Joint Committee have already made sound recommendations.

We leave it to you as to whether such arguments have practical merit or purely theoretical merit at the current time, such that they would modify your view of the safety of using the App.

The real message we want to get across here, is don't take everything at face value, make your own risk assessment. A lot of things painted as the end of the world are a storm in a teacup. People whose role it is to identify and report risks will identify and report risks.

The App is being developed in a time of crisis where time is of the essence and whilst this certainly doesn't give carte blanche to ride roughshod over privacy and human rights legislation it is important to appreciate the accelerated time scales involved in both the App development and privacy compliance. Risk shouldn't be seen as one dimensional and it's important to look at both the likelihood and impact of risks in the context of a single individual, the Isle of Wight Trial, and a national rollout and to remember many risks apply to both the centralised and decentralised approach. We should contrast those risks with the benefits, whilst at the same time implementing the necessary measures to mitigate them over the lifetime of the programme.

Other Worries

Finally, we wanted to touch on some other worries we have seen on social media. Again this section is our opinion rather than a technical view.

FAQ: I don't believe it uses Bluetooth but secretly sends my location and other data.

Our view is that there are many security specialists already analysing the operation of the App and its source code has been disclosed publicly inviting scrutiny. We are confident that the App works using Bluetooth, does not pass location data other than the first half of your postcode, and does not pass other data that it shouldn't.

FAQ: GCHQ are involved. The App is "a vehicle of state oppression"

We suspect we won't be able to persuade people who believe this to change their mind. There are two schools of thought here, one which is that the involvement of the NCSC and GCHQ means that the security and encryption is robust and this should in turn give confidence in the App, and the second which is that their involvement is a shady government scheme of some kind. We know what we think and to be honest even if GCHQ do see the data, at an individual level what are they going to do with it!

FAQ: “What if a malefactor can get access to the data”.

It is true to say that the most likely cause of a data breach is through correlation of the central data with other data used in the broader track and trace programme with a less probable subsidiary threat from a sophisticated hack.

In many ways we have already covered the protections and mitigations in discussing the recommendations of the Joint Committee for Human Rights earlier.

Regardless of the likelihood the potential worry for people is what will be the impact. If your data is exposed in the midst of a sea of similar data, it is difficult to conclude that the impact on an individual would be anything other than very minor. A largescale breach is more likely to give a picture more suited to “big data” processing rather than to attack any individual (unlike say a hack of a commercial company for names, addresses and credit card details), so our assessment for the risk to an individual remains low in the absence of information to the contrary.

FAQ: I don’t wish to use the App as the contract was awarded without due process and I don’t want Dominic Cumming’s mates to have my data.

Our view is it was awarded to a pre-existing Government IT contractor as a matter of expediency and it has rapidly resulted in an App coming to trial and ahead of the Apple/Google solution at a time of national crisis when every day may count. When developing software in our own experience that had urgency we would naturally go to people we would know could do it who were immediately available. We’re sure there will be retrospective scrutiny of the process but respect that people may not wish to download the App on the basis of their feelings about this.

Conclusion

We hope that readers have found this essay useful in explaining how the App works, the strengths and weaknesses of the centralised versus decentralised model and the potential privacy issues that surround them.

Our view is that a digital App operating within the centralised model undoubtedly brings the best potential for getting on top of the virus especially as part of an overall track and trace programme with notable benefits over the decentralised model including hotspot tracking and identifying “Typhoid Marys” (asymptomatic carriers).

The App is being developed in a time of crisis where time is of the essence and whilst this certainly doesn’t give carte blanche to ride roughshod over privacy and human rights legislation it is important to appreciate the accelerated time scales involved in both the App development and privacy compliance. Risk shouldn’t be seen as one dimensional and it’s important to look at both the likelihood and impact of risks in the context of a single individual, the Isle of Wight Trial, and a national rollout and to remember many risks apply to both the centralised and decentralised approach. We should contrast those risks with the benefits, whilst at the same time implementing the necessary measures to mitigate them over the lifetime of the programme.

We personally have already installed the App and have no concern in using it both during the trial and during a nationwide rollout as we believe the benefits outweigh the risks, however we recognise that some may assess the individual privacy risks differently from us and choose not to use the current App on trial in the Isle of Wight, in which case there is merit in the UK <u>also</u> pursuing a decentralised version built on the Apple/Google platform.

References:

NHSX NHS page and App download page, and FAQs about the App

https://www.nhsx.nhs.uk/covid-19-response/nhs-covid-19-App/

https://covid19.nhs.uk/ (works on the Isle of Wight trial only at present)

https://faq.covid19.nhs.uk/

Ian Levy blog from the NCSC

The security behind the NHS contact tracing App

https://www.ncsc.gov.uk/blog-post/security-behind-nhs-contact-tracing-App

Reincubate blog

Staying alive: background tracing & the NHS COVID-19 App

https://reincubate.com/blog/staying-alive-covid-19-background-tracing/

Apple/Google

Privacy-Preserving Contact Tracing

https://www.Apple.com/covid19/contacttracing

Parliamentary Joint Committee on Human Rights

Human Rights and the Government's Response to COVID-19: Digital Contact Tracing

https://publications.parliament.uk/pa/jt5801/jtselect/jtrights/343/343.pdf

ICO document on expectations on App development

https://ico.org.uk/media/for-organisations/documents/2617676/ico-contact-tracing-recommendations.pdf

Open Letter by 170+ Academics

https://drive.google.com/file/d/1uB4LcQHMVP-oLzIIHA9SjKj1uMd3erGu/view

Legal Opinion from Matrix Chambers

https://www.matrixlaw.co.uk/news/legal-advice-on-smartphone-contact-tracing-published/

Analysis of the NHSX Contact Tracing App "Isle of Wight" Data Protection Impact.

https://osf.io/preprints/lawarxiv/6fvgh

NHSX App Source code

https://github.com/NHSX

Version Changes

V1.0 published on 9th May 2020

V1.1 published on 11th May 2020

- Typographical and minor text tweaks and layout corrections
- Improvements to the centralised model description and diagram to correct an over-simplification.
- Addition of the link to the official NHS FAQ in the references

V1.2 published on 11th May 2020

- Rewording of information about the Apple/Google approach to remove a possibly incorrectly presumed dependency of the contact tracing API on features of a new operating system.