[Intro Slide]

My name is Eleanor Saitta, and I'm going to talk today about surveillance.

How Cities Are Watched

There are many kinds of surveillance and tracking that a modern city dweller interacts with. We have been watching people for a very long time, and every new form of media, every new technology, eventually begets its own kind of surveillance. The list you see here is almost necessarily incomplete.

[Next Slide: Surveillance Type Slides (7)]

To make this a bit more real, let's take a minute to look at how many ways an average city-dweller interacts with surveillance and monitoring systems.

Narrative

You wake up and turn on the lights and the smart electric meter notes the change in power use. While you make coffee, you check your web mail. Your ISP screens and stores all of your web traffic. You reply to an email, and an intelligence agency stores it and keyword-matches the contents. Your mail provider analyses the information and sells it to an ad company. You turn on your phone and it registers your location with the E911 database; the phone company also records it. The intelligence agency cross-references the two new text messages that come in with your email. You take a shower and the gas and water meters note the usage. You get dressed and get on the subway to go to work. Along the way, you appear on cameras run by your landlord, the transit agency, several nearby merchants, the police, and the traffic agency. The time you swipe your card to enter the subway is stored, along with your identity. While you wait on the platform, you get a call from your boss. The numbers, time, and duration are logged for billing and possible law enforcement use. The phone company doesn't capture the voice data, but your employer does, as does the microphone attached to the camera you're standing under.

The train arrives, and you walk by an air-quality monitoring box as you get on, whereupon the transit authority tracks the train's movement while onboard cameras record you. As you walk out, the destination of your journey is noted on your metro card, and separately, the turnstile increments the number of people leaving the station in this fifteen-minute period.

On the way into the office, in addition to passing through the vision of a couple dozen more cameras, you stop to buy a croissant. You pay by credit card, and the transaction is recorded by the merchant, credit card processor, the credit card company, and the issuing bank, the last three of whom all run it through fraud detection systems and sell the data to marketers. Your mobile registers your changing location as you walk past a weather station, and its radio signal is briefly evaluated by an RF threat detection system run by the national security agency. A cab, with its own video system and GPS tracker drives by, and plate and face detection systems recognize the car and driver. It backfires in a frenzy of data collection for its emissions system, and a nearby gunshot-detecting microphone evaluates the noise before throwing it out as a false positive. The car itself is logging speed and engine data, plus acceleration, which the police or insurance companies may use after an accident. It drives over a temporary traffic load sensor and past a red-light camera, and, less obviously, a chemical sniffer looking for explosive residue.

You walk into your office, through an alarmed door and past a window with a glass-break detector, and also past a heat rise/smoke detector. You show up on another camera as you swipe your badge to get in, which is logged, as is the time you log in to your computer. You load up a web page, which is checked by a corporate web filter.

You've been awake for less than two hours, and half a dozen government agencies have logged data about you, plus at least three times as many private organizations.

[Next Slide: Structures Around Surveillance]

Analysis

Now that we have some idea of the scope of the issue and just how pervasive it is, let's take a step back and look at some of the structural issues around surveillance to better understand its position within the larger socioeconomic frame.

Secondary Uses

One of the distinct complications of surveillance work is the same property of digital information that makes it fundamentally a more useful medium for its intended purpose: it can easily be retransmitted and aggregated. The physical installation of sensors is visible and obvious, but once those sensors exist, repurposing that data can be done in a much less visible manner. For instance,

when a municipality decides to install a camera system, they may first present it as a "traffic" camera, theoretically only placed to watch for accidents and review traffic levels. However, the police often very quickly shift those same cameras over to more general use. This has been true from the first CCTV installations in Munich in 1958 onward. This pattern repeats for many different technologies—mobile phone location tracking, for example. Originally, in the US, the E911 location reporting feature was designed to let emergency operators know where users with mobile phones were located, something that was becoming an increasing problem. However, police and intelligence agencies rapidly adopted it. Last year, one US mobile phone operator reported more than 8 million cell phone GPS location data requests from the police, not including non-GPS tower-based location data requests, 911 calls, intelligence agency operations, or civil lawsuits.[1]

Buying and Sharing Data

The largest forms of secondary use all involve buying, selling, and sharing data. This is, of course, a fundamental feature of a modern consumer society. For instance, the very concept of a credit rating is founded on companies sharing financial information with each other, through mediating rating agencies, for both consumers and enterprises. This same thing now happens with other, less clear-cut information, like the results of data profiling performed by marketing agencies. An amazing amount of information is included in these profiles, down to mood data reported by voice stress analysis features built into interactive voice response systems—your profile with some data clearinghouse may literally say that the last time you were on the phone with your credit card company, you sounded upset. In much of the world, these data clearinghouses are largely unregulated.

Similar things exist within the state—in the past ten years, governments have built large inter-agency databases, and they collect all sorts of information, not always just about convicted criminals. National intelligence agencies have also taken to buying the databases of civilian clearinghouses, as they can't economically collect all of that data on their own. In addition, this allows them to get around rules about who they can and can't watch, but also means that they get information that hasn't been properly vetted and that was never meant to be used as evidence—like those voice stress systems, really just intended to help with call routing.

---

[1] http://www.eff.org/deeplinks/2009/12/surveillance-shocker-sprint-received-8-million-law

Sunk Cost

This kind of data sales and the entire secondary use system are driven by, among other things, the expense of gathering data. Once a city has spent millions of dollars on a CCTV system, there is immense pressure on the city to make the most of that investment, and this leads to second uses of that data being set up. These secondary uses may not have much direct benefit to the general population, but they show the city is making the most of its investment. Some of these systems may even be directly unwanted by the residents. A similar pressure exists in the commercial world, where companies that have accumulated data for legitimate primary business reasons look for ways to defray the cost of that data collection. Similar pressures result in things like cooperative networks of public and private CCTV cameras.

Opportunity Leads to Abuse

The same pressures end up being behind many of the cases where the government uses surveillance abusively. It is rare that the governments of "free" nations specifically build systems to track, for instance, political dissidents. However, it is a regular occurrence that systems built for some legal purpose, whether they are wiretap systems with insufficient safeguards or databases only intended to track known criminals, are used to monitor, track, and harass dissidents and both personal and political enemies. The more global a system's reach, the more damage it can cause; furthermore, the more easily data can flow, the harder it is to ensure appropriate controls at all levels. The progress of the technology of surveillance is directly at odds with the prevention of abuse; even assuming its first purposes are reasonable.

Equality versus Aggregation

A large class of initiatives which do or have the potential to invade people's privacy are based on making data which is already stored and at least semi-public digitally available. The argument, of course, is that there is no actual change in how much privacy someone has when this happens, but this is pretty obviously functionally false. Privacy is not an absolute, binary function, but a gradation across different types of information, different amounts of data available, and the degree of difficulty of access. Changing any of those has an immediate and concrete effect on individual privacy. This, alone, assuming we consider invading the privacy of others to be basically evil, gives the lie to Google's claims to not be evil—they have done more than any other company to reduce the difficulty of getting access to information, any information, regardless of its effect on privacy.

Aggregation and data profiling make things even more complicated. When companies or agencies make an effort to collect large swaths of information from both public sources and private ones and create a profile out of it, they can learn far more than one might expect. Faint traces of information can be statistically telling, and while not always accurate, they can often be surprisingly so. Economies of scale help this process. To build such a system to investigate one person would be horrendously inefficient, but to do so to investigate a hundred million, while still expensive, is often a sound investment.

This difference in economies of scale is one of the reasons that the future where we all go along with having no privacy is fundamentally unworkable, even assuming there was no intrinsic need for privacy. It doesn't result in equality because differential resources create a power dynamic in who can know what—the results of those data profiling systems is not public record. Furthermore, organizational privacy and personal privacy are not the same thing. Corporations can be hurt monetarily, but this does not cause the same kind of psychological harm that violations of privacy can cause to people. The legal system is also taking very good care of the kind of privacy that corporations care about, thanks to extensive lobbying around patent, trademark, copyright, and trade secret laws.

[Next Slide: Surveillance Can Be Good and Bad]

Good and Bad

Now that we have a better handle on some of the things that make surveillance complicated, let's look at how some specific examples can affect people. While we've mostly looked at surveillance as negative, there are ways that it can actually be very good, depending on who you are and what your relation to the surveillance is. Because the effects of any individual surveillance system are so complex, we're going to look at specific facets of different systems. Some of the real-world implementations of systems we cover here have other, more complicated effects than what we discuss, but that's not our focus right now.

Saved By Surveillance

There are a number of different kinds of surveillance systems that purport to be lifesavers of one kind or another, and many of them actually are. One good example here is a variety of medical monitors that the medical community is developing and implanting, especially pacemakers. Just getting full-time EKG tracking data can be useful enough, even when data is only offloaded in batches, but monitors that can alert doctors and paramedics immediately in

case of a cardiac event can provide a massive advantage for survival. Other examples along similar lines include fall detection and alerting devices for the elderly and for solitary industrial workers, and various car safety systems that allow remote diagnosis and reconfiguration and can summon the police and paramedics to the vehicle in the event of a serious accident. Most of these systems are expensive. Manufacturers package them in luxury automobiles, market them to large enterprises, or they require very good insurance or private medical care. They are luxury surveillance.

Marketed To

One of the largest categories of modern surveillance is that done for marketing purposes. Is this good for the person being marketed to? That's a complicated question. To some extent, corporations do market profiling so they can manipulate the subject into spending money on something, regardless of their prior desires. There is also an element, however, of showing the subject the things they are most likely to want. To some degree this can be seen as a service, ensuring the subject is aware of things that may meet their desires, but this is incomplete as the advertising also attempts to manipulate and create that desire. This communication of desires isn't a one-way street from the company to the subject, however. The companies that employ market profilers want to know what will sell, and they make their product development decisions based on that information. As one moves further up the scale of purchases in price, the potential pool of buyers shrinks. At the high end, each individual in a marketing pool has an increasing voice about the products they want to see made. This is still a form of surveillance, but it becomes more conversational, more of a mass negotiation for those who have enough money.

Data Protection for the Rich and Famous

Even once data has been captured by some system, it is not necessarily treated the same for everyone. Many records systems literally contain specific features designed to protect certain people's privacy more than others. We know this, among other reasons, because people are regularly fired for peeking where they shouldn't—at the medical records, tax returns, and private police reports, among other things, of the rich and famous. In the case of private organizations this is easier to explain, as the rich are in a position to simply demand this treatment from everyone they do business with and receive it. Indeed, having someone on staff to simply hunt down and remove or have sealed unwanted information is probably worthwhile for them. The case of state records, however, is much more troubling. The state should be blind to such things but rarely is.

Fraud Detection

One of the common forms of non-marketing economic surveillance is fraud detection for credit and debit cards and checks or bank transfers. This is another form of surveillance working explicitly in favor of the user, as their interests are mostly aligned with those of the bank. Of course, to be eligible for this kind of surveillance, you have to have a credit card and use it. If you have the money, you can pay more to be more closely protected, which is to say more heavily watched.

Surveillance in the Common Good

There are few categories of surveillance that seem to help everyone fairly equally, like environmental surveillance and especially air and water quality testing. However, even as these kinds of sensors become more ubiquitous, their distribution is not even. Comprehensive environmental quality testing isn't cheap, and different cities can afford it to different levels. Even within cities, sensors aren't necessarily deployed evenly. Denser areas are easier to test, and among those areas, the city may monitor those with the most social (or monetary) capital more tightly. The disparity can become even larger if citizens there opt to supplement public testing with private—very common for things like home lead and asbestos testing. The same logic applies to traffic counting systems used to measure road repair needs and determine funding. Although, in theory, the transit agency should spread these systems equally across all areas in a district, this system is subject to social pressure like any other. Of course, there other cases where corporate or state interests are not aligned with those of the people, such as with corporate environmental liability or when the City of New York banned private air quality testing in the aftermath of the 2001 attacks while debates raged over the long-term effect of dust from the collapsed buildings and the city's liability for it.

Personal Protection

Another category of embedded surveillance system is the alarm, or really the entire gamut of personal protection systems. One can have one's house, car, laptop, mobile, and even person watched, for the appropriate fee. How much good these services will do you varies, among other things by how much one wants to pay. A simple alarm system for your house that calls the police after the fact at one price, a silent alarm with cameras at another, private armed response at a third, and on-site guards at a fourth. Surveillance most strongly benefits those who pay for it—in fact, it is often a race to the bottom, where the

one house on the block without an armed response team is robbed, and everyone else lives in paramilitary camps.

Surveillance as Crime Stopper

While it's only a small fraction of the surveillance in a modern city, surveillance to attempt to stop crime is the main thing people think of when they think about urban surveillance: CCTV cameras mounted on rooftops or overlooking back alleys or places people congregate, especially young people. The "ring of steel" around the City of London, and the same thing being built and rebuilt around lower Manhattan. Airport security scanners.

The efficacy of cameras, especially ones not monitored live, in actually stopping crime is in severe doubt—many studies, especially in the UK, have shown little or no effect on crime, and it is apparent that cameras have no effect on terrorism either. Police and municipalities are now turning this argument around to support the installation of smarter connected cameras (a form of so-called "soft infrastructure", or networked infrastructure), which will theoretically allow the police to intervene in real time. There is still no evidence to support the assertion that this will make a difference, of course, let alone that it would be more effective than other less invasive tactics.

Regardless of their utility, what is most interesting from our perspective is the social division cameras create, or rather reproduce and highlight. The camera creates three explicit social roles in its interaction, namely the watcher, the watched, and the bystander. The distinction between the last two categories is especially interesting. A camera installed for security is not intended to catalog all actions. Rather, it is intended to capture the subset of actions that the watcher deems interesting. Cameras thus select out a subset of the people who pass through the frame, the actual or merely stereotypical committers of those actions, the potential offenders. This is reflected in both the attention that the watcher pays to the video and the positioning of the camera.

When cameras are intended to prevent crimes against people the bystanders are the protected class, setting an even more stark opposition into effect. For the watched the cameras represent hostility and suspicion, and for the protected they represent security and freedom from trouble. For the watcher, they represent the ability to choose who is in which class and which incidents they wish to respond to and which they wish to disappear. This is especially relevant in the case of police-run cameras observing police actions—the police in the frame should in theory belong to the protected class, but by fraternity with their fellow police, they very dangerously belong to the category of the

watchers. This is demonstrated nowhere better than on occasions when police brutality should have been captured by police cameras, but the footage mysteriously disappears in case of a trial.

We should note that the protected are without exception most likely to belong to a higher socioeconomic stratum. Not only will the well-off be perceived in any given situation as less likely to be the offender, but they will spend more time in watched areas—it is no accident that some of the most heavily watched territories in the world are the homes of the global financial markets. There are even, in some cases, setups that allow the privileged to avoid personal scrutiny while still benefiting from the scrutiny of others, whether through diplomatic privilege, the now-defunct "Clear" system that allowed frequent travelers in the US to avoid much of airport security, or simply by being driven into a private garage in their office building in a limo with tinted windows, frequently driven by an armed off-duty police officer.

A further interesting note in this category is the collusion of surveillance technology in restricting speech in public-private spaces, like shopping malls— here the division is not simply between social categories but between those who wish to express ideas inimical to that particular environment and those either willing to remain silent or whose voice is in sympathy with management.

A Consistent Pattern

As we've seen, the effects of surveillance can be both good and bad. Surveillance, as an activity, is very much a direct embodiment of the power structures of the societies that perform it. The technologies that implement it render these social structures in steel, glass, and silicon, and embed it into the fabric of our cities at ever deeper and more ubiquitous levels.

We have seen a consistent pattern by which people with money, power, or social status can either act to be less watched in bad ways and more in good ways, or enjoy greater protection as a secondary effect of their standing— another reflection of the ambient social privilege through which they move, often quite unaware. This shows no signs of stopping. On the contrary, many of the smart city initiatives are attempting to push much of the beneficial intelligence of the new soft infrastructure into mobile devices, including basic functions of the modern city, like environmental monitoring—things which are good surveillance. While smart cities may be good in the end, these devices are expensive luxury gadgets and availability will always propagate down from the top, repeating the pattern again.

[Next Slide: Fixes]

Fixes

So, what can we do about all this? Both the growing surveillance state and the inequality in surveillance must be addressed, in equal measure. Hackers and makers are in a unique and interesting place here, having both the technical background and (often) the sociopolitical understanding to see the full breadth of the problem, plus an existing semi-organized community that can work together, albeit in a modern, distributed-swarm sort of way. We are both a very privileged group and an occasionally marginalized one; we can reach out to others in both directions.

We can do four things—document the problem, raise awareness of it, work to change existing projects directly, and subvert the system. In many ways, this is a catalog of tactics for provoking any kind of social change, although many of the details are obviously different. As with other issues, there are existing groups of people working in this space, and working with them will let you get more done.

Documenting the Problem

The details of the ways different groups watch us are often not public, and certainly not aggregated. Helping to find and collate this sort of information is a basic first step for any kind of further work. This can be as simple as helping to make and update maps of all the security cameras in your neighborhood. On the other hand, it can be more involved, whether that means doing background research on surveillance programs and filing Freedom of Information Act requests (or your local equivalent) or reverse-engineering surveillance systems and documenting them.

An interesting example of the latter happened recently in Seattle, when during a wrongful arrest case involving a hacker[2], the police weren't forthcoming with video evidence. After determining that the surveillance management system in question kept secure logs of all media deletions, he asked the police to back up their assertions of the video deletion with the log data. He proved that the video still existed, and it was made public to the court. Now, lawyers around the world dealing with the same (widely deployed) system and other similar ones can better make existing surveillance tapes work for their clients, removing one

---

[2] Rachner vs. Seattle Police Department

aspect of the police's power to choose when surveillance is presented as evidence.

Raising Awareness

It's not enough for just the hacker community to understand how we're being watched, any more than it's enough for that knowledge to exist in little fragments in our community instead of as a cohesive image. The rest of the world needs to understand, too. We've been telling people around us and the public about computer security and online privacy for a long time, and it's time we started speaking more frequently and more directly about offline privacy. The same issues we see online are leaking into the world more and more, and people need to understand how this hurts them and society in general—how being watched by a video camera won't keep them safe, but it will make a space more hostile to many people and encourage some kinds of harassment.

Working Directly With Existing Projects

Most new surveillance projects and the legal structures on which they're built aren't created in a total vacuum, which gives us room to act. Most of this is the usual stuff of civic activism, whether that means protest, contact with lawmakers, or raising the issue of surveillance at planning meetings. One of the important things to note is that modern surveillance straddles the line between living in the cloud at the national or Internet level and being very specifically located in the real world, and working at both levels is necessary. Furthermore, this kind of intervention doesn't have to be limited to state surveillance. Talk to people who run the businesses you patronize about the ways they use surveillance.

Subverting the System

Subverting the system is something we're very good at, and it can take a whole lot of different forms, from pranks that show the futility or the ridiculousness of the surveillance culture, to practical tools and techniques for avoiding surveillance or countering the power dynamic it creates. Examples of the first category are things like the Tor project for proxying content and fighting traffic analysis, and OTR for secure instant messaging. In the last part of this talk, I'm going to describe a new competition for projects in the second category, that help equalize that power dynamic.

[Next Slide: The Deployable Camera Competition (graphic)]

Competition

The greatest power imbalance in the sphere of surveillance exists between the private citizen and the state, and this is most exacerbated when that citizen is attempting to take action against state violence. Police and military brutality is real, it is horrifying, and it is happening all over the world. The worst excesses happen when no one is watching, and even though surveillance tools are widely deployed throughout the world, there are many countries where surveillance is only ever a tool of oppression.

In many situations, there is no question of whether or not one's image will show up on video and that if that video can be used to incriminate you, it will be. This is especially true for large political demonstrations. However, when the state decides to suppress a demonstration, the state often also suppresses the video. It is questionable at best if getting video out from places like Burma or Iran can make an immediate material difference in those situations, but at the very least it can (and has) swayed public opinion. Even more interesting to us are cases where recourse to the rule of law is possible, but there is still significant corruption—this is a large part of the world, developed and developing.

The obvious way to fix this is just to shoot your own video. Installing a citywide video system isn't a reasonable thing for an individual to do, but there's no need for a network that large. Really, what's needed is just a few cameras with a clear view in the right place at the right time. Phones that shoot video or cheap video cameras can work just fine here, but there are four problems with them. First, the person shooting the video is very easily identifiable in the moment, which is obviously undesirable. Second, hand-held lightweight cameras produce horrible video, compounded by the poor line of sight of someone standing on the ground. Third, the video is in most cases stuck on the camera until it can be uploaded—this is starting to change, but only slowly, and it generally relies on the cellular network which may be unavailable in many situations. In addition, having potentially incriminating video on your camera can be just as bad as being seen shooting it. Fourth, there's no way to collate and manage the video once it's online. There are plenty of generic ways to host the video, of course, but most of them aren't designed to allow cross-references between videos, controlled release, etc. It's also worth noting that every jurisdiction has slightly different rules of evidence that must be followed to use the video in court. All that said, phones do have the advantage of a large existing base.

In the interests of leveling the playing field a bit, I'd like to announce a competition to build deployable video cameras—a combination of off-the-shelf

components, open hardware, and free software that can be deployed in an urban environment. There will be three categories to compete in: aerial (UAV) cameras, static cameras, and software-only solutions to run on existing mobiles to provide discrete live uploading and appropriate video management.

The goal of the competition is as much to spur work on these and similar devices as it is to develop immediately fieldable solutions. The competition is open to all, but hackerspaces and similar groups are especially encouraged to enter. The only firm requirements are that everything developed for the competition must be released under open source/open hardware licenses, and that the resulting system should be legal in its chosen jurisdiction. For instance, in the US this would mean FCC-compliant radios and no ability to record audio.

[Next Slide: The Deployable Camera Competition (criteria)]

The competition will be judged on the following eight criteria:

Functionality:
- Deployability
- Operating life (ruggedness & battery life)
- Video quality
- Network integrity
- Privacy preservation & verifiability

Utility:
- Ease of use (end unit)
- Integration (software)
- Cost
- Constructability
- Documentation quality

Bonus Points:
- Real-world involvement
- Production chain

I'm not going to go into the details of scoring, etc., right now, as all of that information will be available on the competition web site that I'll give in a minute. Two notes, first. One, while cost is definitely a criteria for the competition, it's expected that the first generation of devices may be more expensive than desirable for actual deployment, but it's also assumed that prices will come down in later generations. Second, teams are strongly encouraged to get involved with groups that might use the cameras in the field during the team's design process, and both that involvement and feedback from the groups will be taken into account during judging. There will be a preliminary evaluation of progress in six months, before the CCC Congress this winter, where we'll see if

the teams need more time to complete their work. For more information, go to http://sldrc.com/projects/deployable.

[Next Slide: The Deployable Camera Competition (QR)]

Questions?

[Next Slide: Questions]

[Next Slide: See Something, Say Something]