

Multiple Objectives of Lawful-Surveillance Protocols (Transcript of Discussion)

Joan Feigenbaum(✉)

Computer Science Department, Yale University, New Haven, CT, USA
joan.feigenbaum@yale.edu

It's great to be back in Cambridge.

I want to talk about the supposedly competing objectives of personal privacy and national security. More generally, I'm interested in the alleged tension between proper handling of sensitive data and pursuit of criminals and terrorists. Many people claim that these are irreconcilable objectives, but I don't believe that they necessarily are.

The topic fits into this year's SPW theme of multiple-objective security.

Let me take you back to three years ago, when we had the first Cambridge Security-Protocols Workshop after the summer of Snowden. At that 2014 SPW, I gave a talk that was essentially an angry lament about the mass surveillance that had been revealed by Snowden in tremendously dramatic fashion. One of the bullets that I had on my slides read "all around catastrophic failure of institutions and individuals."

That was my description of the surveillance morass. One of the "institutions" that I singled out as having failed was us: the crypto- and security-research community. The failure that I was lamenting in that talk – our failure – was that we as a group had not really stepped up and made forceful, principled statements opposing the kind of mass surveillance that Snowden had revealed. That was true at the time; it's a little bit less true now. My co-author Jérémie Koenig and I said in that paper that the antidotes to mass surveillance and passive acceptance were mass encryption and active protest. When I finished my lament, Jérémie took over, explained that the "feudal Internet" (a term that I believe was coined by Bruce Schneier) had enabled the surveillance morass, and advocated grass-roots, decentralized cloud services.

That was I at the 2014 Security-Protocols Workshop. Later the same year, another I (actually not literally I but Bryan's and my graduate student and co-author Aaron Segal) presented a paper¹ on "privacy-preserving surveillance" at the 2014 USENIX FOCI Workshop. The view expressed in that paper was *not* "All surveillance is terrible. There shouldn't be any." Rather, it was "There are bad guys out there, and there is a role in society for law-enforcement and

¹ A. Segal, B. Ford, and J. Feigenbaum, "Catching Bandits and *Only* Bandits: Privacy-Preserving Intersection Warrants for Lawful Surveillance," in *Proceedings of the Fourth USENIX Workshop on Free and Open Communications on the Internet*, 2014.

intelligence agencies; they have to surveil some of the people some of the time.” Of course, that does not mean that they are supposed to surveil everybody all the time, which is what I had been lamenting at SPW 2014. Is there a privacy-respecting way for them to identify targets? Not just to surveil known targets but to *identify* targets and to obtain actionable information without intruding on all of the rest of us.

Back in 2014, I assumed that the knee-jerk reaction of anybody in the crypto and security community would be “Sure there’s a way.” In the framing that I used in my 2014 SPW talk, this is certainly one front on which our community has not failed at all. For decades, we’ve been doing work on secure multi-party protocols, private information retrieval, and related cryptographic techniques for mining a single pertinent fact from a distributed data set without learning anything else about that data set. Enabling law enforcement and intelligence to obtain one particular piece of information that they need but not have everybody’s private data revealed to them: That’s exactly what we know how to do, right? We’ve certainly been claiming we know how to do that for many years.

In our 2014 FOCI paper, we considered lawful, accountable, privacy-preserving surveillance. The idea is to combine cryptographic protocols (SMC, PIR, *etc.*) with black-letter law. (I’ve always thought that was an interesting phrase, because it sort of sounds like the opposite of what it is; but, apparently “black-letter law” means known, open, public processes that the voters can read, can understand, and then can challenge through the political process. That’s what we want.) More generally, we advocated combining technical protocols with legal and social protocols, and we sought to build into these protocols the kind of things that are always recommended in crypto papers: Limit the scope of the warrant; distribute the power to authorize a surveillance operation so that no one party has too much power over who gets surveilled and who doesn’t; build in oversight, *e.g.*, public reporting of statistics on how much surveillance has been authorized. What could a FOCI audience possibly think was wrong with that?

Our motivating example in the 2014 FOCI paper was the so-called “high-country bandits” case. In 2010, there were three old-fashioned F2F bank robberies in Arizona and Colorado. In three banks in three different towns, a gang of bank robbers with guns stormed in and said “Stick’em up. Give me the cash.” They took the cash and left – I don’t think there were any fatalities. Because it was an inter-state crime, the FBI pursued it, and they got a tip at some point that one of the bank robbers was talking on a cell phone during the heist. So the FBI got three cell-tower dumps: all of the cell-phone numbers that sent or received calls at the times of the robberies via the cell towers nearest to the banks that were robbed. That’s metadata on a total of about 150,000 users. The FBI intersected the three sets, and it turned out that the intersection contained a single phone number. They went to the carrier that served that number, got the name and address of the customer, and arrested him. Sure enough, he was one of the robbers; he ratted out his friends, and that was the end of high-country bandits.

Success, right? When I get to that point, most people say, “Wow that’s great.” But I say, “Well, they succeeded in catching the bandits, but there’s a problem. What about the 149,999 innocent bystanders whose cell-phone metadata were sucked in by the FBI?” Perhaps the FBI claims that it has destroyed those data; I don’t know whether they made that claim, but we should not just take their word for it even if they did. Moreover, the possibility that data about innocent people – people never even charged with crimes, much less convicted of them – can be retained and misused by law enforcement is not a hypothetical problem. Remember the controversy about stop and frisk in New York. The NYPD had a policy of stopping people on the street who they thought “looked suspicious” and frisking them for weapons, drugs, or something else illegal; the police department claimed that stop and frisk was a success and pointed out that many of the people on whom they found something illegal turned out to have outstanding warrants against them. Ultimately the policy was declared by a judge to be unconstitutional; the main objection to it was its disparate racial impact. But there was at least one other major problem that the judge and civil-rights advocates called attention to: The police department was keeping records of the descriptions, names, and addresses of completely innocent people who were stopped, frisked, discovered not to be doing anything wrong, and let go without being charged. A lot of these people were later questioned when a crime occurred and their descriptions matched the suspect’s. Once the information was in a government database, there was a significant chance that it would be used. So we want to keep irrelevant personal data out of government databases.

It’s important to note that our motivating example in this work is a scenario in which law enforcement *already has lawful access to a lot of personal data*. We’re not saying, “Let’s enable pursuit of criminals and terrorists by using fancy protocols to justify law-enforcement access to data they’re not getting now.” Rather, we’re asking, “How can law enforcement perform the same task that it is performing now in a less privacy-destructive manner?” In particular, how can privacy of innocent bystanders be maintained?

It’s also important to note that the goal in this scenario is to identify a target (or a small set of targets). The government seeks access to sensitive data about a large set of people; most of the people are “untargeted” data subjects or, as I’ve been calling them, innocent bystanders. A small number (perhaps just one, as in the high-country bandits case) are targeted data subjects – there is evidence that they have committed crimes or, more generally, are directly relevant to an investigation. Why can’t the government agency just get a warrant for data on the targets? Because these are “unknown targets.” That sounds like an oxymoron (how can someone be both “unknown” and “targeted”?), but it is not. The government can describe the data subjects very precisely by saying that they are the only people who were at these k places at these k times; if the targets happened to be using cell phones or otherwise leaving electronic bread crumbs that mark their presence, then intersecting k sets of data should identify them. But the government does not have any PII on these people when it goes to get a warrant; it can’t get a warrant to track “Mr. Smith’s phone” or to track

“phone number 917-359-4081.” It needs what’s referred to in the US as a John Doe warrant: permission to track or search a person or people who fit a precise description that has been presented to a judge but whom the agency seeking the warrant cannot (yet) identify. Technically speaking, we want to identify the unknown, targeted subjects of a “John Doe” warrant *without* identifying all of the unknown, untargeted people whose data may be indistinguishable from those of the targets before the warrant is executed.

In our FOCI 2014 paper, we pointed out that the high-country bandits could have been captured without the FBI’s obtaining cell-phone numbers of any untargeted users. Privacy-preserving set-intersection is a well studied problem. In this particular application, we used a variant of the Vaidya-Clifton protocol; it works on sets of *encrypted* data, outputs the cleartext of data items in the intersection of all of the input sets, and leaves the items not in the intersection encrypted. Our proposal was to have repositories store call records in encrypted form – specifically to encrypt them using the public keys of multiple authorities all of whom would have to participate in the execution of the set-intersection protocol in order for it to run to completion and decrypt the records in the intersection; that is equivalent to “distributing the power to authorize a surveillance operation,” which is one of our design principles. Aaron Segal implemented the protocol and found that it could handle a set of 150,000 encrypted call records very efficiently.

Beyond privacy-preserving set intersection, we advocated surveillance regimes that obtain a large set of encrypted data about both targeted and untargeted users, feed it into a cryptographic protocol that winnows it down to the records of users targeted by the John Doe warrant, and decrypt only those records. As I said earlier, how could a FOCI audience object to any of that?

In fact, many in the FOCI audience objected to everything we proposed. We’ve heard similar objections from many people in the crypto- and security-research community in the intervening three years. Now that I know that the our ideas are so controversial, I figure that a Cambridge SPW is the perfect place to present them.

Frank Stajano: Those operators: Do they have access to plain text?

Reply: I’m not sure exactly what you mean by “operators.” In the bandits example, the plain texts are phone-call records; they’re produced by the phone networks as a byproduct of network operation. So, of course, the phone companies have access to them. But no one else need have access to them. The encrypted records that may be subpoenaed can be stored by the phone companies themselves or by neutral repositories. Our framework is very general: Large sets of sensitive data about both targeted and untargeted data subjects should be encrypted by multiple authorities and used in cryptographic protocols only if all of the authorities agree that a legitimate warrant has been obtained. Who stores the encrypted records will depend on the use case.

Frank Stajano: Why can't the FBI just say to a phone operator "give me the records of the people who were there at the time"?

Reply: That's exactly what the FBI did, and it wound up getting the records of all of the untargeted phone users as well as those of the single target. We're trying to prevent that.

Paul Wernick: If you get the information from the phone companies, won't they know that the person you're looking for is the target of an investigation by the security services?

Reply: They'll know that the target may be identified by *one* element of the large set of data that they supply, but they won't know which one.

Fabio Massacci: I think there's a big error. You need to tell the phone operator what it's going to be used for. They have to encrypt the raw data in a way that makes it usable in the protocol.

Reply: Good, you're making an important technical point. Data that might be collected in bulk in encrypted form and fed into a cryptographic protocol by government agencies in order to identify targets (or, more generally, to discover something about *a priori* unknown targets) need to be encrypted in a particular way. They need to be prepared to serve as input to a specific privacy-preserving protocol. Notice that I mentioned ElGamal encryption on a previous slide; that's because our set-intersection protocol uses ElGamal and Pohlig-Hellman. This requirement that repositories store encrypted data in precisely the form that will be needed by the protocol that the authorities later execute implies that the government must know in advance what operations it plans to perform on the sensitive data in question. That may ultimately limit the applicability of our whole approach. Still, we do know that set intersection is used regularly by law enforcement and by the NSA. They could be using it in a privacy-preserving manner instead of the privacy-invasive manner in which they're now using it. The same applies to some other common operations.

Go ahead Ross.

Ross Anderson: The problem is that, under the UK Investigative Powers Act, for example, we get something called Internet connection records, which phone companies must support with cleared staff and appropriate enclaves, and these give not just intersections but joins. What the UK government now entitles itself to do is to say, "Tell me all of the websites that have been visited both by wicked Dr. Sierra and wicked Professor Anderson and then show us everybody else in Britain who visited those websites in the last month."

Reply: So that's *this* objection to the entire approach (speaker points to a bullet on her slide), which is that it's a non-starter politically.

Ross Anderson: You should have sold this in the 1990s. (Laughter)

Reply: (Turns to another member of the audience.) Yes, go ahead.

Partha Das Chowdhury: In a place like Kashmir, when you do operations, you have loads of people coming from the street, throwing stones, and preventing the police or the army from targeting the militants. So whatever they (the police) do in such a situation, they do to everyone. They can't target anyone. The same applies in other places; we have loads of Maoists in the eastern and central parts of India. In those places, how would you define your target? Because the young soldier you are sending down, he has a family to take care of; he's not going to listen to "privacy, etc."

Also, government agencies don't want large quantities of data. Forensics is very slow ... getting data, analyzing them. The agents are not willing to deal with large amounts of data.

Reply: It sounds like you're actually saying two different things, the first one of which isn't directly relevant to our framework. We're addressing precisely the scenario in which there *is* a specific target, but it's an unknown target, and the only apparent way to get the data on that target is to sort through a much larger superset. You're saying that there are situations in which you cannot properly target. Fine. There may be situations in which you actually need to surveil a crowd, because there's a threat from an entire crowd. Those are not the situations that are addressed by our framework.

Now to your point that government agencies don't want a lot of data. Perhaps *some* government agencies don't want a lot of data. The US intelligence agencies seem to want every bit of data in the whole damned world.

Partha Das Chowdhury: Forensic tools are very slow; you don't want to use them on a large data set when you're trying to investigate.

Reply: Well, that's one of the arguments that computer scientists make to intelligence agencies: No, you shouldn't be sucking in every single bit that's sent anywhere on the Internet, because you don't have computational techniques that can actually use all those data. But that has been a tough sell. Anyway, in our high-country bandits case, we're not talking about big data: We're talking about 150,000 records; so what?

We skipped over one slide, but I don't think it's all that important. We are actually having the discussion that I wanted to have. Since FOCI'14, I have been very surprised by the knee-jerk negative reaction to the idea of "privacy-preserving surveillance" and by the immense technical pessimism about the possibility of using cryptographic protocols to simultaneously enable legitimate pursuit of targets and privacy of non-targets. I even heard Ron Rivest poo-hoo the idea; he said "I don't know whether exotic protocols should be used for law enforcement and intelligence." These aren't "exotic protocols"! I saw a DIMACS Workshop talk about efficient privacy-preserving set intersection more than 20 years ago! I've had this experience before, often in discussions of Internet voting. Members of my own research community sort of nay say the idea of using our own tools for some social objective by saying, "oh it's never going to work." It drives me crazy.

Ross Anderson: I think that's an unfair criticism of Ron, because he, like I, was an author of the "Risks" report in 1998. When a bunch of people in the crypto-research community were quite happily proving mathematical theorems, we were prepared to roll up our sleeves and get engaged in the struggle. Those of us who have been on the front lines know what this is like. We know what's even remotely likely to fly, and we know what simply doesn't have a chance. So I support Ron 100% on this.

Reply: So ... back to this bullet? (Points to a bullet on her slide) Because technical ...

Ross Anderson: Because Ron, like me, has been in dozens of contretemps with policemen – in the first crypto war, in the second crypto war. We've been in private meetings in Washington, we've been private meetings in Brussels, we've been sitting down with industry people. You know, we've got the form. We know this fight, we know what it's about. It's not about technical control mechanisms. It's about policy.

Reply: Okay, so then you're not saying what Ron said to me in that particular discussion, which was that SMC protocols *per se* are just not usable for this purpose. That nobody could ever implement them and deploy them at this kind of scale or in this kind of situation. You're saying that, politically, they won't fly.

Ross Anderson: Well, it's not where the fight is. You see, no sensible person has got an objection to ... Remember the case here in Britain about ten years ago, in which a man conducted a number of rapes in the east end of London? Nobody was bothered about the police going in and taking cell-tower dumps from the relevant places until they found the guy. In those days, however, if you took a cell-tower dump, it involved serious manual labour by phone-company employees, and it cost the police tens of thousands of pounds. What's changed is that, now, this has all been automated. Government has spent hundreds of millions of pounds in up-front cost to ensure that the marginal cost of getting cell-tower dumps is basically zero.

Reply: All right. That was not Ron's objection in the conversation I had with him, but I understand your objection.

(Turns to another audience member.) Yes. You have to identify yourself.

Tuomas Aura: You are treating surveillance like a logical proposition. I think you have to understand that the ability to find someone depends on data mining and technical forensics. If you are the police desperately looking for clue, you look for the link in forensics. It's based on statistical likeness.

Reply: What you just said is actually related to what you said. (Points to Fabio Massacci.)

Tuomas Aura: It's brittle.

Reply: Yes, there is brittleness here and potential limitation of applicability. That's not what most of the controversy is about. It actually is not true that

privacy-preserving data mining, even privacy-preserving data mining that uses cryptography, can only be applied to deterministic, well defined functions. There is a lot of ongoing work right now on privacy-preserving statistical data mining.

Tuomas Aura: It's ongoing but there are no complete solutions.

Reply: I'm not at all saying this framework is universally applicable. Its Achilles heel might be that, ultimately, there are not robust enough methods for pre-processing data to both hide them – make the whole end-to-end operation privacy preserving – and actually use them for enough data mining operations that anybody would bother.

Tuomas Aura: If you say that this is the limit on the low end, then, unfortunately, these agencies have only one function that they can use.

Reply: No, I'm absolutely *not* saying that.

Tuomas Aura: You will be blocking more creative methods of mining the data.

Reply: First of all, I'm not saying if this framework doesn't apply, then law enforcement cannot do anything. Nor is it true that there's only a small number of deterministic operations that this will work for. But nonetheless, you're absolutely right that, for this whole thing to be interesting, we have to demonstrate its applicability to more than set intersection, contact chaining, and graph searching.

By the way, the differentially private graph-searching algorithms to which these techniques apply *are* probabilistic algorithms.

Okay! I think we're now having the discussion I wanted to have, in which the question is "why this combination of hostility and pessimism?" George Danezis – I was hoping he was going to be here . . . too bad he isn't – burst out of his chair during a workshop talk that I gave and said, "This is really interesting cryptographically, but we shouldn't be working on this stuff! We should be fighting for a world in which there is no surveillance. By anybody, of anything, anytime." I was very surprised that anybody would say that; of course, that was a PETS workshop, and PETS people say things like that all the time. But what has subsequently been revealed to me is that these last two bullets (points to a slide) are a less emotional way of saying the same thing.

On another front, Paul Syverson coined the term "function drag," which is supposed to be the evil twin of "function creep." He cautioned that, someday, phone companies may be able to provide service and bill customers without retaining any records of individual calls. Paul is worried that, if we put in place a system that enables data mining of encrypted phone-call records, the FBI will get used to mining all of those records and won't let a memoryless phone system be built.

Frank Stajano: We can take one more I think.

Reply: Right.

Jonathan Weekes: So, if I understand the process correctly, the phone company collects the data and encrypts it; if the FBI needs it, they get the encrypted data.

Reply: The general framework includes a number of different parties that have different roles to play in accountable, privacy-preserving surveillance.

There are entities that create (or otherwise acquire) sensitive data in the normal course of doing business. In the bandits use case, those are cell-phone companies whose networks create call records in their normal course of operation.

There are repositories that store encrypted data. They could be the same entities that create the data, but they need not be in every use case.

There are multiple authorities that must agree on the legitimacy of a request for data and participate in a cryptographic protocol that ultimately decrypts only a subset of the data. This is how power is distributed so that no one person or government body has too much say over which sensitive data are revealed and when. In the bandits use case, which occurred under the American legal system, it is natural to think of these authorities as employees of different (even competing) government bodies or officials of the three different branches of government. But that's not essential in our framework; the requirement is simply that power be distributed in the sense that it is vested in multiple independent parties. The parties don't even have to be parts of the government.

The data that are stored in repositories are encrypted under the public keys of all of the authorities; that is why all of the authorities must participate in the protocol if the targeted data records are to be decrypted. Encryption of the original cleartext data records might be done by the entities that created them or by some other party – it depends on the use case.

So at the beginning of the process, someone in law enforcement or intelligence requests access to data. In the bandits case, the requester would be an FBI agent, and he would request the intersection of three cell-tower dumps. The requester must go to a judge and get a warrant. If the judge grants the request, he may put restrictions on the warrant; for example, in the bandits case, he might say that the number of records ultimately decrypted and sent to the FBI must be small – at most the number of robbers that held up the banks. The requester submits the warrant to all of the authorities, who must authenticate it. If they all agree that it is a legitimate warrant, they execute the protocol, verify that the output satisfies whatever restrictions the judge imposed, and send the decrypted records to the requester.

I think we don't have any more time for questions now, but we have the whole rest of the workshop for discussion.

Thank you!

<http://www.springer.com/978-3-319-71074-7>

Security Protocols XXV

25th International Workshop, Cambridge, UK, March

20–22, 2017, Revised Selected Papers

Stajano, F.; Anderson, J.; Christianson, B.; Matyáš, V.

(Eds.)

2017, XI, 307 p. 19 illus., Softcover

ISBN: 978-3-319-71074-7