Keep on Blockin' in the Free World
(Transcript of Discussion)

Melanie R. Rieback

Computer Systems Group
Vrije Universiteit
Amsterdam, The Netherlands

I am here today to talk about some of the security implications of Radio Frequency Identification. RFID tags are remotely-powered data carriers that resemble the theft control tags that you might find in a sweater when buying clothing from a store. Like theft-control tags, RFID tags are powered and accessed from a distance using radio waves, but RFID tags differ from theft-control tags in that they tend to have more storage space and processing power. RFID tags have security issues that have been exposed in the past few years. The heart of the problem is that RFID tags don't usually support cryptography, plus RFID application scenarios are usually not conducive to commonly performed security operations like key management; in fact, many security and privacy issues that generally exist in ubiquitous computing reappear in RFID applications specifically.

I will be discussing a new technique called Selective RFID Jamming that is an extension of a concept called RFID Blocking, that was originated by Ari Juels from RSA Security. However, Selective RFID Jamming has a number of features that makes it novel. So what is Selective RFID Jamming? It is a form of off-tag RFID access control. RFID tags can barely support on-tag security mechanisms, which is why it is desirable to take their security management and move it off the tags. Unlike RFID Blocker Tags, Selective RFID Jamming also utilizes a battery powered device. In other words, a person might carry around a small computer like a PDA or cell phone. This kind of active device has sufficient resources to harness traditional security tools like cryptography access control lists, it provides a means to introduce traditional security techniques to the realm of RFID. Another improvement is defense against a differential signal analysis attack that is faced by the RFID Blocker Tag. I'll explain that a bit later.

So first, we will delve into the differences between on-tag and off-tag access control. With on-tag access control, tags decide themselves which queries are authorized, and then respond accordingly (or don't respond). With off-tag access control, a third-party device mediates access between RFID readers and RFID tags; in a similar manner to the RFID Blocker Tag, mediators will determine which queries are authorized, it will jam unauthorized tag responses rendering them unreadable by the querying reader.

On-tag RFID solutions include a kill command, sleep/wake modes, hash locks, varying identifiers (called pseudonyms), and lightweight cryptography or authentication, using reduced AES and NTRU. Off-tag mechanisms include Faraday cages (a.k.a. tinfoil), blocker tags, and external re-encryption (where RFID

B. Christianson et al. (Eds.): Security Protocols 2005, LNCS 4631, pp. 60–68, 2007.

readers periodically re-encrypt tag data). Each of these mechanisms have their pros and cons.

One advantage of off-tag mechanisms is that they can provide access control to low-cost RFID tags. Low-cost RFID tags, including EPC (Electronic Product Code) tags[1], should be cheap enough that you can embed them in everything. This allows tags to be found in a wide variety of real-world objects like consumer items, money, passports, drivers licences, and other identification cards. For low-cost items (e.g. a can of tuna fish), RFID tags should be sufficiently cheap that the incremental cost of the tag will not eliminate the profit margin on the tagged item. In other words, low cost is essential for making RFID-based computing possible, and the "ten cent tag" is a commonly cited goal.

High cost RFID tags may or may not offer privacy protection, at the discretion of the manufacturers. One example of high-cost RFID tags are subdermal RFID tags called Verichips[2]. Trend-conscious clubgoers in Barcelona and Rotterdam get these Verichips implanted into their arms, just below the shoulder. They use these chips to pay for their drinks, and access the VIP areas and hot tub; it is supposed to be quite the thing to do. But unfortunately, Applied Digital designed Verichips to rely upon obscurity for security and privacy protection, based upon their "proprietary readers". If an RFID tag manufacturer neglects to put security on their tags, the consumer needs other RFID security/privacy options. This underscores the utility of off-tag access control.

RFID blocker tags are the best-known example of off-tag access control. However, because they are implemented on an RFID tag, they are subject to all of the limitations of RFID tags, including power limitations, storage limitations, and reliability problems (an RFID tag incorrectly orientated with the reader will not even power-up let alone enforce access control. The blocker tag works by abusing the RFID reader's singulation protocol, that is invoked when there are multiple RFID tags in the interrogation field. Like several other kinds of broadcast media, some RFID tags use a "tree-walk" singulation algorithm to resolve tag collisions. Here is an example: RFID readers may query all tags that begin with a '1'. If the reader receives a collision, it then continues by querying all tags that begin with a '10'. If that collides, it queries again looking for tags beginning with '100'. The RFID reader continues in this fashion, increasing its mask length until it resolves the collision. Simply put, the RFID blocker tag interferes with this process by simulating collisions at every step of the way – it prevents readers from figuring out which tags are present by causing a full tag id namespace traversal.

Our technique of Selective RFID jamming was inspired by the RFID Blocker Tag, but it deviates from it in a number of ways. First of all, it uses a battery powered device, so it does not face the restrictions of a power-limited RFID tag. We are testing Selective RFID Jamming on a platform called the RFID Guardian, that we are currently developing together with the Delft University of Technology and Philips. The RFID Guardian is a battery powered portable

[1] See www.epcglobalinc.org

[2] See www.verichipcorp.com

device that leverages two-way in-band RFID communications. People have previously suggested managing RFID security with a Wi-Fi or Bluetooth PDA, which is a fine idea if you assume every cash register checkout will have Wi-Fi or Bluetooth communications available. However, this will often not be realistic, and so the RFID Guardian exclusively uses RFID protocols to conduct its security operations. More specifically, the RFID Guardian acts like both an RFID reader and an RFID tag emulator (using 13.56 MHz RFID, and ISO-15693 compliant).

For convenience purposes, we believe that the RFID Guardian can be best implemented in existing available personal devices, like PDAs and cell phones. Since Nokia has already put some RFID-enabled cell phones on the market, we do not believe that our vision is far-fetched.

Since it is battery-powered, the RFID Guardian can then perform any number of standard security protocols with RFID readers, (on the behalf of resource-limited RFID tags) which may include symmetric or public key security, and which could involve entire PKIs (if deemed necessary). In other words, the RFID Guardian provides traditional security tools for a non-traditional application scenario (RFID).

Because the RFID Guardian is battery-powered, it also has more than adequate memory for storing possibly complex access control policies. In contrast, RFID Blocker Tags might only have 1K bits of space for a security policy, which severely limits its possible complexity.

Thirdly, with RFID Blocker Tag approach you are likely to end up with not one but many security policies, because each tag has its own policy. It would be a big nightmare to keep these policies updated, and when tags have been deployed you will need updates to take care of the 50% of the tags that will be lost or destroyed. If you use a centralised device that manages the security of the tags within the radio range, you know at least that the security of the tags in near proximity (perhaps 1 meter) are going to be taken care of. Non-mobile RFID Guardians can also be placed in specific locations, to create zones of protection for RFID tags in specific areas (like at home).

Tuomas Aura. So you can have your mobile phone switch to shoplifting mode? [Laughter]

Reply. Indeed, but it's not the only (or necessarily the easiest) way to shoplift.

Another disadvantage of localizing RFID in individual RFID tags is that you are increasing the complexity and price of every single tag. This means if you have a thousand tags, you have a thousand implementations of the same access control mechanism. In contrast, with Selective RFID Jamming, you indeed have to purchase the device (which is a certain financial overhead), but then it can protect thousands of very low-cost RFID tags. So for applications where cost is a show-stopping factor, like supply chain management, centralizing the access control infrastructure is the best approach.

Matt Blaze. So the model is that, I have my phone and anything within RFID range of the phone is protected?

Reply. Yes.

Matt Blaze. So it's not tied to individual tags?

Reply. No.

Matt Blaze. So presumably the model is, I would be required to turn this device off as soon as I walk into a store, because as soon as I walk up to the cash register their systems are going to stop working.

Reply. Well not necessarily. This is the reason why our mechanism is called selective RFID jamming, and that's actually exactly what I'm about to discuss in my next slide.

Matt Blaze. So it is tied to an individual device?

Reply. Yes. So this is where we get into the whole access control part. Now this is going to look really familiar. You have block and pass access control lists, sources, and targets. Just like in a network packet filter. Here the target is a list called MYTAGS, which consists of a list of RFID tags that my RFID Guardian knows belongs to me. The purpose of this list is to make the distinction between jamming queries that are directed towards your tags, as opposed to queries directed towards tags that don't belong to you.

The targeted RFID tag IDs are extracted from the incoming RFID query, and this id value (which may or may not be in the MYTAGS list) is compared to the ACL to determine whether or not it is authorized; if not, the tag response is jammed on its way back to the RFID reader.

Determining the origin of the RFID query has to be handled a different way. The current RFID protocols do not offer room for a source address, and if there one were available, it would be just as easily-spoofed as IP addresses are. So the question is: how do you determine where an RFID query is coming from? The answer is that. Unfortunately, the RFID Guardian won't usually know where a query originates from, but in a small minority of the cases the RFID Guardian might encounter a "cooperative" RFID reader. For example, you might want your RFID tags to have more lenient permissions at home, since it is a trusted environment. In this case, you might install some special backend software on your RFID reader that cooperates with the RFID Guardian (the RFID reader hardware does NOT require modification for this). You could do this for the RFID reader at home, or at your mother-in-law's house. The local supermarket could also offer this as a service. However, keep in mind that while a select number of environments will have RFID readers that actively cooperate with the RFID Guardian, the grand majority of RFID readers will not be either helpful or even aware of what is going on. For RFID readers that are unknown, there should be some default settings in the access control policy.

Here's a basic example. If the RFID Guardian detects a query that targets your RFID tags, you may want to suppress the tag responses. However, if you happen to be at home, then you may allow the queries responses. If you are at Wal-Mart, you could possibly use either pre-exchanged keys or a PKI to authenticate the RFID readers. You can grant authorized read/write permissions if necessary. And for the rest, the RFID Guardian tries to disrupt the workings of the nearby systems as little as possible.

Ben Laurie. So what you're saying is if I want to read your tags I can stand near you in Wal-Mart with my RFID reader. I let you authenticate and then I read your tags.

Reply. That question actually leads us to the issue of authenticated sessions. Once a reader authenticates itself, it needs a way of determining which RFID queries originate from it (and hence are probably allowed). Source authentication is not currently part of RFID protocols, but we're currently taking a look at doing this one layer higher.

Frank Stajano. You mention MYTAGS a lot, but there is no inherent association between me writing this policy and the tags being mine.

Reply. Well the RFID Guardian manages the association between RFID tags and the owner of the RFID Guardian by performing periodic queries to find out what tags are nearby. The RFID Guardian can correlate the results of these queries over time to determine which tags are "affiliated" with you (either knowingly or unknowingly). In such a way, the RFID Guardian maintains a dynamic MYTAGS list. It needs to be this way anyways, because you will not necessarily know what tags you will own at the time that you're writing the policy.

Frank Stajano. So you're not even telling your guardians what your tags are?

Reply. There's several ways that you can establish the "ownership" of RFID tags. At home you might have an RFID system that can backup and synchronize ownership information with the RFID Guardian every night. You can also "acquire" new RFID tags from a store, for example when you are going through an RFID automated checkout. While the RFID reader performing the queries necessary for purchasing your items, the RFID Guardian can glean the purchased tag numbers directly from the queries. Friendly "Guardian-aware" RFID readers might even send explicit ownership information as part of the purchasing procedure. (After authentication, of course.)

However, for tags that are added covertly (e.g. an attacker drops an RFID tag in your handbag), then the only way that you can discover it is by correlating periodic RFID queries. This may happen when you get home at the end of the day, when you discover that you now have one RFID tag than that same morning (that wasn't explicitly added to the ownership list). The frequency of these periodic queries represents a trade-off between privacy, accuracy, and battery power.

Frank Stajano. You are talking here about how to automatically discover which tags are yours, but my point was a slightly different one which is that I as a malicious other person can pretend that, for example, your watch is one of my tags and I can tell my Guardian to jam it. In such a way, you can perform Denial of Service on other people's RFID tags. And surely the ID of the tag is not a secret so I don't have to be the owner of the product to jam a tag?

Tuomas Aura. No. There's nothing to prevent that. Obviously anyone can carry RFID tags, anyone else can jam them.

Mike Bond. It seems that in creating these RFID guardians, these high powered portable computing platforms that are capable of impersonating any tag, that we're putting a tool into the hands of people which is going to destroy the entire binding between a tag and an object, and suddenly I'll be able to walk around with a device and pretend I'm covered in tags if I want to, or I can jam tags wholesale, it just seems that if you look ten years down the line when everyone's got this capability in their mobile phones, and they can run little programs and do whatever they like, then, you know, what do tags mean anymore?

Reply. You've got a very good point.

Mike Bond. So, you know, the observation is that, I think RFID technology does not really scale up to a proper real world environment.

Reply. Yes, well once RFID is deployed on a wider scale I think it's going to have more problems than just rogue RFID Guardians. People who oppose RFID technology might remove or even switch the RFID tags on objects. This leads to a situation where RFID systems have to deal with false positives and false negatives, just like an intrusion detection system. This indeed cheapens the entire RFID infrastructure, if you cannot believe the information that you get from it.

Mike Bond. And you can do it all with your mobile phone, you don't need to actually hack out the tags from your clothes?

Reply. The RFID Guardian, like a lot of tools, has both good uses and bad uses. However, quite frankly if I don't build this kind of an RFID tag emulator, somebody else probably will.

Mike Bond. I think its a *good* tool. [Laughter]

Matt Blaze. It seems like the threat model here is kind of ill-defined because we don't really know how the tags are going to be used or how they can be misused. One of the problems getting the most attention is RFID tags that continue to exist well beyond their necessary lifetime. If they're intended for the supply chain to the consumer, once they reach the consumer, they should die at that point. But they don't, they're still around. They can be read, and an attacker can abuse that to figure out where a person has been walking around, and for all sorts of applications that they weren't originally intended for. That seems like a small subset of the overall problem space, and a solution to that problem seems relatively straightforward, relative to this general solution that you described. However, the RFID tags in my passport must continue to exist because I want to be able to go in and out of other countries without getting strip searched or whatever they do, when you don't have one. So if you take out the supply chain RFID tags that can just be destroyed before they get to the consumer, are you left with a problem that requires this kind of generality, and this kind of centralised device. How many of these other remaining RFID tags am I going to be walking around with?

Reply. Quite a few, perhaps. Let's say you buy a box of cream spinach and it has an RFID tag. In the glorious world of ubiquitous computing, a showcase example is the RFID-enabled microwave that reads the data off the box of spinach, to determine the cooking times. There are also similar projects with RFID-enabled washing machines that warn you when you put your red sweater in with your white socks, and it can automatically determine that your load of laundry requires a low temperature, no bleach. Automating the returning of products to a department store is another proposed use for non-deactivated RFID tags. People may want these tags for their value-added features, so it's not necessarily a matter as simple as killing all RFID tags at the checkout. That is why, in the long run, we need some kind of a solution that allows us to both protect our privacy, and preserve a bit of the functionality is promised by the visionaries behind RFID technology.

Matthew Johnson. Even if you've only got an RFID passport it is still useful to be able to jam it so you can verify it is a government passport if you've got a PKI infrastructure because then you can verify that its a government passport reader that is probing it, and you can block other requests.

Audience. So this is related to the problem of context with the RFID tags, so if I'm in Wal-Mart, I want Wal-Mart to be able to read Wal-Mart tags, not any of my tags. I don't want them doing marketing research as I walk through the checkout.

Reply. Exactly. You can sometimes configure your access control lists to filter blocks of RFID tags IDs, the same way you can filter blocks of IP addresses. For example, very much resemble the Class A, Class B, Class C, and Class D IP addresses. And if some blocks of EPC codes are associated with Gillette Razors or feminine garments, then I can block all of these things. Wal-Mart would also probably have its own manufacturer code, and the product might have a distributor code, item code, and unique item code. You can filter access to RFID tags, based upon any of these criteria.

Audience. But then you have the same problem you have with IPs which is the company gets taken over, and you didn't trust the new company who owns it, and so you start using other tag IDs so you've got exactly the same problem down the line.

Reply. This is why the management of the RFID tag access control lists could be a problem, just like management of network access control lists.

George Danezis. But if you assume that the RFID tags will be used for high level things like washing socks and cooking spinach, then you have a much better ability to actually identify which tags are yours, which tags should not be revealed, which tags are sensitive. If you assume that, you also can assume from a security point of view that you have much richer information to make decisions.

Reply. As in much of ubiquitous computing, the context of the RFID Guardian owner tends to be fuzzy and ill-defined. Can the fact that your RFID tag is spinach or a sock help you make access control right decisions and is there a methodical way to represent all of this, and to gain all this context information at the right time? Maybe. However, we would much prefer to keep our idea of context as simple as possible,. However, it's potentially open-ended what context you can store and respond to. It is an entire subfield in itself, actually. But you're right that context can be useful if you know how to harness it correctly.

Matt Johnson. But on the other hand, other readers should not be able to query that same context, like the manufacturer and the model of whatever it is you're wearing.

Reply. Context in this case should help determine the amount of access that will be granted to that reader. If singulation starts occuring in the proximity of sensitive articles, than the RFID Guardian might start jamming singulation at that moment. All of this can be represented as a policy.

Audience. But that means that you'll jam anybody else around you if you have some kind of clothing?

Reply. One of the big unresolved issues here is the denial of service, because our system has the potential to harm its environment, by interfering with other RFID systems. This may even cause legal problems. But our hope is that, if you make the jamming procedure selective enough, we'll be able to prevent interfering with the RFID systems around us.

Another disadvantage of the RFID Guardian is that it is a single point of failure. If somebody steals your Guardian, you've got a problem. We think it's a good idea to use PIN codes to lock the user interface, so we can provide at least some kind of a barrier against attackers extracting the information from stolen RFID Guardians. However, once the device is physically in the attacker's hands, well you have problems anyways. Another precaution is to store as little data in your RFID Guardian as possible. For example, your RFID Guardian might only want to keep the key information for tags that were present when you left the house that morning. This minimizes the consequences of losing your Guardian.

Another problem is RFID readers with massively directional antennas. In other words, an RFID reader with a big Yagi antenna will perform its query silently to the RFID Guardian. Unfortunately, the RFID Guardian can't enforce what it can't hear. Quite frankly, I'm not sure what we can do to solve this. At least it's a comfort that the attacker is going to look pretty silly walking around with that big antenna.

Frank Stajano. What do you mean you can't even hear a query?

Reply. Since the radio waves are directional, the RFID tag would be able to hear it, but the RFID Guardian (in a different location) wouldn't be able to hear the query. And if the RFID Guardian cannot hear the original RFID query because it's so directional, it has no way of determining if it is unauthorized, and stopping it. However, for this attack to work, the attacker needs to know the exact location of the RFID tag that it is querying.

Mike Bond. If you were able to put tags to sleep, then maybe the solution is when the tags come into the range of the RFID Guardian, they can fall asleep and then the device just pretends to be the tag until they are released again.

Reply. Yes, definitely. What I've discussed today (the Selective Jamming) is only one fraction of the total functionality that you can actually implement on an RFID guardian. In July, I will be presenting a paper at the ACISP Conference[3], that also discusses topics like context-based tag activation and deactivation. This may even be preferable to doing Selective Jamming, assuming that sleep/wake functions are available on your RFID tags, because of Selective Jamming's possible legal issues. So you're right.

[3] Melanie Rieback, Bruno Crispo, and Andrew Tanenbaum. RFId Guardian: A Battery-Powered Mobile Device for RFID Privacy Management. (Australasian Conference on Information Security and Privacy - ACISP, July 2005).