Relay-Proof Channels Using UWB Lasers (Transcript of Discussion)

Alex Shafarenko

University of Hertfordshire

This talk is about the mechanics of security as well as protocols for security; what I am trying to do is to work out some general principles of a certain technology that is necessary for a certain type of security protocol.

This is an authentication problem with a twist. There's a prover and a verifier, they talk to each other, they use standard protocols for the prover to prove its identity to the verifier. The twist is that both the prover and the verifier have spatial coordinates, and the goal is not just to verify that the prover is who he says he is, but also that he is there in person.

The verifier verifies the identity of the prover, but also the prover has some assurance that the verifier is talking directly to the prover. So there are two physical entities involved in the protocol, whether it's explicitly by actually finding where they are, or implicitly by providing assurance that they are where they are assumed to be. This turns out to be an incredibly subtle problem. It's not a new problem, there's been research in this area, but the way that this tends to be looked at is distinctly logical and not technological, and the trouble with logic is that it is, by necessity, a model, and the model throws something away, and people have to support the model with technology finally. And that's when you discover that you threw away something that needs to be there in the model. If I manage at the end of this talk to convince you that there's something fishy about all these systems then I have achieved my main goal, but I might be able to do a little more.

So why this is important is due to the so-called relay attack, which probably you know. What happens here is that the genuine principals, Bob the verifier and Alice the prover, have two impostors, two men-in-the-middle, spliced in, that are two agents of Moriarty, Mort and Cove, and although Alice believes she is talking to Bob, in fact she is talking Cove, who is relaying every message from Alice to Mort, who is talking to Bob, and vice-versa.

There's nothing wrong with the messages sent, the integrity is not compromised, authenticity is not violated. What is violated is the binding between the message and the principal, because Alice is in the next room, and Bob faces somebody who appears to be Alice, does as Alice would, except Alice is not there, it's Mort, who is an agent of Moriarty.

But why is it so important? It is important because sometimes material goods are involved in this transaction, like I am going through the door, this is the door, this is me going, except it's not me, it's Mort who goes through this as a result of this authentication transaction. And I am actually going through a door somewhere else, into an entirely unimportant place, not knowing that I am actually letting somebody in elsewhere.

B. Christianson and J. Malcolm (Eds.): Security Protocols 2010, LNCS 7061, pp. 47–53, 2014.

Now if I could include spatial coordinates in the protocol, then that attack would not work, because I would be saying, I'm Alice located here, and Bob would check the coordinates to see that whoever appears to be talking to him is at that location. But what reliable method do we have for coordinate verification? Obviously if you give me three numbers, that's not much use for me, they have to have landmarks, as part of the scene, and I have to reference the principals with the set of these landmarks, and these landmarks should be sufficiently prominent, and sufficiently dense in the environment, not to allow Mort to mingle with Alice somewhere, and confuse Bob so that he can't distinguish one from the other.

George Danezis: Isn't the case that if you just include coordinates in the messages, you can't actually authenticate to something that is further away? So if my car knows where it is, because it has some GPS in it, and then I'm opening the door with a hand-held remote, and that door relays the message to a door that is further away, if within the encrypted shell of the messages both sides introduce coordinates ...

Reply: Absolutely, then the attempted relay would not work.

Bruce Christianson: Provided you have a foolproof unspoofable unbreakable un-denial-of-service-able trustworthy GPS within your security envelope.

Reply: But what is a reliable method to do coordinate verification? What if I think that you have this set of coordinates, in fact you have a different set of coordinates, just because my coordinates system is off, and it is off enough to allow as the impostor to share the square that's recognised as one coordinate on the coordinate system.

George Danezis: Ah, I was a bit confused about what the problem was.

Reply: Now if you use electronic means of positioning, they have to be authenticated, effectively you introduce a third party, you don't want to do that, the whole point of this type of protocol is not to rely on third parties, you want to be able to verify locally that the principals are actually bound to the cryptographic entities that represent them. You could instead focus on the communication technology and use it in order to achieve the properties that you want. So you have your options, you can make relay strictly impossible, or you can make relay prohibitively expensive, or alternatively you may just make it detectable, in which case it's ineffective. You don't have any other options if you want to look into verification.

So start from the top here, what is it to make a relay impossible? Well you can use a physical principle, for instance, the speed of light is limited, so any delay would lengthen the communication path, therefore it will increase the communication delay, and there are so-called distance bounding protocols, Kuhn etc, which are based on this principle. The trouble with this whole approach of distance bounding is that what we measure is not just the propagation delay, we measure the sum of the propagation delay and the processing delay, whatever time it takes to run the tiny little protocol. So they simplify that protocol tremendously, in fact all that most distance bounding protocols do is an XOR with an incoming message, which doesn't take much time.

But if you are operating in a very small confined volume, you will need to differentiate between communication over say ten foot distances and twenty foot distances, right, and we are talking about units of nanoseconds. And what these people don't allow for is that delays in the receiver and transmitter of that order are not to do with the distance, they're to do with electronic processing, you have to charge up a line to trigger a latch at the end of it, and that's not the propagation time, it's the time associated with the capacitance of that line, which could be significant. So when you are in a nanosecond range, the variation of delay could be significant. Now Moriarty has a lot more resources than Alice, so Moriarty could have a huge parallel receive and processing system, and could try to reduce that capacitance delay so much that it will mask the difference in propagation time.

There are protocols that try to accommodate those possibilities, there are some partial solutions, but they're never completely satisfactory. OK, another option is make relay detectable, how do you make relays detectable, this is an old problem, how do you make a copy distinguishable from the original. Now if we're talking about copies of bits, that's pretty hard isn't it, because all zeros are the same. There is technology called quantum fibre where you have individual photons propagated in a coherent quantum mechanical state, and if you try to measure that photon then you change it. So if you try to relay it, you have to first measure it, and then reproduce it, and that changes the original, which is detectable. That is in a price bracket that I don't want to consider for supermarkets, for instance. It is a possible way of going about this, but I don't think this is necessarily the cheapest way. What I would like to focus on is something that is practical, cheap, and reliable. If Moriarty has all the resources of the world, and I want to be be a hundred years future-proof, then I fail, but if I am developing this technology for today, and maybe five years from today I redevelop it quickly and cheaply, then maybe that will work. Most security assurances are cost-based anyway.

What we need for a practical approach are some universal principles, so that whenever the culture changes, and we want new technology to support this approach, then we can check it against certain principles on a checklist, and if we satisfy those principles, then we are more or less assured that this would work as advertised. So are there such principles? OK, this is the core strategy, I call it OWM, overwhelm Moriarty. There's Moriarty somewhere, he's trying to listen on all communications to relay them, we're trying to make it very hard for him to do that. How do we overwhelm Moriarty? We, the verifier, offer a massive challenge, not a few bits but a few tens of terabits, of which the real challenge that the honest receiver should receive is a tiny proportion, a tiny fraction, that only the true Alice knows where to find, because that information is part of the shared secret between the prover and the verifier.

Now how does Alice get to that portion of data? Alice would use passive selection of this significant part sent by the prover, passive selection that does not involve detection, amplification or retransmission. Passive selection does not distort anything, it is totally safe to use. For instance, when you listen to the

radio, you tune your receiver, you're choosing the wavelength that you receive, you're not distorting any other wavelengths, and you're not using any machinery for signal amplification in order to select. Moriarty does not have this luxury, Moriarty does not know where the signal is, where the actual challenge lies, because the actual challenge is small. So Moriarty would have to relay the whole huge amount of information without changing any tiny bit of it. And that would require a huge communication capacity. Communication capacity is not the same as channel capacity, channel capacity refers to the ability of channels to sustain a certain number of data bits per second. Communication capacity includes the capacity to receive, detect, communicate, and then return a message, so includes all processing both ends. And whereas you can have a significant channel capacity, communication capacity would be more limited.

So it looks like we have two requirements here. Requirement one is to prevent passive relay. Otherwise Moriarty does not even have to be there in the scene, because whatever signals are sent from the verifier to the prover may simply be received further away, and I call this a passive relay, something that acts as a relay due to the core properties of the original channel, so you have to prevent the passive relay. So examples of passive relays include mirrors or waveguides, between the verifier with a laser beaming light at the prover, so these things should be prevented.

They could be prevented passively or actively. Passively by putting things in the proverbial Faraday cage, I say proverbial because what people use this phrase for is not exactly what Faraday meant it to be, Faraday cage is an isolation unit in security models. So how do we isolate this completely so that the source of the signal and the receiver are in some closed volume, or we actually physically block the signal outside a certain volume. But there are also active ways of preventing passive relays. One active way which I find very useful is when you create a noise. Just create a lot of noise outside the area of communication, if it's optics we're talking about, bright light everywhere, except in the little tiny volume of space where you have the prover and the verifier talking to each other. So if Moriarty wants to passively eavesdrop on that he will have to contend with this huge interference. You can also focus a signal on one point in space. If you focus light on one point and put a mirror there you will disperse that light all over the space.

Requirement two, the transmitter must be able to achieve a greater dynamic range than any feasible receiver, and that sounds paradoxical because people think that a transmitter and receiver share the same kind of technology, in fact they don't, receiver technology is much more subtle and much more difficult than transmit technology. A radio 4 transmitter on the Isle of Wight here radiates one Megawatt of radio frequency power in long-wave, and you can build a circuit that emits only 1 nanowatt, and you can use them together on slightly differing frequencies, that gives you 12 orders of magnitude, 150 dB dynamic range. Now if you can build a radio receiver that can receive 150 dB dynamic range, i.e. a waveform accuracy of more than 24 bit, you will be a rich person very soon, because all I know is possible at the moment is 60 to 80 dB maybe.

So how do we do that, what sort of graded dynamic range? You can build different physical parameters, you can have verifiable limits on the receiver size, that's the easiest of all, whatever the receiver technology, if you demand that the prover produces a miniature receiver, say you've given them a slot in which to plug it, then you will limit the ability of that receiver to receive very much, because receivers have to be large if they have some reasonable dynamic range. Physical limits to linearity I've already mentioned, if you have to contend with huge signals and small signals, you will distort either the huge signals by saturation, or the small signals by limited challenge capacity. That's the more subtle one, limited challenge capacity because the sensor has a transmitter, OK, so you have a sensor, it receives signals from the verifier, but then if you are Moriarty, there's a huge demand for information that you have to relay, and the stream that you have in this relay system to the transmitter can't be narrowed by the channel if we're going to sustain that level of communication.

So these are the principles. Now I'll show one example that illustrates all these principles in one go. Suppose we decide to use an optical channel between the verifier and prover to make the communication unrelayable. So we are sending the challenge from the verifier to the prover over a beam of light. Now to make the transmitter overwhelm the receiver I use all sorts of wavelengths first of all, it is quite feasible for visible light (400 to 800 nanometres) to have a laser that is capable of radiating 400 nanometres broadband light, so from red to blue. Here I will have about 400 colour channels, each colour channel will carry easily 100 megabits per second. The information will be noise, except for certain very narrow spectral windows, the position of these windows will be governed by the shared secret that the prover and the verifier have, OK. So this is the window for the green channel and I see a signal here, this is the window for the red channel, I see no signal here, and this is the window for the blue channel, I have a signal here, so I have received one green, zero red and one blue, that's the code that I have received. If I can collect about 100 bits over a second, I'm OK. Over the same second Moriarty will receive 100 billion times 400, about 5 TeraBytes. If he has to relay that, I don't think that's easy.

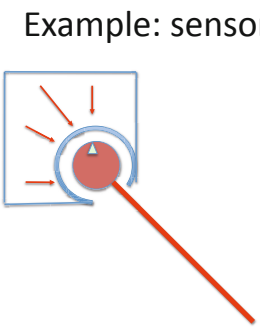

Fig. 1. The prover

Now this is the prover's device, it's a pen with a nib, the nib is one cubic millimetre in size. The verifier has a lot of space available, so the verifier has lots of lasers shining at different angles, and the prover has installed a sensor in one point here, but where that point is Mort doesn't know. So we have wavelength, time and angle. That, in my view, can give me two orders of magnitude more, so that would probably push me up to half a PetaByte per second, communicated over that short distance. If you're Mort you'll have to retransmit that. Now Mort can't do that for the following reasons. Mort can't do it over the radio, it really doesn't carry that kind of bandwidth. Mort can't do it over another beam of light that she shines from her device to Cove's fake verifier. The reason for that is that it wouldn't travel in empty space because of the huge bandwidths, it will actually disperse, you can't focus it here. Here it doesn't need to because Alice's sensor is thrust against Bob's light source, there's no propagation.

Jonathan Anderson: What if Mort's sensor in fact is a whole big bundle of fibres, and at the end they all go in different directions and so Cove can re-create ...

Reply: One cubic millimetre, and a whole bunch of fibres, which you have to coordinate, you have to make a parallel beam of light, this is focused now on this point. So you have a tiny little microwave, and behind it one fibre, yes, you can have one fibre, but I can make this volume as small as I want because this is light, so we're talking here, 10, 20 microns. Even if Mort has a perfect fibre channel, because Moriarty is a genius, right, then where does this fibre go to? Mort is standing in front of a till and she has this huge cable full of optical fibres coming out from under her petticoats, and trailing along the floor to Cove's trouser leg. It's not really practical. I work with Aston University Photonics Group, they are actually my experts in the physical side of things, and they assure me that there's no such thing as mobile fibre optic communication, they said that if that was possible a lot of people would be a lot richer. At the moment it's practically impossible.

This is my device, this is the area that we've flooded with light to avoid free space propagation. This is my little cavity with a nib. And I use every physical parameter I can. My main ones are time and bandwidth, that I can always do, I can have a lot of different wavelengths. Now the problem is, whenever I show a diagram like this I need to prove to you that I have passive filtration facilities, remember the principals have to be passive because otherwise they wouldn't reduce the effect of noise and extract the signal. Do I or do I not have them? In fact I would not have been able to do this thing only two or three years ago, the reason for it is that in astronomy there's a similar problem. You have a telescope looking at a far galaxy in the sky. The light from that galaxy is in fact almost completely masked by the scattered light of the sky, right, and this scattering only happens in very narrow spectral windows that correspond to spectral lines of the primitive elements there, oxygen, hydrogen, maybe molecules like water, etc, so there are about 150 lines in total that all the blue sky light energy sits in. So astronomers had to set their telescope in the mountains or in outer space because they couldn't deal with that, because the amount of light that they get

from the sky exceeds by 20 orders of magnitude the weak light from the stars. So then they did not have good optical filters, but now they do.

I have a publication from 2009 on my desk where they report a filter with 102 dB efficiency, and the spectral windows that they deal with are exactly 100 1nm channels, that's what I want to do. So they can now slice portions of 1nm of visible spectrum and suppress everything else by a factor of 102 dB, which is about four billion. So we can have a passive filter for this system sitting here, the filter by itself is very portable, essentially a piece of portable fibre, and there's a diffraction grating inside, and that grating is very high tech. To calculate where the elements of the grating need to go requires a supercomputer, and they're very, very impenetrable computational schemes, they have only just found a way of implementing it, many months of computing because this is a very poorly defined problem.

Now all I need is to wind this fibre into a small volume, put it behind the sensor here, the sensor is just a lens essentially, and then here I have only 1 millimetre width, one channel, 40 Petabits per second, I use a timer to slice up little pieces of it, and then put these pieces together. That's my message challenge. And I reply by radio or by, I don't know, by voice, it doesn't matter, because the channel is already unrelayed. If it's unrelayed all one way, that's enough to bind the principals to the endpoint. That's it, OK, thank you.

George Danezis: There are related approaches using spectrum, that effectively have similar properties, where you share a key ...

Reply: Indeed, this is spread spectrum essentially.

George Danezis: So you could have used it maybe over RF as well?

Reply: Well you can use it over RF, but that's not enough, because RF can be relayed over optics. I wouldn't have been saying that five years ago, but now we have software defined radio, which can sample a large chunk of radio spectrum with a digital converter. We didn't dream about that five years ago, that wasn't even remotely possible, now it's just a commodity problem. So we can't use radio any more as it's now relayable, even the microwave portion is generally relayable over optics, but optics cannot be piggybacked on anything else yet, I mean, I'm not aware of anything. Maybe X-rays.

Bruce Christianson: Preventing relaying is a matter of having an appropriate Faraday cage to enclose the protocol. Low energy frequencies are easy to stop: wire mesh will stop RF (including microwaves), and black paper will stop visible light (including IR and UV). Higher energy frequencies (such as X-rays) have a low background level, so it's easy to detect if they are being used with sufficient bandwidth to relay the entire challenge[1]: as Alex pointed out earlier, for a Platonic Faraday cage relay-detection suffices[2].

[1] Remember, Moriarty is evil and doesn't care about fatalities.

[2] We might also need some Tupperware to slow down the nano-bots.