# Remote Power Analysis of RFID Tags

Yossef Oren

# Abstract

We describe the first power analysis attack on passive RFID tags. Compared to standard power analysis attacks, this attack is unique in that it requires no physical contact with the device under attack. The power analysis can be carried out even if both the tag and the attacker are passive and transmit no data, making the attack very hard to detect.

As a proof of concept, we use power analysis to extract the kill passwords from Class 1 EPC tags operating in the UHF frequency range. Tags from several major vendors were successfully attacked. Our attack can be extended to HF tags and to remote fault analysis.

The main significance of our attack is not in the discovery of kill passwords but in its implications on future tag design – any cryptographic functionality built into tags needs to be designed to be resistant to power analysis, and achieving this resistance is an undertaking which has an effect both on the price and on the performance of tags.

# Acknowledgements

I am truly thankful for having the opportunity to work on such an exciting topic of research. The results presented here are the result of collaboration with many people, and would not have been possible if not for the generosity and guidance of those who helped me along the way.

First of all, I am indebted to my dear wife Michal for encouraging me to take on M.Sc studies and for enduring two years of living on scholarship. I thank her and both sets of grandparents for the countless hours they invested.

Simon Krausz found enough time in his impossible schedule to pass my idea through its first trial by fire. Oren Zarchin and his colleagues at Prof. Moti Heiblum's lab taught me how to use a spectrum analyser. Oded Smikt, the department's engineer and tinker (and master of the department's apparently endless stockrooms) eliminated the MacGuyver element from my lab setup and helped me get some results instead of a nasty electric shock. Daniel Dobkin provided me with many crucial hints regarding my research, both through personal communications and through the WJ Communications online forum. Mickey Cohen reviewed the initial version of this work and helped make it more palatable to people outside the cryptographic community. Ari Juels provided me with the idea for the attack on Generation 2 devices and helped me understand its consequences. Eran Tromer jump-started my project more than once with his many helpful suggestions and constant encouragement (I'm just glad his optimism was justified). I would also like to thank Eli Okon for helping me set up the lab, Gideon Yuval for pointing out the password guessing attack on TENEX and Prof. Oded Goldreich for the "shidduch". I thank the Feinberg Graduate School and the Department of Computer Science of the Weizmann Institute for funding my studies through the past two years. I would also like to thank my roommates for coping with the noise, wierd gear, wandering screws and perfectly safe radiation. There are many others – I thank them all for kindly sharing their knowledge, time and equipment.

I would especially like to thank Prof. Adi Shamir – for constantly challenging me with impossible goals and helping me meet them, for providing his ideas and his experience and most of all for his enthusiasm. It was truly an honour and a pleasure to work with Prof. Shamir.

Finally, I would like to dedicate this work to the people of Gush Katif. I hope that we will all be offered consolation through the building of Jerusalem.

Rehovot, August 2006, Av 5766

# Contents

# List of Figures

CHAPTER 1

# Introduction

## 1.1. The RFID Tag – the World's Simplest Computer

The continuing advancement in the field of computer engineering results in a steady stream of exciting new applications for computing systems. One such new application is the field of *pervasive computing*, which attempts to integrate computers into the environment, transforming everyday objects found in home and office environments into aware, intelligent and connected computing devices.

The RFID tag is one step toward achieving this goal. While various flavours of

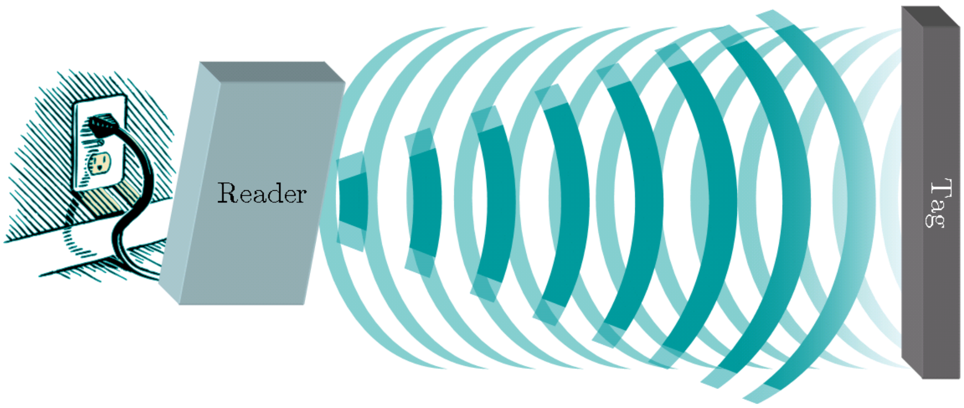


FIGURE 1.1. The general structure of an RFID system

| Property | Possible values | In this work |
|---|---|---|
| Power source | Internal battery (active) or externally supplied (passive) | Passive |
| Tag-reader link | Inductive (near field) or radiative (far field) | Far field |
| Operating frequency | Low Frequency, High Frequency, Ultra High Frequency | Ultra High Frequency (900MHz) |
| Air interface | Proprietary, ISO/IEC 14443, EPCGlobal | EPCGlobal Generations 1 and 2 |
| Tag capabilities | Read-only memory, read/write memory, microprocessor | Read/write memory |

TABLE 1. Classification of RFID tags

*medium*. The reader generates a powerful *electromagnetic field* around itself and the tag responds to this field. In passive systems, such as the one attacked in this work, placing a tag inside the reader's field (commonly referred to as *illuminating* the tag) also provides it with the power it needs to operate. As stated before, the tag is usually attached to a physical object while the reader is connected to a powerful computer or to the network. A reader may, in general, communicate with many tags simultaneously.

**1.1.2. A Taxonomy of RFID Systems.** RFID tags can be classified according to a variety of parameters, as summarized in Table 1 on page 12. Some tags are *active*, containing an internal battery to provide them with power. Most are *passive*, relying on the reader to provide them with operating power through the field it generates. Some tags rely on *inductive coupling* to link to the reader. These tags are identified by coil-shaped antennas and have a very short operating range (usually a few centimeters, depending on the frequency). Other tags use *radiative coupling* (also known as *electromagnetic coupling*) and are identified by dipole antennas, shaped more or less like a straight line. Radiatively-coupled tags typically have a longer operating range of several meters, with some types of active tags achieving ranges of 100 meters or more.

The *operating frequency* of a tag is influenced by its coupling method. Inductively-coupled tags use the *low frequency* or *high frequency* bands – roughly 30KHz to 30MHz, with wavelengths of 10 kilometers to 10 meters. The wavelengths used by inductively-coupled tags are chosen to be much higher than the distance between the tag and the reader. Radiatively-coupled tags use the *ultra high frequency* band – roughly 300MHz to 3GHz, with wavelengths of 1 meter to 10 centimeters. The operating frequency has an effect on the ability of the tag to work in RF-hostile environments such as near liquids and metals or inside the human body.

Radiatively-coupled tags offer a higher read range than inductively-coupled tags for the same reader power. This is because magnetic field strength decreases in proportion to $r^3$ while electromagnetic field strength decreases in proportion to $r^2$[**55**, p. 43]. In addition, while electromagnetic connections have no absolute bound on their range, magnetic induction simply does not work unless the magnetic field lines of the tag and coil intersect, which bounds the distance by $\frac{c}{2\pi f}$ (according to [**55**], about 3.6 meters for standard 13.56 MHz tags). On the other hand, inductively-powered passive tags enjoy a relatively abundant power supply, while radiatively-powered tags are expected to work with less than one milliwatt of power (see Subsection 3.1.1.3).

Tags and readers communicate using a standard *air interface* protocol. Some vendors have a proprietary air interface which is supported only by their own hardware. Others comply to international standards. The main standard for high frequency tags is *ISO/IEC 14443*[**13**], while the main standard for ultra-high-frequency tags is the *EPCGlobal* standard suite[**5, 19**]. The EPCGlobal air interface has gone through two generations of standards, both of which are covered in this work.

Finally, there are different *capabilities* for different tags. The simplest tag is a *1-bit* tag, which merely announces its presence when illuminated by a reader. These tags are actually quite common and are used in theft prevention scenarios. The tags covered in this work are slightly more advanced, containing several hundred bits of read/write memory and a simple protocol to control them. At the top end of the spectrum there exist tags with full-fledged *microprocessors.* These tags are in many cases standard smart cards with an added contactless interface. Due to their relatively high power requirements, contactless smart cards are usually designed for inductive coupling.

Our work focuses on passive UHF tags adhering to the EPCGlobal standard. These tags are commonly called *Electronic Product Code (EPC) tags.* These tags are passive and are radiatively coupled. They work in the UHF band (900MHz) and contain a small amount of read/write memory. EPC tags were designed as part of a global initiative to replace the common optical bar codes found on marketed goods with an RF-based electronic version. The EPC system also describes a wide-ranging technical and business-oriented infrastructure that supports this transition[**18**]. EPC tags improve on optical barcodes by offering an increased read range and more reliable data transfer, but more significantly by expanding the name space for product codes from the existing 47 bits[1] to 96 bits or even more. This larger name space can allow items to be tracked according to their individual identity, not only according to the class of product they belong to.

[1]14 decimal digits, to be exact

The economics of the EPC system were designed with the notion of tags 5 costing cents apiece. As of late 2006, vendors are beginning to reach this price point for large volume purchases. A single tag is much more expensive – we spent nearly $3 per tag for the ones used in our experiments. An EPC reader usually costs between $500 and $800, including antenna.

**1.1.3. The Case for Security in RFID – Present and Future.** One may assume that the current crop of EPC tags seems too simple to protect – after all, they are merely an upgrade to the optical bar code, which obviously has no security measures. The authors of [**46**] challenge this assumption, noting that the increased reading range of a tag, combined with the increased name space, severely compromise the privacy of individuals bearing tagged goods. This risk stems from the fact that an individual can be implicitly tracked by the specific ensemble of items he is carrying on his person. We present a more detailed survey of the security risks of RFID tags in Section 2.4. Making the data stored on RFID tags secure and trustworthy is an important concern for today's users of RFID.

In the future, the continued growth and development of the field of pervasive computing is expected to further enhance the capabilities of RFID tags as well as their popularity and their extent of deployment. Aided by the inevitable phenomenon of feature creep, future tags can be expected to contain more sensitive data and may also have the ability to make crucial decisions based on this data. As the capabilities of RFID tags approach those of smart cards, the need for cryptographically enhanced security and privacy will become even more apparent.

## 1.2. Side-Channel Cryptanalysis

Cryptanalysts try to devise methods for attacking secure systems. There are two main approaches to cryptanalysis – *mathematical cryptanalysis* and *side-channel cryptanalysis*. The difference between the two will be explained below.

Any interactive system can be defined in general by its official external *interface*. This interface specifies the *inputs and outputs* to the system. It also specifies the *behaviour* of the system when presented with different inputs. This convention holds both in the software world and in the physical world. For example, a public-key digital signature server receives a message as an input, calculates a cryptographic signature for this message using the server's private key, and finally outputs this signature; a padlocked door receives a series of dial settings as inputs and unlocks itself when the correct sequence of dial settings is punched in. Mathematical cryptanalysis attempts to attack a secure system by making use of weaknesses in the formal description of the system.

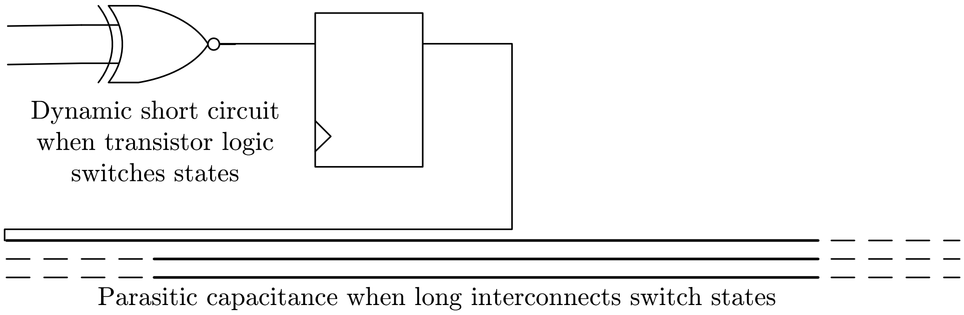


FIGURE 1.2. Sources of dynamic changes in power consumption in CMOS circuits

In addition to their official outputs, most secure systems also provide auxiliary, or *side-channel*, outputs as they work. Referring to the previous example, the signature server may take a different amount of time to sign different messages; the padlock may emit a series of clicks and whirrs as the dial is turned. Armed with knowledge of the internal workings of the device under attack – acquired by use of inside information, by reverse engineering or simply by educated guessing – cryptanalysts can now find correlations between the secret information encapsulated by the system and these side-channel outputs. Side-channel cryptanalysis focuses on finding ways of compromising secure systems based on these correlations. Turning again to our examples, the timing information of an SSL server was used by [**4**] to recover its 1024-bit private key in two hours, while [**3**] shows how many standard safes can be cracked by observing the amount of mechanical resistance the safe dial offers as the attacker spins it.

**1.2.1. Power Analysis.** One very effective method of side-channel cryptanalysis is called power analysis. Power analysis focuses on relating changes of power consumption to changes in the internal state of a cryptosystem.

Before we explain the method of operation of power analysis attacks, we will briefly review the internal structure of Complementary Metal Oxide Semiconductor (CMOS) circuits, the technology used to fabricate most low-power devices on the market today. A CMOS *integrated circuit* (IC) consists of internal state registers, some logic circuitry that makes use of these registers, and an interface that connects these functions to the outside world. Both state and logic are implemented by networks of transistors or *gate elements* (GEs). The gate elements are connected by *interconnects*, which are strips of metal running through the integrated circuit. More complex systems (beyond the scope of this work) may also use dynamic random access memory (DRAM) to store their state.

As stated in [**6**], there are three main contributors to power dissipation (or, equivalently, to current draw) in a CMOS device: *leakage current, direct-path short* *circuit current* and *loading capacitance current*. The leakage current is a constant dissipation which is a result of the manufacturing process of CMOS and is not influenced by the internal state of the tag. The other two contributors are dynamic: The direct-path short circuit current is a rush of current that occurs whenever transitions in the CMOS logic result in a temporary short circuit between the IC's power supply and the ground; The loading capacitance current results from the fact that a circuit's interconnects behave like capacitors and thus require a charge/discharge current when they change state. According to [**1**], about 15% of the dynamic power consumption of typical devices results from dynamic short circuits, while 85% is the result of parasitic loading capacitance. Both dynamic sources of power consumption manifest themselves only when state bits in the IC flip their values.

These dynamic properties mean that, in general, a active CMOS device consumes more power than an idle device, allowing an attacker to learn exactly how long certain operations take and raising the possibility of timing-based side channel attacks. With sufficiently sensitive equipment, an attacker can also estimate how many individual bits flip at every point of time, allowing even more powerful attacks. To mount a power analysis attack, the attacker places a sensitive current probe between the device and its power supply, then measures the change in power consumption over time. This attack is especially suited to smart cards, a popular class of secure devices which is designed to be tamper resistant but still has its power supplied by an outside party.

A more detailed survey of power analysis attacks and the work done to prevent them is available in [**1**] and in [**28**].

**1.2.2. Protection from Side-Channel Attacks.** As mentioned previously, side-channel attacks gain their strength from the correlations between a system's secret data and its side-channel outputs. Side-channel countermeasures are generally designed to minimize this correlation. There are several approaches toward this goal, including masking the secret with some random data, attenuating the side channel by some form of shielding and using specially designed components with less side-channel leakage. We survey several countermeasures suitable against power analysis attacks in Section 5.2.

## 1.3. Our Contribution

In this work, we show how power analysis, a form of attack which typically requires physical access to the device under attack, becomes a remote attack when we apply it to passive RFID devices. We show how to construct a lab setup that can perform power analysis over a distance and demonstrate the effectiveness of the attack by recovering the kill password of an EPC tag. We also present an existing power analysis countermeasure which fits into the tag manufacturing process and effectively protects against our attack.

## 1.4. Structure of this Document

The structure of the remainder of this work is as follows: in Chapter 2, we review some related work in the fields of power analysis attacks and attacks on RFID tags. In Chapter 3 we review the theoretical aspects of our attack, including a short survey of the EPC standard. In Chapter 4 we discuss the practical issues related to our attack and present our results. Finally, in Chapter 5 we discuss the implications of this attack, review several types of countermeasures to protect against it and present several new ideas for extending our results.

CHAPTER 2

# Previous Work

This section will discuss previous work related to our results.

## 2.1. Password Guessing Attacks

The results presented in this work make use of power analysis to try and guess a password. Using the fact that we can guess the password one bit at a time, the time required to search the password space is reduced from exponential to linear, and thus it is easy to attack arbitrarily long passwords. The first documented use of this attack as a way of guessing passwords in a computing environment was in the TENEX operating system, circa 1970 [**32**, Section 2.1].

## 2.2. Power Analysis of Smart Cards and Other Cryptosystems

The capabilities of power analytic attacks were first demonstrated in an academic setting in [**28**]. Power-analysis attacks lend themselves naturally to smart cards, since the internal state of smart cards is protected from outside inspection by various tamper-proofing methods while their power supply is run from an external line and, as such, can be delicately monitored by an attacker without tripping the tamper protection. There have been many follow-up works to [**28**], exploring both the capabilities of power analysis and the cost involved in preventing them. Power analysis has been used to extract the keys from smart cards using secret key ciphers such as DES [**15**] and AES [**14**], as well as devices using public key cryptosystems such as RSA [**37**]. In [**46**] the authors suggested that RFID tags may be vulnerable to power analysis and fault attacks. In [**42**] the authors presented an RF front-end for an inductively coupled contactless smart card, remarking that "Contactless smartcards are especially susceptible to power analysis because the power signature of a transaction is actually broadcast in the air".

## 2.3. Remote EM-based Attacks

Several other research works present attacks mounted by a remote attacker armed with a directional antenna. Two interesting results are discussed below.

**2.3.1. Screen Sniffing.** Cathode ray tube (CRT) displays operate by scanning a single electron beam of variable intensity over the entire screen in a predetermined pattern of lines. The signal controlling this intensity is called the luminance signal. This luminance signal is subjected to a very high level of amplification before it is used to manipulate the strength of the electron beam. In [**54**] Wim van Eck demonstrated that this signal leaks out of the display unit in the form of electromagnetic radiation. This signal can be then intercepted and used to reconstruct the image shown on the display. In a demonstration for the BBC's "Tomorrow's World" show in 1985, van Eck used a van-mounted VHF antenna to intercept the screens of computers in the Scotland Yard building several tens of meters away from the attacker.

In [**31**], this attack was further extended to modern LCD displays. The attacks of [**31**] operates in frequency ranges and signal envelopes similar to the ones used in our attack, and thus his estimates on the usable attack range were useful to us as well.

**2.3.2. Smart card Emanations.** Since any conducting wire inside a computing device can be considered as a transmitting antenna, several researchers have attempted to monitor the electromagnetic emanations of the internal circuits of a smart card and thus deduce its internal state without tripping the tamper-protection safeguards. These attacks are usually carried out by placing a short coil antenna directly above the smart card, exactly over some location of interest. In [**36**], the authors demonstrated a different remote attack on cryptographic smart cards using electromagnetic emanations. In this attack the receive antenna was not located directly above a specific spot on the smart card, but rather at a distance of 2 meters from the card. The attack was performed in an *anechoic chamber* - a special environment designed to minimize the interference caused by external radio frequency sources and by multipath propagation of the intercepted signal. Our attack is different than the attack in [**36**] in that it does not monitor the electromagnetic emanations of the circuits inside the device under attack, but rather presents an indirect way of monitoring its actual power consumption. Our method of attack apparently has better range and more resistance to noise than the attack in [**36**], since we were able to mount it in an electromagnetically noisy lab environment and without explicity filtering against multipath effects.

There are also different countermeasures to be employed against these two attacks. On one hand, our attack can be prevented by making the device resistant to power analysis, as we will describe in Section 5.2, while the authors of [**36**] argue that some power analysis countermeasures will not help against their attack (since even after masking the power consumption of the device as a whole, there may still be relevant data in the power consumption of specific parts of the device). On the other hand, surrounding the chip at the heart of the tag with EM shielding (without, of course, shielding the antenna) will protect against standard EM-based attacks but will not protect against our attack.

## 2.4. Attacks on RFID Tags

**2.4.1. The RFID Threat Model.** The current lines of attack against RFID tags are derived from the capabilities of today's tags. At present, the most common functionality of a tag is to provide a static payload (identifying the item to which it is associated) to the reader. Either reader or tag may be required to authenticate themselves, and the signal exchanged between reader and tag may be encrypted. In this scenario, an RFID adversary may desire either to prevent the tag and the reader from communicating, to masquerade as the tag or as the reader, or to bypass the channel encryption. The general threat model is surveyed in detail in [**46**]. One specific threat which we address in this work is the case of an adversary disabling or rewriting tags at will.

Generally speaking, the technological situation is bringing RFID tags ever closer to having all properties of a standard reconfigurable computer, both in terms of the quantity and quality of the information it may store and in terms of its computing power. The threat model will obviously evolve as the capabilities of RFID tags grow, finally converging into the standard threat model for a personal computer.

**2.4.2. Current Attacks on RFID Tags.** This section presents a short survey of some attacks against the current line of passive RFID tags. Since this work focuses on physical layer attacks, we will present three works in which tags were attacked at the physical layer – **zapping, jamming** and **skimming**.

**Zapping attacks** attempt to incapacitate a single tag, rendering it unable to communicate with any reader. An attacker with physical access to the tag can achieve this goal by cutting the antenna apart using a pair of scissors – the stubs of the antenna which remain connected to the tag will provide it with much less power than a full-sized antenna, dramatically reducing its read range[**23**]. A more advanced attacker can try and create an electromagnetic pulse (EMP) which will overwhelm the tag's reciever circuit and render it unusable. In [**39**] the authors show how to create an "EMP gun" using a disposable camera's flash circuitry and demonstrate its use in disabling ISO/IEC 14443 RFID tags. The authors also note that their EMP gun has the same destructive effect on personal computers, portable music players and pacemakers.

**Jamming attacks** are active attacks which attempt to disrupt the tag-reader communications in a certain location. In [**22**] the authors present a "blocker tag" – a special tag designed to prevent any other tags being read in its vicinity. The blocker tag is designed to participate in the EPC singulation protocol (see [**5**, subsection 4.2.2]) and to answer positively to all reader inventory queries, in fact creating the false impression that all $2^{96}$ possible tags are present in the reader's vicinity. Since the blocker tag understands the EPC protocol, it can be designed to have a more benevolent behaviour, perhaps blocking only a certain subset of the ID space (for example only medicines) or disabling itself when being presented with a properly authenticated reader. While this attack works in the application layer, another attack, presented in [**17**], achieves similar functionality at the physical layer. The attack of [**17**] prevents tags from even hearing the reader by creating a competing signal in the same frequency range as the reader. This attack is made easier by the fact that while the RFID reader performs frequency hopping to help it share the air with other devices, the RFID tag listens undiscriminately to all signals in its range.

**Skimming attacks** allow the attacker to impersonate a tag. These attacks are commonly launched against tags designed to offer proximity-based priveleges, such as allowing a car to start only when the key is present, or allowing restricted access to a facility only to persons injected with a subdermal RFID tag[**8**]. For the simplest tags without a challenge-response protocol, it suffices to record the signal backscattered from the tag under attack as it is queried by a reader and replay on demand [**56**]. This can be performed even without understanding the bit structure of the tag-reader channel, as long as the recording is done at sufficient resolution. More robust tags with challenge-response protocols can be subjected to **relay attacks**. This sort of attack is carried out by two colluding attackers, as shown in [**26**]. One attacker (the "ghost") is located next to the reader and another (the "leech") is next to the tag. The tag-reader data exchange is relayed between the two colluding attackers by an external channel.

CHAPTER 3

# Theoretical Background

This section will describe the theoretical aspects of our work. It will describe the family of tags we attack, and how we intended to attack them.

## 3.1. The EPC Standard Family

The RFID system we considered was the EPC system, used in tags attached to items of merchandise in retail and other supply-chain scenarios. The standards body governing the concepts of this system is EPCGlobal, a not-for-profit organization formed around MIT's Auto-ID Center in 2003[**18**].

The EPC standards define the capabilities of the tag and reader and determine how the two should communicate. The communications protocol consists of two layers:

- The **physical layer** (also called the **air interface**) defines the radio characteristics of the protocol, including the way bits and symbols are represented and the allowed frequencies and power levels used by the reader.
- The **application layer** defines the set of commands the reader and tag should exchange. It also defines how the tag should respond to the commands it receives.

The EPCGlobal architecture defines in [**45**] 6 classes of tags, ranging from Class 0 to Class 5. Class 0 tags are the weakest – they are read-only devices capable only of emitting a certain fixed ID they were assigned when they were manufactured. Class 5 tags are the most powerful, being for all practical issues full-fledged portable computers with support for the EPC air interface. The most common tag class, and the one discussed here, is the Class 1 tag. It defines a passively-powered tag with no computational resources other than a small amount of read/write memory and the logic required to access this memory.

The Class 1 EPC protocol went through two generations. The first generation was defined by MIT's Auto-ID center in [**5**]. While it was never formally ratified, the Generation 1 protocol was accepted by the industry and was deployed in hundreds of millions by late 2005 [**9**]. The parts of the protocol which were vague or incompletely specified were implemented arbitrarily by manufacturers, leading to

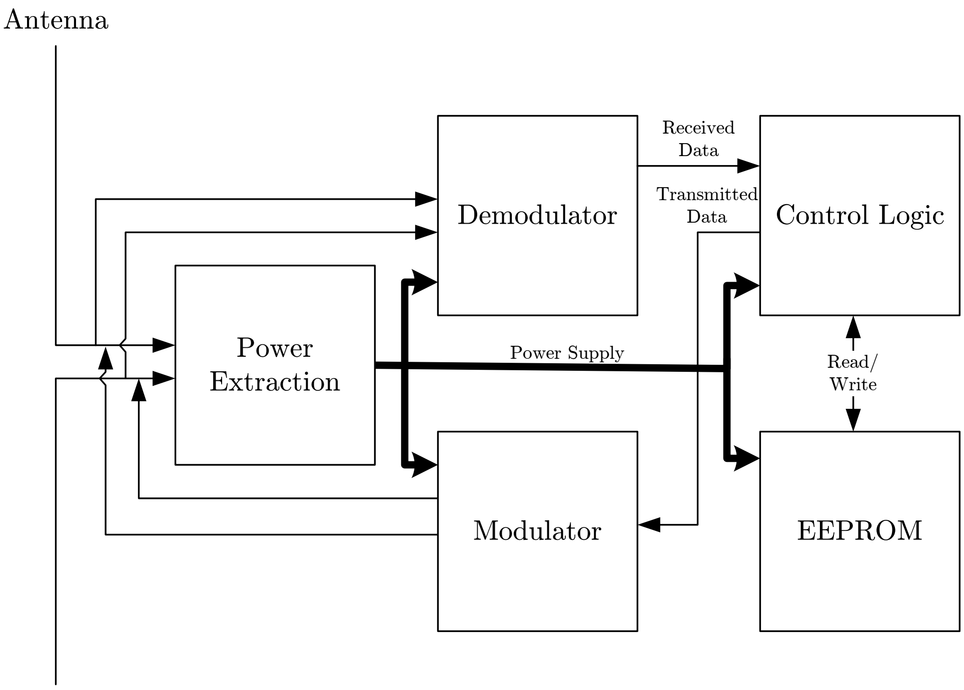


FIGURE 3.1. The general structure of a tag

some incompatibility problems and to the emergence of a de-facto standard based on imitation of the leading vendors' implementations (see for example [**50**]).

The current recommended standard is called Class 1, Generation 2 (C1G2), and is defined in [**19**]. While both generations of the Class 1 protocol share a common air interface, their application layer is quite different[1].

We will now discuss the parts of both protocol layers which are relevant to our attack.

**3.1.1. The Physical Layer.** The aspects of the physical layer we will discuss are the frequencies and power level used, how power is supplied to the tag, and how tag and reader communicate.

3.1.1.1. *The General Structure of a Tag.* Figure 3.1 on page 24 shows the general structure of a passive UHF tag, as described in [**24**].

The tag's *antenna* is its link to the outside world. The signal incident on the antenna contains both the power required to operate the tag and commands from the reader. The *power extraction circuit* converts the low-amplitude AC input to a 5V DC voltage which is provided to all other modules. The *demodulator,* or decoder,

[1]The common air interface was chosen to allow Generation 1 readers to support Generation 2 tags without changing the hardware. The MPR6000 reader we used was actually supplied with Generation 1 firmware and was upgraded to Generation 2 using software from WJ's web site.

extracts the data from the incoming signal and forwards it to the *control logic.* The control logic module, which is in charge of implementing the tag's command set, is generally constructed of a *finite state machine* that has access to *electrically erasable and programmable memory (EEPROM).* The EEPROM is used by the tag to store long-term, non-volatile data such as the tag's ID. The *modulator* is used to convey responses from the tag back to the reader. This is done by dynamicaly changing the impedance of the tag and thus the magnitude or the phase of the signal it reflects back to the reader's receive antenna (see Subsection 3.1.1.5). Recalling our cryptanalytic intentions, it is apparent that while all six modules mentioned above have an effect on the transient power consumption of a tag, the secret information we are after is contained only in the power consumption of the control logic, the EEPROM, and the bus connecting the two.

3.1.1.2. *Frequencies and Power Levels.* The EPC system operates in the industrial-scientific-medical (ISM) frequency band. The ISM band is unique in the sense that low-powered transmitters can be used in it without being individually registered and licenced. The exact choice of frequency for the ISM band varies between countries, but it is generally between 860 and 960MHz.

In a passive RFID system, the electromagnetic field generated by the reader is used to power the tags it interrogates. The amount of power available for the tags decays quadratically as the distance between the tag and the reader grows, and all tags require a certain minimal amount of power to operate. Thus, for any fixed tag design, the maximum transmit power a reader can use immediately determines its maximum range.

The different national standards institutes define the maximum power allowed for an RFID reader. As shown in [**12**, subsection 4.2.5.2], this maximum power limit is commonly measured using the e*ffective isotropic radiated power* (EIRP) rating. EIRP is defined in [**11**] as "the power that would have been radiated by an isotropic antenna with the same power density as the real antenna in the direction of maximum gain".

Practical antennas do not have a uniform power distribution in space, but rather follow some *radiation pattern.* For example, a dipole antenna generates a strong field on the plane normal to its axis and incident on its center, while it radiates no power along the axis of the dipole itself. The EIRP measurement method means that users cannot try to increase the range of their readers while staying within the regulated power constraints by using a high-gain, highly directional antenna[2]. Figure 3.2 on page 26 shows the radiation patterns of several types of antennas,

[2]A passive tag's radiated field is always much weaker than the reader's field. However, the tag's dipole antenna is nearly isotropic while the reader's antenna is usually directional. We found this creates points in space where the tag's signal is much stronger than the reader's signal, making our attack easier to carry out (see Subsection 4.2.2).

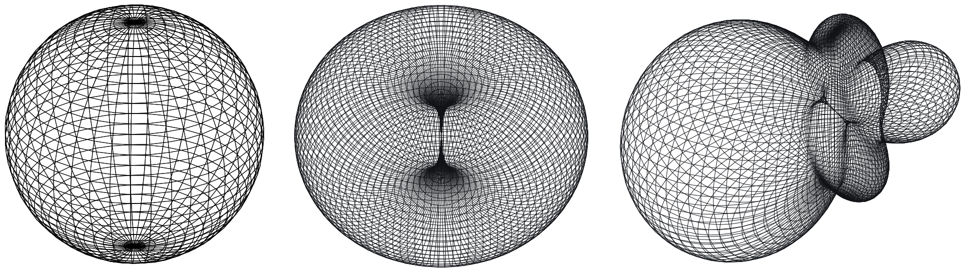


FIGURE 3.2. The radiation patterns of different types of antennas

plotted using EZNEC [**33**]. The contour line in the figure represents areas in space with similar power density.

The exact definition of the allowed frequencies and power levels for EPC tags varies between countries. In the USA the Federal Communications Commission (FCC) defines a frequency range of 902–928MHz and a maximum EIRP of 4W. In most of Europe the European Telecommunications Standards Institute (ETSI) defines a frequency range of 865.6–867.6MHz and a maximum EIRP of 3.2W[**2**]. As of late 2006, RFID regulation in Israel was still a work in progress – the suppliers we contacted stocked both FCC and ETSI-compliant tags.

EPC tags transmit their responses to the reader by modeulated reflection of the reader's signal. Due to this property, the operating frequency of a tag-reader system is completely determined by the reader. Thus, a single tag can be made to be both FCC- and ETSI-compliant.[3] In general, a tag operated outside its recommended frequency range will still work, but its usable read range will be dramatically lower.

3.1.1.3. *Power Supply to Passive Tags.* The carrier wave generated by the reader is used both to provide the tag with power and to send it data. While this subsection will focus on power extraction, the next subsection will deal with the data payload carried by the reader's signal.

According to Maxwell's equations, placing a conducting antenna in a variable electromagnetic field causes a current to flow through the antenna. Conversely, passing an alternating current through a conducting antenna generates an electromagnetic field around it. In standard applications of radio-frequency data transfer, such as FM radio, the signal induced on the antenna by the external electromagnetic

[3]Tag vendors may have an incentive to create tags which are not usable throughout the world. Limiting a tag to a specific market and frequency range allows vendors to tune the tag to a narrow frequency band. This lets the tag operate at a higher "Q factor", allowing a more efficient energy transfer into the tag circuit and thus improving its usable range. According to [**11**], higher-Q tags can also use shorter antennas than low-Q tags, improving their form factor.

field provides the data signal (for example, the music on the radio), while power is provided by another source. Passive tags, on the other hand, have the ability to extract their operating power from the current induced on their antennas.

The reader powers the tag by generating a sine wave with the approximate frequency of 900 MHz. If the tag antenna is properly *matched* to this signal, a standing wave will develop on the tag's antenna. This standing wave is an alternating voltage differential which causes current to flow through the tag circuitry.

As stated in Subsection 3.1.1.2, real-world antennas are not isotropic, meaning that their orientation in space has an effect on the power of the signal they receive. Tag makers are usually interested in maximizing the tag's usability, regardless of its orientation, so tags usually use *half-wave dipole* antennas, which are relatively omnidirectional. These antennas are constructed from a straight segment of conducting material with a length equal to half the wavelength of the signal (in our case approximately 15cm). Vendors usually *meander* the antenna in a somewhat curved pattern to minimize the effect of the dipole's "blind spots" directly along the antenna's axis.

To maximize the transfer of energy from the antenna to the IC, the tag's internal circuits are designed to expose a *purely resistive* load toward the antenna. This is done by cancelling out any capacitive or inductive loads generated by the tag's internal circuitry by an additional network of coils or capacitors.

We can now use the standard microwave engineering equations to discover the power available to the tag.

To see the total power available to the tag circuit's power extraction module, we can measure the proportion of the sphere of energy surrounding the reader which is captured by the tag antenna's effective area. This gives us the formula

$$P_{\text{Tag}} = P_{\text{Reader}} \cdot \frac{A_{\text{Tag}}}{4\pi r^2} \tag{1}$$

The effective area is a measure which depends on the antenna's geometry and the incident wavelength. For a tag with a half-wave dipole the effective area is defined in [**12**] as $\frac{\lambda^2}{2\pi}$. Assuming a 900MHz signal, for which $\lambda = \frac{c}{9 \cdot 10^6} = 33.31\text{cm}$, we arrive at $A_{\text{Tag}} \approx 88.3\text{cm}^2$. Assuming a tag located 3 meters away from an ETSI-compliant reader with an isotropic antenna transmitting at 3.2W, the total amount of power available to the tag is approximately $250\mu W$.

This power is presented to the tag as a standing wave $U_0$ generated on the tag's antenna. According to [**12**, p. 125], the amplitude of this standing wave is defined as

$$U_0 = \ell_0 \cdot E \tag{2}$$

where $\ell_0$ is the *effective length* of the tag's antenna and $E$ is the electric field strength of the incoming wave. The effective length is a function of the antenna's effective surface area $A_{\text{Tag}}$, of the wireless medium's *characteristic wave impedance* $Z_{\text{F}}$ and of the antenna's *radiation resistance* $R_{\text{R}}$. The strength of the electric field is a function of the reader's effective isotropic transmitted power and of its distance from the tag:

$$\ell_0 = 2\sqrt{\frac{A_{\text{Tag}} \cdot R_{\text{Rad}}}{\pi Z_{\text{F}}}},\ E = \sqrt{\frac{P_{\text{Reader}} \cdot Z_{\text{F}}}{4\pi r^2}}$$

Combining the two equations we obtain

$$U_0 = \sqrt{\frac{A_{\text{Tag}} \cdot R_{\text{Rad}} \cdot P_{\text{Reader}}}{\pi^2 r^2}}$$

For a standard half-wave dipole antenna $R_{\text{Rad}} \approx 73\Omega$[**12**]. Assuming that $A_{\text{Tag}} \approx 88.3\text{cm}^2$, we arrive at the final approximation of $0.258\frac{\sqrt{P_{\text{Reader}}}}{r}$ Volts. Returnimg to the previous example, the tag described there can expect a voltage differential of about 0.15 Volts.

To reduce costs of tags and enable high-volume manufacturing, tag circuitry is usually constructed from older-generation CMOS silicon, which requires a 5V DC operating voltage. To convert the low-voltage AC signal received on the antenna to a reasonable DC source, the tag makes use of a circuit called a *charge pump*. As shown in [**24**], this circuit rectifies and amplifies the standing wave into a higher-amplitude DC voltage and finally uses it to charge a capacitor that powers the rest of the circuit. The power extraction process is inherently lossy, meaning that typically only 15% of the power incident on the tag's antenna will actually be available to power the tag's internal circuitry (see [**24**, subsection V]).

3.1.1.4. *Data Transfer from Reader to Tag.* The EPC protocol is based on a reader-talks-first methodology, in which the tag may only transmit data as a direct response to reader queries. The communication protocol is *half-duplex*: at any stage in time only one of the two devices may transmit data.

In both generations of the EPC protocol the reader sends the tag commands in the form of packets, each of which consists of a sequence of symbols. Each symbol consists of a series of binary bits. The symbols have variable lengths, generally consisting of 2 to 8 bits. The different lengths for each symbol form a basic Huffman encoding that allows common commands to be sent and handled more efficiently. Each bit is sent by varying the strength of the reader's field between two levels over time, using a scheme called *pulse amplitude modulation* (PAM) or *amplitude shift keying* (ASK) [4]. The changes between the high and low levels are not abrupt in practice. Instead, the output signal is low-pass filtered, resulting in a gradual

[4]The Generation 2 standard offers two additional modulation methods, SSB-ASK and PR-ASK, which are not discussed here

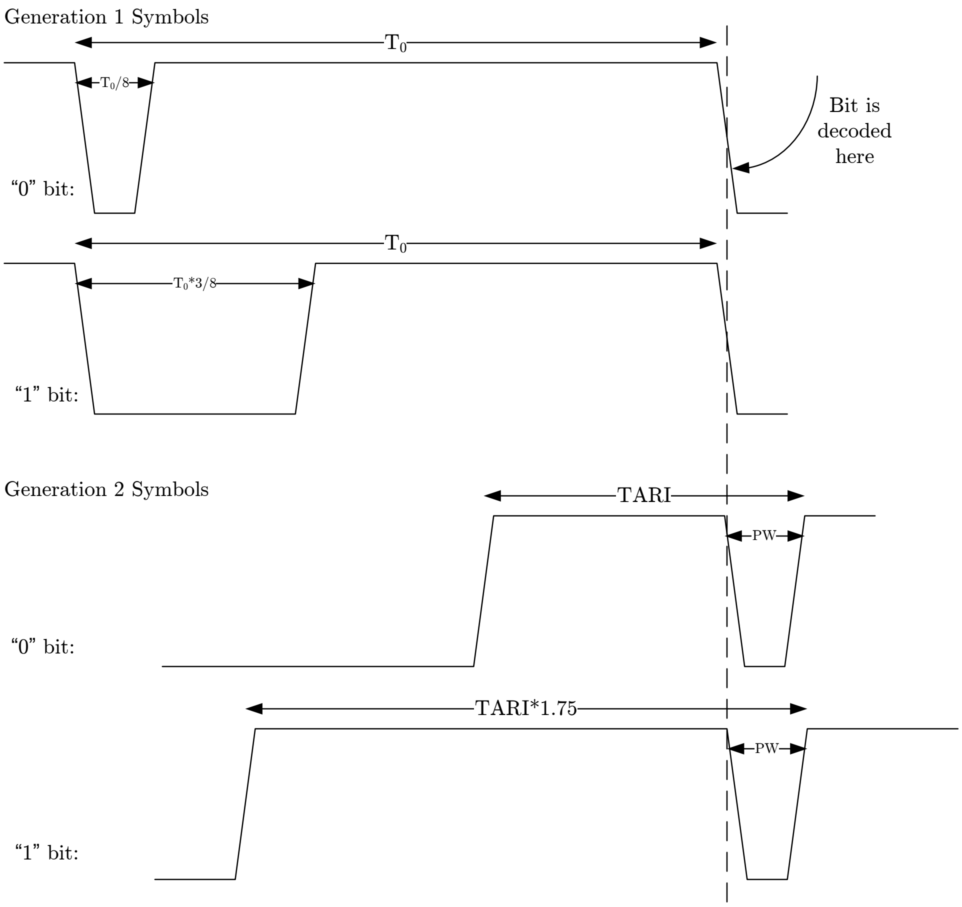


FIGURE 3.3. Bit shapes of Generation 1 and Generation 2 symbols

change in power levels. This ensures a narrow reader bandwidth, which is required by regulations.

Both generations of the EPC protocol use similar bit shapes, as shown in Figure 3.3 on page 29 (see [**5**, subsection 5.3] and [**19**, subsection 6.3.1.2.3]):

The values of the parameters $TARI$, $T_0$ and $PW$ used in the diagram are defined by the protocol and can very between different regulatory domains. Note that in both protocols the value of the bit can be measured by calculating the distance between a rising edge and the following falling edge. Thus, actions that must be carried out once per bit are probably performed as soon as the falling edge is detected by the tag.

Because the tag extracts its operating power from the reader signal, there is a correlation between the symbols transmitted by the reader and the amount of power available for the tag. Parts of the signal in which the data modulation requires a lower amplitude supply the tag with less power, so both protocols specify symbols which have a relatively high duty cycle. In fact, Generation 2 high-powered operations that write to the EEPROM (tag write and tag kill) require that the reader provide the tag with a long stretch of unmodulated carrier wave (CW) while the command executes (see [**19**, subsections 6.3.2.10.3.3 and 6.3.2.10.3.4]). This CW sequence offers the highest amount of possible power to the tag, ensuring these two power-hungry commands will execute completely.

The protocol allows some variability in two parameters of the shape of individual bits – the symbol length and the depth of modulation. A shorter symbol length allows a faster symbol rate but increases the bandwidth of the reader signal. A deeper modulation increases the difference between high and low power levels, increasing the reliability of the tag's demodulator and thus allowing an increased range, but again at the price of a higher-bandwidth signal. In general the FCC allows shorter symbols with deeper modulation to maximize speed and range, while ETSI uses longer and shallower symbols to minimize the effect of the EPC system on other ISM band users (see [**5**, subsection 5.3.1]).

As stated in the previous section, the tag is powered by the reader's signal. Since this form of modulation causes the amplitude of the reader's signal to vary in time, the power extraction module of the tag uses an intermediate power storage in the form of a large capacitor. Generally speaking, the capacitor is charged during periods in which the reader's signal is at its high level, providing power to the rest of the tag's circuit while the reader's signal is relatively weak.

3.1.1.5. *Data Transfer from Tag to Reader.* Upon receiving a command from the reader, the tag can send a response to the reader using the *backscatter modulation* principle. As noted in Subsection 3.1.1.3, any current flowing through the tag's antenna immediately causes an electromagnetic wave to be generated around it. By controlling the current flowing through the antenna, the reflected field can be modulated and thus used to convey meaningful information to the reader.

Assuming a fixed input voltage on the tag's antenna terminals, the tag can control the current flowing through the antenna by changing the impedance exposed to the antenna by the tag's internal circuitry. Changing the *resistance* of the tag's internal circuit has an effect on the *amplitude* of the reflected field. Changing the *reactance* of the circuit has an effect on the *phase* of the reflection. As demonstrated in [**24**, subsection 3.B], the tag can switch rapidly between two impedances and thus modulate its response to the reader. To switch between the two impedance values, the tag uses transistors to connect or diconnect an additional subcircuit consisting of a capacitor (for changing phase) or a resistor (For changing amplitude) in parallel to the other tag functions.

The use of phase modulation is recommended by the Generation 2 protocol designers since it allows the tag's control logic to receive the same amount of power

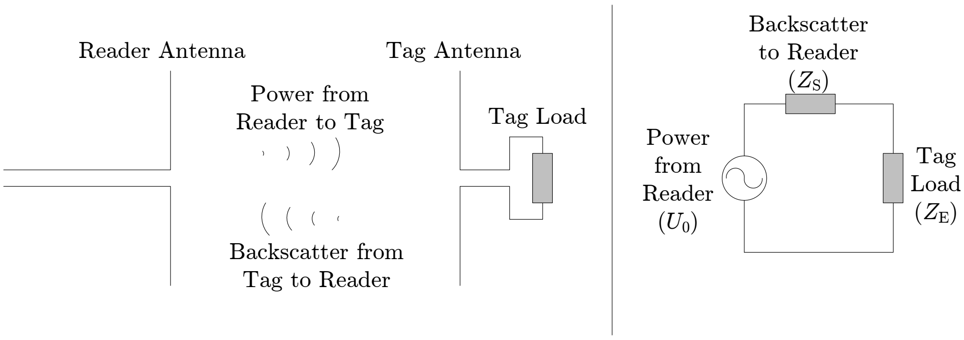


FIGURE 3.4. The reader-tag channel and its equivalent circuit

regardless of whether the tag is transmitting a '1' bit or a '0' bit (see [**24**, subsection IV.B]). In practice, this form of phase modulation can also be detected by an AM receiver at the reader because the phase differences cause interference (either destructive or constructive) with the reader's transmitted signal, resulting in different amplitudes for different phase differences. Because the intereference can be either destructive or constructive depending on the relative positions of the tag and the reader, the tag cannot tell beforehand whether reflecting a stronger signal will make the reader's received signal weaker or stronger. To get aronud this obstacle, the tag modulates the response data into the timing of the changes between high and low states, not into the high or low values themselves. This is actually a basic form of frequency modulation, commonly known as *frequency shift keying* (or FSK). From this point we will focus on changes of amplitude due to a varying resistive load, since this effect is more relevant to the attack we present.

The exact relation between the tag's internal impedance and the strength of its reflected field can be derived by observing the *equivalent circuit* of the tag-reader system. As shown in Figure 3.1.1.5 on page 31, the system can be viewed as an *alternating voltage source* $U_0$ representing the electromagnetic field falling across the dipole antenna, a *complex impedance* $Z_E$ representing the tag's effective internal loading and another *complex impedance* $Z_S$ representing the signal transmitted through the antenna and into the air. Assuming the case of a *matched circuit*, in which the antenna is properly tuned to the frequency of the reader's signal, the impedances become real, Ohmic loads, which are marked $R_E$ and $R_S$. In the equivalent circuit representation, power falling on $R_E$ is used to power the tag, while power falling on $R_S$ is actually radiated from the antenna. While $R_S$ is generally a constant depending on factors such as the shape of the antenna and the wavelength of the incident signal, $R_E$ is a time-varying quantity depending on the tag IC's internal state.

The relation between $P_{\mathrm{E}}$ and $P_{\mathrm{S}}$ (the power consumption of $R_{\mathrm{E}}$ and $R_{\mathrm{S}}$, respectively) is calculated using the standard voltage divider equation:

$$\begin{aligned} P_{\mathrm{S}}(t) &= I(t)^2 R_{\mathrm{S}} = \left(\frac{U_0}{R_{\mathrm{S}} + R_{\mathrm{E}}(t)}\right)^2 \cdot R_{\mathrm{S}} \\ P_{\mathrm{E}}(t) &= I(t)^2 R_{\mathrm{E}}(t) = \left(\frac{U_0}{R_{\mathrm{S}} + R_{\mathrm{E}}(t)}\right)^2 \cdot R_{\mathrm{E}}(t) \end{aligned} \tag{3}$$

$U_0$ is determined by the strength and wavelength of the reader's field and by the properties of the tag's antenna and is independent of the tag's power consumption.

A plot of the relation between the internal resistance $R_{\mathrm{E}}$, the absorbed power $P_{\mathrm{E}}$ and the reflected power $P_{\mathrm{S}}$ is shown in Figure 3.5 on page 33. The graph is normalized such that $P_{\mathrm{S}} = P_{\mathrm{E}} = 1$ when $R_{\mathrm{S}} = R_{\mathrm{E}}$. We can make several observations on this graph. First, the absorbed and reflected power do not sum to a constant along the graph. This counterintuitive result stems from the fact that some power is *transmitted* through the tag without being absorbed or reflected. Two interesting end-cases for this relation are when $R_{\mathrm{E}} = 0$ (short-circuit), in which the received signal is completely reflected, and when $R_{\mathrm{E}} \to \infty$ (open circuit), in which it is completely transmitted. In both of these situations zero power is available for the tag's internal circuitry. The maximum effective power available to the tag is found when $R_{\mathrm{S}} = R_{\mathrm{E}}$, indicating the power extraction can never reach an efficiency of more than 50%.

This detail of the EPC protocol is actually quite significant in our attack – the backscatter modulation method, used by the tag to send data to the reader, is also how power consumption data is leaked to the adversary. To get a taste of this effect, assume that the tag's resistance at the idle state is equal to $R_{\mathrm{S}}$. A momentary power draw caused by a transition in the control circuitry causes the tag's effective resistance to lower momentarily. Turning again to Figure 3.5 on page 33, we can see this will cause the tag's *working point* to move left on the graph, resulting in a momentary surge in the amount of reflected power.

**3.1.2. The Application Layer.** The two generations of EPC protocols define many commands which can be exchanged between tag and reader. Most of them are outside the scope of this work. We will survey two interesting areas of functionality – the tag singulation protocol and the kill command. More commands are described in [**5**] and [**19**].

It is important to recall throughout this subsection that the primary requirement from EPC tags is for them to provide their 96-bit ID (or payload) to the reader, so that the item to which they are attached may be identified. A secondary requirement is for them to be programmable by a reader – for a fresh tag to be provided with an ID or for an existing tag to be rewritten.

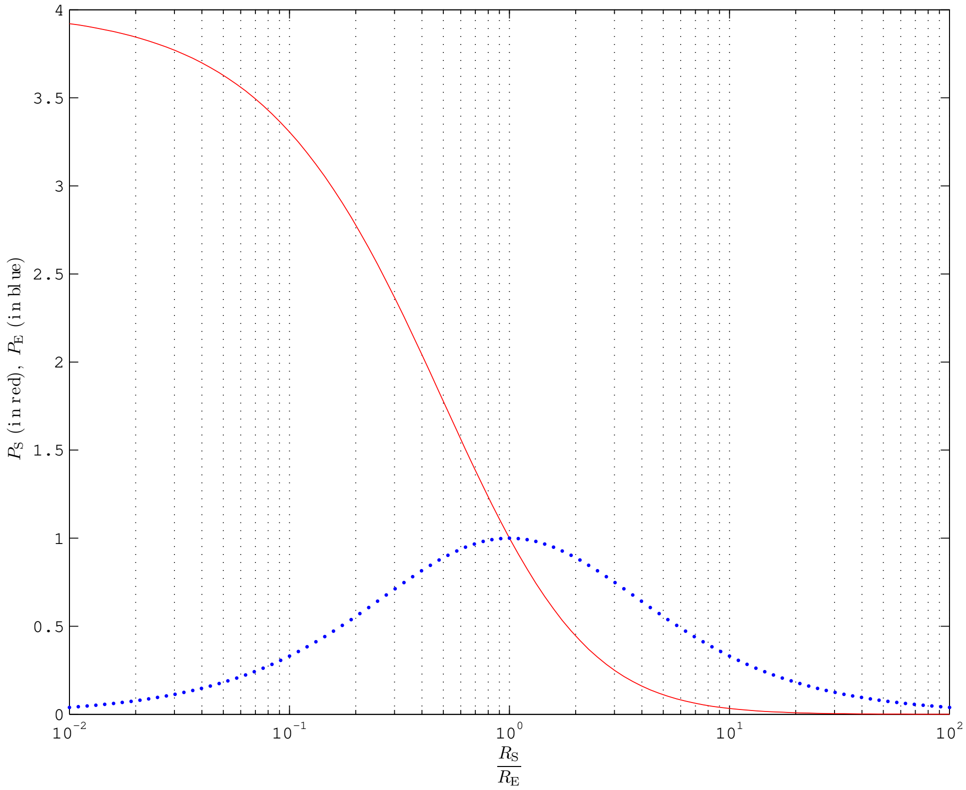


FIGURE 3.5. The relation between internal resistance and reflected power, based on [**12**, p. 124].

3.1.2.1. *Tag Singulation.* Radio is inherently a broadcast medium. At any time the radio environment may contain many devices, all sharing the same wireless channel. This is especially the case when considering the primary application of EPC tags in supply chain management (supermarkets and warehouses), in which hundreds of tags may be present simultaneously in the field of one or more readers.

The **singulation protocol** is designed to allow the reader to select a single tag and communicate with it exclusively[5]. There are different singulation protocols for Generation 1 and Generation 2 tags.

The Generation 1 singulation protocol is defined in [**5**, subsection 4.2.2]. The result of a successful Generation 1 singulation is usually the tag's 96-bit payload. The payload also serves as the unique address of commands sent to the tag. There are three variants of the singulation protocol, designed to address different population densities and security concerns (see [**46**]). Assuming the tag's payload is

[5]While both generations of the protocol also offer commands that address targeted groups of tags according to some selection criteria, these commands are outside the scope of this work. The kill command we discuss is always addressed to a single tag.

known beforehand, it is not generally necessary to perform singulation against a Generation 1 tag.

The Generation 2 singulation protocol is defined in [**19**, subsection 6.3.2.10.1]. As opposed to the Generation 1 protocol, the end result of a Generation 2 singulation is always a 16-bit *random handle* and not the tag's payload. This identifier is generated afresh each time the tag is powered up, so a Generation 2 tag may not participate in most commands unless the reader explicitly performs singulation to discover the tag's random identifier.

3.1.2.2. *The Kill Command.* The kill command is designed to irrevocably disable a tag and render it unusable. This kill feature was designed as a privacy benefit, in response to concerns raised by various organizations. The kill function prevents the contents of the tag from being disclosed after it has left the supply chain and then being used to track the individual bearing the item.

Since the Generation 1 protocol was never formally ratified, vendors ended up misapplying the kill command's original intentions. Instead of irrevocably disabling the tag, most vendors chose instead to delete all tag data upon receipt of a kill command (see [**50**, p. 26]). This means that the tag still participates in protocol commands, but the ID it sends out is an all zero string. Of course, this means the bearer of the tag can still be tracked due to the *existence* of the blank tag on his person. The adversary can even rewrite the blank tag with a unique value and track this new value in the future.

The Generation 2 protocol strictly demands that tags be completely disabled when they receive a kill command. A dead Generation 2 tag should not respond to any command sent by the reader. In practice some tag vendors still disregard this requirement, but the brand we tested implemented it[6].

The Generation 1 kill command, as defined in [**5**, subsection 4.2.2] (and elaborated upon in [**50**, pp. 26]), is shown in Figure 3.6 on page 35. Both this figure and Figure 3.7 on page 36 omit for clarity several implementation-related fields, such as parity checks, message headers and preamble sequences.

An execution of the Generation 1 kill command consists of a single packet being sent from the reader. The packet specifies the tag's entire 96-bit payload, a checksum, then the 8 secret bits of the kill password, another checksum and finally the pattern for binary '1' repeated for 30 milliseconds. If the tag's payload matches the specified payload, the kill command is correct and the checksums match, the tag should erase all of its non-volatile memory and respond to no further commands from any reader. Otherwise, it ignores the command. In both cases, the tag is completely passive during the entire process and does not send any reply indicating the success or failure of the kill command.

[6]Please see Subsection 4.3.1 to see why even this requirement is still not enough to ensure privacy.

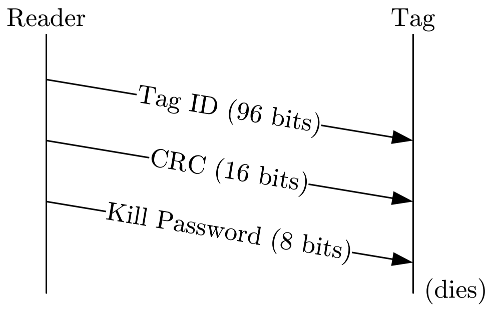


FIGURE 3.6. The Generation 1 kill command

There were several problems with the Generation 1 implementation. The first and most easily noticeable problem is the small key space – with only 256 possible kill passwords, it is trivial to conduct a brute-force search for the kill password of the tag and disable it. Another more subtle problem was discussed by [**46**] and is related to the assymetry in signal strengths between tag and reader. Since the reader emits a very powerful signal, it is reasonable to consider a passive adversary who can listen only to reader commands, but not to tag responses. During the kill command the entire contents of the tag memory are broadcast by the reader, allowing such an adversary to learn of the dead tag's former identity from a potentially much larger distance, even beyond the detection range of the tag. Finally, we must recall that vendors implemented this kill command improperly, severely compromising its security benefits.

The Generation 2 kill command was designed to solve these problems. The key space was changed from 8 bits to a more adequate 32 bits, raising the time of a brute-force attack from under a second to around 8 months. The Generation 2 protocol is designed to accomodate the assymetric channel between tag and reader – the reader never broadcasts incriminating data such as the tag's EPC code or the kill password itself. Finally, the protocol strictly defines that a killed tag should be honestly and truly dead, preventing tag vendors from merely clearing their tag memories.

Execution of the Generation 2 kill command is more complex than the previous generation. This is because tags must be singulated before being sent the kill command, and also because the entire 32 bits of the password are not sent in a single iteration, but rather in two. This design choice was made in order to trim 16 bits from the amount of internal storage required in the tag, since the tag has only to remember 16 bits of cover coding instead of 32, as shown below.

The Generation 1 kill command, as defined in [**19**, subsection 6.3.2.10.3.4], is shown in Figure 3.7 on page 36.

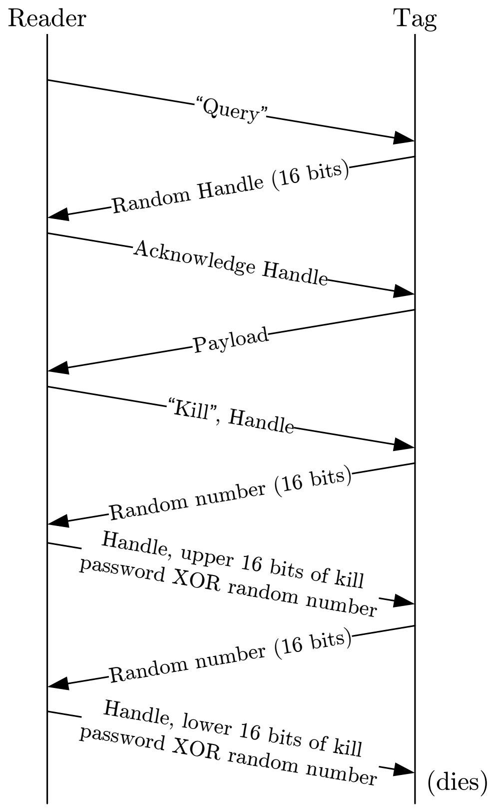


FIGURE 3.7. The Generation 2 kill command

First, the reader and tag perform the singulation protocol common to all Generation 2 commands. The singulation protocol ends with the tag identifying itself with a 16-bit random handle. Next, the reader requests 16 random bits from the tag, and responds with the first 16 bits of the kill password, XORed with the random bits the tag just sent. The protocol continues with the tag sending the reader an additional 16 random bits and the reader replying with the second half of the password. If all 32 bits match, the reader is expected to send a long stretch of CW which provides the tag with sufficient energy to delete its long-term storage. It can be noted that the tag acknowledges each 16 bit segment in turn, but only checks for correctness after the entire 32 bits of the password have been sent. This prevents the attacker from launching a trivial form of a meet-in-the-middle attack taking $2^{16}$ attempts on average to brute force the two halves of the password seperately.

The form of data exchange in which one party's transmission is XORed with the random challenge sent by the other is called *cover coding* in the EPC parlance. Cover coding makes it neccessary for an adversary to capture both the reader's transmitted request and the tag's response before it can discover the kill password, protecting against asymmetric attackers.

Careful observation will show that the tag can be designed so that it never uses much more than 48 bits of temporary storage throughout the kill protocol (16 bits for the handle, 16 for the cover coding, 16 bits for CRC calculation), and that 32 of these bits are explicit outputs of the internal 16-bit random generator and are not further manipulated. Random access memory is a very scarce resource in low-cost RFID tags, costing at least 6 gate elements per bit. Minimizing the gate count is one of the most effective tools RFID circuit designers have to minimize the cost and increase the range of their devices. This fact leads us to believe that tags do not store the entire 32 bits of the kill password in memory, but instead check the first 16 bits and carry only a single bit of state (whether this half of the password was correct or incorrect) into the second 16 bits. This behaviour should be also observable by power analysis (although we did not have the time to run this experiment), theoretically allowing the attacker to launch a $2^{16}$ time meet-in-the-middle attack.

## 3.2. The Parasitic Backscatter Channel

As stated in Subsection 3.1.1.5, momentary changes in the internal resistance of the tag result in changes to the strength of the tag's reflected field. This backscatter channel is used by the tag for intentional communication with the reader. We set out to examine whether we could also observe the minute changes in internal resistance which result from the internal state of the control circuits, thus enabling a power analysis attack from a distance.

**3.2.1. Estimating the Power Consumption from the Reflected signal.** The power supplied to the tag by the reader is shared by two consumers – the power reflected by the tag and the power it consumes internally. Because of this fact, the reflected power $P_{\mathrm{S}} = I^2 R_{\mathrm{S}}$ is a function of the tag's internal power consumption $P_{\mathrm{E}}$. Taking (3) and solving for $R_E$, we obtain:

$$R_E(t) = U_0 \sqrt{\frac{R_{\mathrm{S}}}{P_{\mathrm{S}}(t)}} - R_{\mathrm{S}} \tag{4}$$

Assuming $U_0$ is known and $R_{\mathrm{s}}$ is constant, (4) gives us a direct way of obtaining the power consumption of the tag by measuring its reflected power.

There are several simplifications that have to be noted at this point. First, we assumed the tag's load is purely Ohmic. This may not be true, but it is certainly a good enough approximation of the *instantaneous* resistance of the tag. Second, we assume $U_0$ is well known. In fact, $U_0$ is the time-varying field generated by the

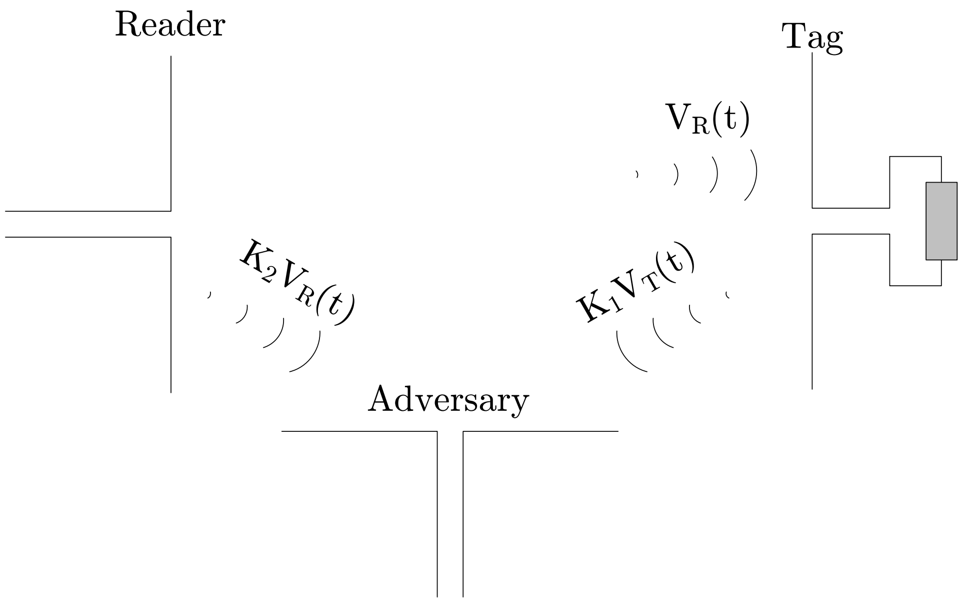


FIGURE 3.8. The multiple sources of the adversary's trace

reader and may contain noise or undesirable artifacts. Finally, it assumes we can accurately measure the power reflected from the tag in the presence of the much stronger signal generated by the reader itself. As we will see, these simplifications do not prevent our attack.

**3.2.2. Methods of Attack.** All of our attacks had a common structure. Using a cooperating reader, we sent a series of kill commands with incorrect passwords to the tag under attack. We then measured the power reflected over time from the tag, taking care to minimize the effect of the reader on our measurements. Taking several such traces and comparing them, we tried to measure the effects of different password values on the shapes of the traces.

While the signal reflected from the tag has a strong dependence on the tag's power consumption, it also has other external influences which are not found in conventional power analysis traces. To show this fact, let us derive a simple expression for the reader signal $v_R(t)$, as incident on the tag's antenna:

$$v_{\mathrm{R}}(t) = a_{\mathrm{R}}(t)\cos(2\pi f_c t + \varphi) + n(t) \tag{5}$$

In this equation the reader signal $v_{\mathrm{R}}(t)$ is a sinusoid with carrier frequency $f_c$ (typically in the area of 900MHz for UHF tags), amplitude modulated by the varying signal $a_{\mathrm{R}}(t)$, and finally suffering from some additive noise. As mentioned in Subsection 3.1.1.4, the changes in $a_{\mathrm{R}}(t)$ are used by the reader to provide the tag with commands and with their parameters. While (5) should also include some path losses due to the distance between the tag and the reader, we chose to absorb them into $a_{\mathrm{R}}(t)$.

This received signal is backscattered by the tag with a varying reflection coefficient determined by the tag's power consumption, as dicussed in Subsection 3.2.1:

$$v_{\mathrm{T}}(t) = K a_{\mathrm{T}}(t)\, v_{\mathrm{R}}(t) + n_{\mathrm{T}}(t) = K a_{\mathrm{T}}(t)\, a_{\mathrm{R}}(t) \cos(2\pi f_c t + \varphi) + n(t) \tag{6}$$

Finally, this signal is received on the attacker's antenna combined with the reader's signal:

$$v_{\mathrm{A}}(t) = K_1 v_{\mathrm{T}}(t + \varphi_1) + K_2 v_{\mathrm{R}}(t + \varphi_2) \tag{7}$$

Where the phase differences stem from the different distances the reader and tag signal have to travel before they reach the attacker. While there may be some additional data encoded in the phases, our AM receiver was not designed to make use of them, so we ignore them from this point on. Substituting (6) into (7), we see that the intercepted signal is actually an amplitude-modulated version of the reader's signal, which is itself amplitude modulated:

$$v_A(t) \approx K_3 \left(K_4 + a_{\mathrm{T}}(t)\right) a_{\mathrm{R}}(t) \cos(2\pi f_c t + \varphi) + n(t) \tag{8}$$

We can now pass this signal through our AM demodulator to receive our amplitude trace:

$$T(t) \approx K_3 \left(K_4 + a_{\mathrm{T}}(t)\right) a_{\mathrm{R}}(t) \tag{9}$$

The value of $a_{\mathrm{T}}(t)$, which is somewhere inside the above equation, is functionally equivalent to the traces provided as an input to conventional power analysis attacks, and our goal is to extract it.

Let us first examine a reader which is transmitting a sinusoid of constant amplitude $\left(a_{\mathrm{R}} \equiv \mathit{const}\right)$. This mode of transmission is called carrier wave, or CW, in the EPC specifications. In the case of a reader transmitting a CW signal a wideband AM receiver tuned to $f_c$ would be enough to extract the value of $a_{\mathrm{T}}(t)$ and send us on to the power analysis part. A long stretch of CW, however, is rarely found in the attack-worthy parts of current protocol implementations[7].

We are forced, then, to find a practically plausible way of extracting the tag-contributed signal $a_{\mathrm{T}}(t)$ from the above signal. We are fortunate in the fact that, compared to $a_{\mathrm{T}}(t)$, the reader signal $a_{\mathrm{R}}(t)$ is both slower-varying and more predictable. Several approaches to extracting $a_{\mathrm{T}}(t)$ from the combined trace are

[7] One can argue that in the case of a processing-intensive task, such as an AES encryption or a modular exponentiation, the reader will actually try to provide the tag with as much power as possible, so stretches of CW will probably be the norm and not the exception when attacking cryptographically-enabled tags.

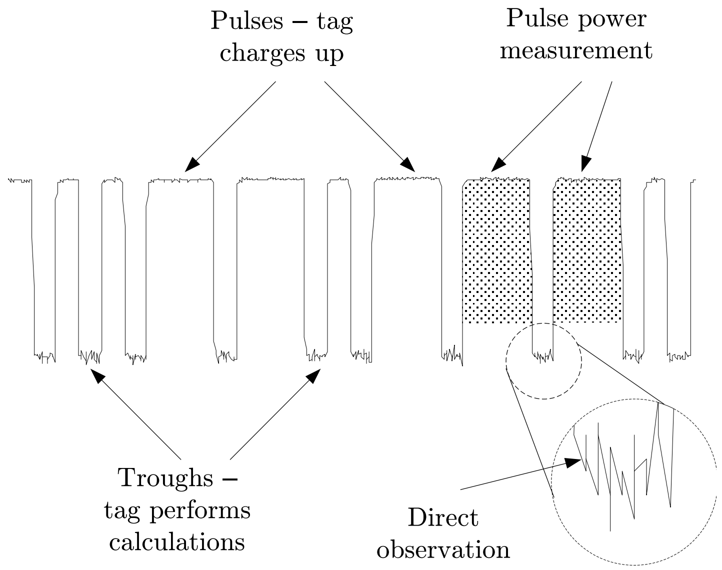


FIGURE 3.9. The attack methods compared

presented below, as well as advanced approaches we did not have the resources to try.

**3.2.3. The Direct Observation Attack.** In this attack we attempted to directly capture $v_{\mathrm{A}}(t)$ and then analyze it offline. Our wideband reciever performed AM demodulation for us, leaving us with a trace of the form $\left(K + a_{\mathrm{T}}(t)\right) a_{\mathrm{R}}(t)$.

The main problem with this attack is an instrumentation problem – the intercepted signal has a very large amplitude range, most of which is caused by the fluctuations in the reader signal $a_{\mathrm{R}}(t)$, while modern digital oscilloscopes only have about 1% accuracy in the vertical scale. This means that we had to choose between capturing the whole gamut with a high measurement noise or limiting the measurement to parts of the vertical scale and risk losing meaningful data. As we will see later, for our attack it sufficed to look only at the tops of the peaks of the original signal.

**3.2.4. The (Theoretical) Differential Observation Attack.** The instrumentation problem encountered in the direct observation attack could be solved if we could somehow cancel out the effect of $a_{\mathrm{R}}(t)$ on the trace. In the following subsection we use $T_1 = \left(K_1 + a_{\mathrm{T}}(t)\right) a_{\mathrm{R}}(t)$ to indicate the AM-demodulated trace received by the adversary.

First, assume that the attacker can explicitly generate $a_{\mathrm{R}}(t)$. This is not so far-fetched – the EPC protocol is well known and the sequence of bits sent by the reader is either fixed or easily predictable. The exact shape of the waveforms generated

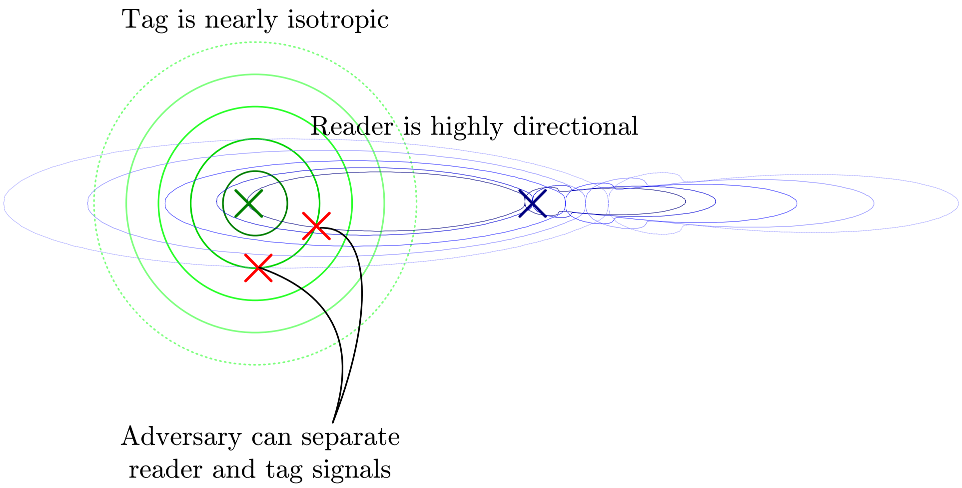


FIGURE 3.10. Using the directionality of the reader to reconstruct the reader signal

by the specific reader under attack can be easily recovered by the adversary by monitoring the reader when no tag is present. Since the reader is usually in a fixed location and has a robust power supply, there is only a very low amount of variation in the shapes of the signals it sends out, allowing many traces to be averaged together to arrive at a reliable estimate. One should also note that the regulatory demands on the reader's transmitted bandwidth are very strict – only $100KHz$ in some cases [**19**, Annex G] – so the amount of noise which cannot be filtered out is minimized.

Assuming the adversary can directly estimate $\hat{a_{\mathrm{R}}}(t)$, the attack can now be performed on the signal $T_1 \cdot \frac{1}{\hat{a_{\mathrm{R}}}(t)}$. The arithmetic operation performed on this waveform is simple enough to be carried out by an external circuit before entering the digital oscilloscope, thus minimizing the dynamic range of the captured signal and maximizing the scope's vertical sensitivity. We must note the price to be paid by taking this approach in terms of the measured noise. $T_1$ is an AM-demodulated version of a signal corrupted with additive noise. We now note that this trace $T_1$ is multiplied by $\frac{1}{\hat{a_{\mathrm{R}}}(t)}$, and thus the noise is also subjected to this amplification. At points in time in which the reader transmits a very weak signal the noise will be subjected to a high degree of amplification and the final trace will be very unreliable.

In some cases the adversary may be unable to predict $\hat{a_{\mathrm{R}}}(t)$. This may be because of the limited resources of the adversary. It could also be the case that the reader's signal is sent from a less predictable source, such as a portable reader or even a fixed reader in a crowded, changing environment. In this case the previous approach can be simulated by using a pair of antennas, each located in an area in which the reader's field is recieved with different magnitude. Because readers are usually meant to interrogate only tags in a specific location (such as tags passing through a gate or on a conveyor belt) they are usually quite directional, meaning that the distribution of their transmitted power in space will have large fluctuations even over a short distance, as illustrated in Figure 3.10 on page 41. In our tests we could see a 27dB (x600) difference in the magnitude of our reader's field between two places located 10cm apart. The tag's dipole antenna, on the other hand, has a relatively uniform power distribution on all locations equidistant from the dipole, as long as the receiving antenna is oriented in parallel to the tag.

Using a pair of antennas will leave the adversary with two traces of the form $T_1 = \left(K_1 + a_{\mathrm{T}}(t)\right) a_{\mathrm{R}}(t)$ and $T_2 = \left(K_2 + a_T(t)\right) a_{\mathrm{R}}(t)$, where $K_1$ and $K_2$ hopefully differ by orders of magnitude. By solving the two eqations for $a_{\mathrm{T}}(t)$ we obtain:

$$a_{\mathrm{T}}(t) = \frac{T_2 K_1 - T_1 K_2}{T_1 - T_2} \qquad (10)$$

If we assume that $K_1 \gg K_2$, we can use the following approximation:

$$a_{\mathrm{T}}(t) \approx K_1 T_2 \cdot \frac{1}{T_1 - T_2}$$

This mathematic manipulation can still be applied before the digitizing step, allowing the approximation of $a_T(t)$ to be captured directly using the scope's highest measurement sensitivity. One nice feature of this equation is that it does not require the adversary to estimate neither $K_1$ or $K_2$ beforehand.

Note that the drawback of the previous approach manifests itself even more powerfully in this case. Since both $T_1$ and $T_2$ are corrupted by noise, which is uncorrelated to $a_{\mathrm{R}}(t)$, the value of $\frac{1}{T_1 - T_2}$ will have strong fluctuations when $a_{\mathrm{R}}(t)$ is low and cause the approximation of $a_{\mathrm{T}}(t)$ to be even less reliable as $a_{\mathrm{R}}(t)$ grows weaker.

The differential approach can also be enhanced using an array with more than one antenna and DSP beamforming techniques.

**3.2.5. The Pulse Power Attack.** This attack is based on the observation that significant decisions about the correctness of the password are made once per reader bit. Both generations of the air interface use pulses of differing widths to differentiate between 1 and 0 symbols. The tag's demodulator's decides on the value of the bit at the falling edge, which incidentally comes at a time when the tag's internal power storage is relatively full (see Subsection 3.1.1.5). It is reasonable to assume that bit-dependent computations are then performed at the trough between two consecutive pulses, at which time the tag receives very little power from the reader and uses its internal storage to power itself. We can assume, then, that the tag will attempt to replenish this internal storage during the next pulse it receives, and that it would be "thirstier" if it had to flip the values of many bits during the previous trough. Integrating the power consumed by a tag over the period of an entire pulse will then give us an indication of how hard the tag worked after the previous falling edge. Because it measures over a relatively long period of time, this attack is less sensitive to noise, again at the risk of losing some data. We also believe this form of attack is the most easily adaptable to low-cost attack devices.

**3.2.6. The Probing Attack.** In this form of attack, we illuminated the tag with a *probe signal* consisting of a pure sine wave of constant amplitude at one frequency, while performing a normal transaction with a reader tuned to a slightly different frequency. The amplitude of the probe signal was made as low as possible, so that it in itself will not be enough to provide the tag with power. If the reader and probe frequencies are set far enough apart, the amplitude of the bounced probe signal will only indicate the power consumption of the tag without including any residual data from the reader. This allows us to get a lower dynamic range and thus capture the entire reflected waveform at high vertical accuracy. This attack has the disadvantage of requiring additional equipment and of announcing the presence of the adversary. Our results do not make use of the added power offered by this attack, although it seems to have practical advantages, especially when looking into time segments with low or unstable reader power.

CHAPTER 4

# Our Attack in Practice

This section will discuss the physical aspects of our attack and present our results. We will begin by describing the physical and logical structure of our lab setup, along with the design choices leading to this final setup, then present our results.

## 4.1. Lab Setup

**4.1.1. Physical Setup.** As discussed in the previous chapter, our attack required us to send reader commands to a tag and measure the tag's backscatter as accurately as possible. Since the lab was being built basically from scratch, we also needed the lab setup to be relatively inexpensive. To achieve this we ended up renting or borrowing most of the equipment, and even manufacturing some of it in-house. An additional requirement was to have a setup which is as portable as possible, in case we needed to run tests in a nearby anechoic chamber. We also needed the ability to carry out long unattended test runs by creating automated scripts.

Our final lab setup, built to satisfy these requirements, is shown in Figure 4.1.1 on page 46. From right to left, the figure shows the wideband receive antenna (Fratcus EZConnect) ④, the tag ③, the directional antenna (MaxRad MP9026) ②, the digital oscilloscope (Lecroy 9304C) ⑥ on top of the wideband receiver (Agilent-HP E4405B-AYX) ⑤, [**12**], and finally the PC containing an internal RFID reader (WJ MPR-6000) ① and our data collection software ⑦. The HP E4405B is actually a spectrum analyser with a baseband output, but we used it only as a very sensitive AM receiver. For the probing attack the setup was augmented with a HP 8530 swept signal generator (not shown), configured to send out a sine wave of constant amplidute at 900 MHz.

A logical view of the setup can be found in Figure 4.1.1 on page 46. The PC was connected to the reader via its internal PCMCIA bus. The reader was connected to the antenna using a coaxial MMCX cable. The directional antenna, tag, and, wideband antenna were located on a lab table at a distance of about 1m. The wideband antenna was connected via coaxial BNC cable to the wideband receiver's RF input, as were the receiver's baseband video output and the digital oscilloscope's

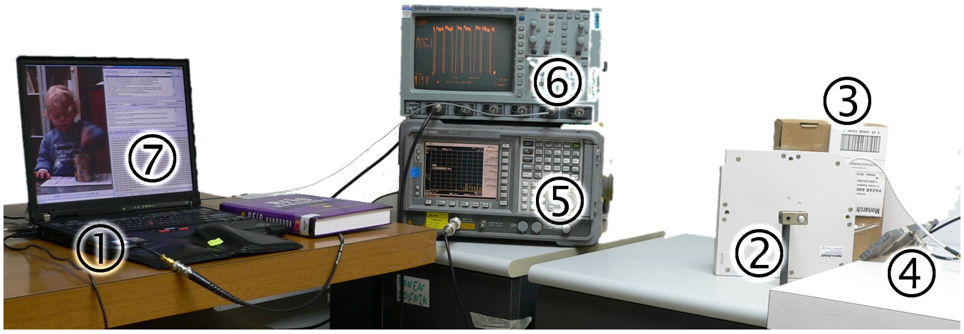


FIGURE 4.1. Our lab setup

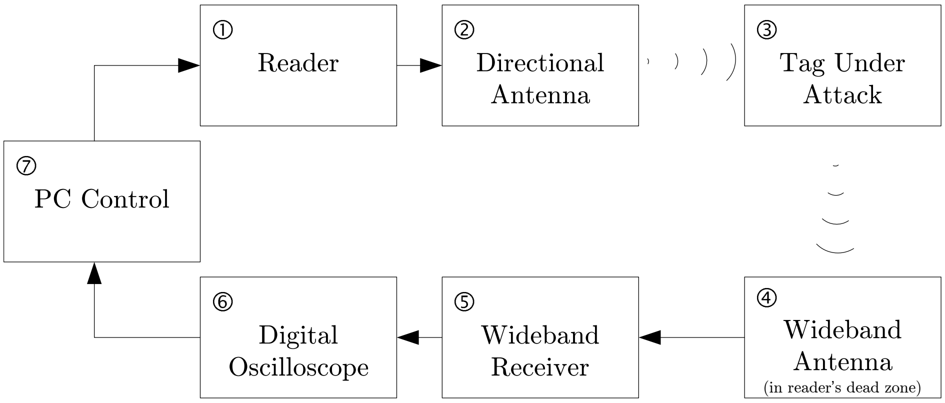


FIGURE 4.2. Block diagram of lab setup

signal input. Finally, the PC was connected to the scope using both a modified Centronics parallel cable (for sending external triggers) and a serial RS-232 cable (for retrieving traces).

We will now describe each component in this setup. In Subsection 4.1.2 we will describe how these components worked together to help us carry out our attack.

4.1.1.1. *Wideband Receiver.* Since we were interested in extracting data from an amplitude modulated signal, the central piece of equipment in our setup was the wideband AM receiver. Recall that an amplitude modulated signal has the form $a(t)\cos(2\pi f_c t)$. In general, wideband AM receivers are differentiated by two main characteritics: their *frequency range* and their *resolution bandwidth.* The *frequency range* of an AM receiver determines the allowable frequencies of the carrier wave $f_c$. In our case the carrier frequency was 860–960MHz, as defined by the EPC standards. The *resolution bandwidth* of the receiver determines the maximum possible bandwidth of the data bearing signal $a(t)$. In our case, the resolution bandwidth requirement was derived from the advertised clock rate of generation 2 EPC tags, 1.92MHz[**47**], which meant that we needed at least 4MHz of resolution bandwidth (and preferably much more) to be sure that our signal captured all data. FM radio, for comparison, has a maximum bandwidth of 75KHz [**52**].

Table 1 on page 48 describes several alternatives for use as a wideband reciever. The most straightforward (and most expensive) option was a wideband TEMPEST receiver such as the Dynamic Sciences R-110B. Designed specifically for EM attacks, these receivers have a very high resolution bandwidth and very low internal noise levels (see [**31**, subsection 2.2] for more praise of these devices). Unfortunately, these devices are very bulky and expensive and generally require additional instrumentation (PC with data acquisition card, spectrum analyser) to be truly effective. In addition, the more advanced TEMPEST receivers are classified as munitions by the US authorities and require a special export licence.

On the cheap end, one very interesting option was to use a tuner box taken out of a broken-down TV or VCR[1]. UHF television has similar characteristics to our signal - the bandwidth allocated to a single station is 8MHz, and the frequency band is 400MHz-860MHz – just below the output range of our reader[**53**]. The authors of [**41**] demonstrated that a tuner box can be pushed to work reliably at frequencies as high as 1GHz. We actually performed some initial experiments with a tuner box taken apart from a discarded VCR, provided by a helpful local TV repair shop. There are several advantages to using this tuner box, as opposed to building our own wideband receiver – it has a shielded enclosure that keeps noise down, an industry-standard antenna input jack, and a frequency range compatible with UHF RFID tags. It also has a price that can't be beaten. Its disadvantages – a relatively clumsy set-up requiring three laboratory power supplies, an intermediate-frequency output which is still beyond the range of low-cost scopes and requires an additional downconverter[2] to come down to baseband. Perhaps most significantly, using a TV tuner would have introduced an additional unknown quantity to an already risky endeavor. We think a low-priced attack setup can certainly be built around such a device, now that the existence of the parasitic backsatter channel is established.

Our final choice for the wideband receiver was a HP/Agilent spectrum analyser. The model we picked was the mid-range portable E4405B model, with the additional AYX option which provided a baseband output we could connect to a digitizing scope. The spectrum analyser met our performance requirements at a much cheaper

[1]According to our helpful TV repairman, the most fault-prone component of a TV is its cathode-ray tube, while the most fault-prone component of a VCR is its tape handling mechanism. Most discarded TVs and VCRs have fully functional tuner boxes.

[2]The baseband analog television signal has a non-trivial internal structure, consisting of a pair of video and audio signals which are each individually subjected to frequency modulation. For that reason, it would have been problematic to keep the tuner box inside the TV and then use the TV's own internal down-converter and demodulator.

| Property | TEMPEST Receiver[**31**] | Spectrum Analyser | Tuner Box[**41**] |
|---|---|---|---|
| Model Compared | Dynamic Sciences R-110B[**21**] | Agilent/HP E4405B-AYX[**51**] | Panasonic NV-7200[**34**] |
| Frequency Range | 1KHz–1GHz | 9KHz–1.5GHz | 80MHz–920MHz |
| Resolution Bandwidth | 50Hz–200MHz | 10Hz–5MHz | 8MHz |
| Output Format | Baseband | Baseband | 47MHz Intermediate Frequency |
| Price in late 2006 | $130000 | $29000 ($1K/month lease) | Free (broken-down VCR required) |

TABLE 1. Our wideband receiver, compared to expensive and inexpensive alternatives. Our requirements were a frequency of 900MHz and a resolution bandwidth of at least 4MHz

price and with higher portability than the TEMPEST receiver. We were even able to lease this spectrum for a monthly fee, further reducing our total costs.

There are some advantages to using a spectrum analyser as opposed to the TEMPEST receiver. Most notably, the spectrum analyser has the ability to function as a, well, spectrum analyser, a fact that was instrumental whe trying to understand the frequency-hopping behaviour of the reader and the directionality of the antennas.

4.1.1.2. *Transmit and Receive Antennas.* The reader antenna we used was directional, as are the antennas attached to most deployed readers. Directional antennas tend to have *null zones*, in which the power of the field radiated from the antenna approaches zero. The tag's antenna, on the other hand, is a relatively omnidirectional dipole. As long as we placed the adversary's antenna in one of the reader's null zones, we would practically pick up nothing but the tag's signal. This fact allowed us the luxury of attaching an small omnidirectional wideband antenna, properly located, to the adversary's receiver.

For our attack, we used the pair of antennas supplied to us with the MPR-6000 reader development kit. One antenna (the MaxRad MP9026[**16**, p. 56]) was a highly directional panel antenna. We attached it to the reader. The other (the Fractus EZConnect[**43**]) was an omnidirectional chip antenna. We attached it to the adversary's receiver. Both were specified to work in the frequency range of 902–928MHz, but the small chip antenna had better performance than the panel antenna outside this advertised frequency range.

A practical attacker working in an RF-saturated environment will most probably attach a directional antenna to his receiver. Most directional antennas force some trade-off between directional gain, bandwidth and the size of the antenna. However, the only RF sources found in our lab setup were the tag and the reader.

The 900MHz signal we were dealing with was influenced strongly by the physical environment. Placing a hand near the lab table was enough to offset the experiment's result. We even had to replace a flexible tripod we had originally used for locating the tag in space with a cardboard box, due to the strong reflections from the tripod's metal base. This sensitive behaviour all but forced us to use automated measurement tools and run unattended tests.

4.1.1.3. *RFID Reader.* The RFID reader we picked was the MPR-6000 from WJ Communications. This RFID was chosen for its tight PC integration, which helped us run our tests more efficiently. While most readers we could find were designed to be standalone wall-mounted devices, complete with rugged packaging and a very narrowly-defined programming interface, the MPR reader is plugged into a computer's PCMCIA port and can be controlled relatively well by the PC. A developer's kit we ordered contained the reader, a few sample tags and two antennas - one directional and one omnidirectional. Sadly, the kit we received was assembled in 2004 and did not have generation 2 tags or the firmware required to support them. We were able to upgrade the firmware with help from the manufacturer, and generation 2 tags were provided to us by several local suppliers.

The reader identifies itself to the PC as an extra serial port, using a well-described command language to send and recieve commands to tags [**20**]. The reader's firmware offers some encapsulation of the low-level EPC protocol. For example, sending a Generation 2 kill command requires only a single command to be sent to the reader's firmware, with the firmware performing the singulation and going through the kill protocol. The reader was also supplied with a demo application that could read, write and kill tags. Examining the debug outputs of this program and mimicking its behaviour was very useful in the development of our own control software.

In compliance to FCC regulations[**40**, §15.247], the WJ reader uses frequency-hopping spread spectrum modulation to minimize its impact on other users of the unregulated ISM band. This meant that the reader progresses through a sequence of carrier frequencies, sending each command on a different pseudorandomly chosen carrier frequency. This behaviour was very confusing for us at first, since this progression was over a frequency range of 25 MHz, while our receiver had a resolution bandwidth of only 5 MHz, resulting in some commands we completely missed. Fortunately, we discovered that the pseudorandom sequence had only 50 entries and that the reader switched frequencies only once every user command, even though the firmware may translate this command into any number of over-the-air transactions. To allow all measurements to be performed at a constant frequency we simply followed each relevant command we wished to send by 49 dummy commands (Generation 0 reads, to be exact).

4.1.1.4. *Digital Oscilloscope.* The digital oscilloscope used in our experiments was a LeCroy 9304C[**7**]. We were lent this scope by Prof. Amir Yacoby, to whom we are grateful. This scope offers a 50,000 sample memory at its top sampling rate of 100 Megasamples per second, for a total of 500$\mu$Sec per trace. The scope can be controlled by a PC over a serial port, at speeds of up to 56Kbps. The PC can use this interface to send commands to the scope and to download trace data. The traces were provided in a proprietary format, and we wrote a small script to translate them into Matlab-compatible files.

The main drawback of the 9304C was its low vertical resolution. The scope we used has an 8-bit ADC, allowing only 256 possible output levels. Any value falling between two sampling points will be subjected to a rounding error, also called the *sampling noise.* The average rouding error in this case is approximately 0.2% of the total vertical range chosen for the trace. In our experiments we could see that the amplitude of the parasitic backscatter signal was about 0.6% of the reader signal's amplitude. Thus, we were forced to choose either to zoom in vertically on the trace or to capture the entire vertical range and suffer an unbearable amount of sampling noise. In both cases we risked losing so much information that the trace would be useless. The documentation of the scope suggests that the effective vertical resolution can be raised either by averaging many traces together or by passing the trace through a low-pass filter. The drawback of the filtering approach is that it lowers the effective sampling rate of the scope to the cutoff frequency of the filter. The scope has several built-in filters which trade off an effective resolution gain (in bits) versus a reduced scope bandwidth. Since the scope has a relatively low sampling rate of 100 Megasamples per second, we could only gain one more bit of vertical resolution using this method without going below our required effective sampling rate of 10 Megasamples per second. The results contained in this thesis reflect the result of averaging many traces without additional low-pass filtering.

The relatively small sample memory of the scope, together with the need for averaging many traces of the same event, required us to have good control over the exact point in time in which the scope was triggered. We initially achieved this objective by connecting a custom-made cable from the PC's parallel port to the scope's external trigger input. At some point in time after the command was sent to the reader, we programmed the controller software to send a pulse to the scope through one of the parallel port's data lines. The exact value for the delay between the command and the trigger was specified as part of the control script.

This approach was problematic for several reasons. First, we could not achieve the sub-millisecond time precision required to propery align a 500$\mu$S sampling window. There were simply too many sources of random time jitter in our setup – the controller software was run on a multitasking Windows machine with limited time resolution and arbitrary scheduling, and the WJ reader's firmware and serial interface introduced additional uncertainty. All in all, only 30% of the traces recorded through this method were of any use for our calculations. A more fundamental problem with this approach was the intimate connection between the reader and the adversary – a connection not likely to be found in a real-world attack scenario.

We improved on this situation by making use of a scope feature called "smart triggering". In smart trigger mode, the scope can be defined to trigger when a certain condition on its inputs is satisfied. In our case, we counted the amount of transitions between low and high levels on the input. The smart trigger was not fazed even by the Generation 2 pseudorandom cover coding (see Subsection 3.1.2.2) since both '1' and '0' Generation 2 symbols have 2 rising edges per bit (see Subsection 3.1.1.4). We still needed some guidance from the PC, in the form of a signal from the parallel port, so that the scope would be able to tell apart the "real" command from the 49 "dummy" commands we used to get around the frequency hopping. However, the time variation of this external trigger did not affect the accuracy of the smart trigger, bringing us to a nearly 90% yield in later experiments[3].

Moving the traces between the scope and the PC was the most time-consuming part of our experiment. While the scope captured 50,000 8-bit data points per trigger, these traces were represented very inefficiently in the scope's proprietary format, ending up with 160K of data per trace. This data was transferred to the PC via a 56Kbps serial connection, requiring about 20 seconds to transfer each trace and a total of 2 hours for every batch of 300 good traces.

More expensive and better capable scopes do exist, albeit at a higher price than our budget allowed. A high-end modern scope offers a top sampling rate measurable in tens of gigasamples per second (which translates into better vertical resolution at lower sampling rates, aided by a strong filter). These scopes also have much deeper sample memory (as many as 64 million samples) and more advanced triggering options which could have done away with trigger hints from the PC altogether. Some scopes even offer a gigabit Ethernet port, offering a potential data rate of 1Gbps and an expected sub-millisecond transfer time per trace.

**4.1.2. Logical Setup.** Our attack was composed of a succession of identical experiments. Each experiment consisted of sending a kill command from the reader

[3]There were still some unusable traces in cases where the Windows scheduler actually got around to executing the parallel port trigger only after the radio-interface command was finished.

(with an incorrect password – we wanted to keep the tag alive), demodulating the response of the tag using the spectrum analyser, capturing the baseband signal using the digital oscilloscope and finally transferring the capture to the PC. In initial tests the scope was triggered by a signal sent from the PC while further on we configured the scope to trigger on a specific wave shape. Each attack consisted of about 300 such experiments and each experiment took a bit under 25 seconds, most of which was spent transferring data from the scope to the PC through the slow RS-232 serial port. This gave us a total time of just under two hours per attack. Considering the fact that a kill command takes about 10 milliseconds to execute, the net time of each attack (which could be easily achieved with a more integrated attacking device) was only a few seconds.

After the data has been transferred to the PC, we loaded the samples into Matlab, normalized and aligned them, and finally analyzed them both visually and via a suitable program.

The programming effort related to this project was divided into two main areas of functionality:

- A custom-coded Visual Basic application to control the reader
- A set of Matlab and perl scripts to manipulate the results

The source code for both parts is included in the companion CD. All in all, the project required a total of about 1000 source lines of code. We used the open source RCS system[**35**] for revision tracking and source control.

4.1.2.1. *The Reader Controller.* The controller program was designed to offer easy access to the MPR-6000 internal reader and allow the results of tests to be efficiently gathered and processed. We chose to write the controller in Visual Basic since VB offers excellent string handling along with access to internal Windows functionality such as parallel port I/O and sub-millisecond timers.

The controller is scriptable – it exposes an object model that can be accessed by other programs written in Visual Basic, and includes a built-in Visual Basic script compiler, based on [**10**]. This allowed us to write relatively complex unattended runs that could cycle through various scenarios.

To avoid the frequency-hopping aspect of the reader, the controller performs 49 dummy commands (Generation 0 reads) after every real command. It also has the ability to send a trigger to the digital oscilloscope after a microsecond-accurate delay (subject to jitter – see the previous subsection).

4.1.2.2. *Matlab and perl Scripts.* Since our adversary model relied on radio interception instead of connecting directly to a device's power supply, our signal was typically noisier than common power analysis traces. To overcome this fact, we collected between 80 and 300 traces per experiment and averaged them.

To properly align the traces we wrote a set of Matlab functions that applied some heuristics to the traces. We had both to decide whether the traces are worthy of being used at all (recall that only 30% of the initial traces were good enough to use), and to find the best relative displacement to combine them. We ended up trimming the top and bottom halves of the trace, normalizing it and then finding the displacement that would provide a maximum cross-correlation with a reference signal. Using this algorithm on a dual-processor Xeon server running Linux and Matlab R13, we could align approximately 5 traces per second.

We also wrote several basic shell scripts to convert the trace data from the proprietary LeCroy format to a format Matlab understands.

## 4.2. An Attack on Generation 1 Tags

**4.2.1. Objective.** The attack on Generation 1 tags was carried out in several steps:

First, we demonstrated that the signal backscattered from the tag has "interesting" information. Next, we showed that this signal is affected by the tag's power consumption in a measurable way. Finally, we showed that the contents of the tag's internal memory had an effect on its power consumption, thus opening the way to power analysis attacks.

**4.2.2. Test Execution.** The experiment began by locating the directional antenna's null zone. This was achieved by sending a constant signal through the reader antenna and moving the attacker's antenna around the reader in a circle with constant radius until we found the direction with the minimal receive power. We then affixed a Generation 1 tag to a cardboard stand and placed it on a lab table facing the directional transmit antenna. Our assumption was that the less power available to the tag, the more significant (relatively speaking) would be its parasitic backscatter. To reach the point at which the least power is available to the tag, we varied the distance between tag and reader and the reader's transmit power until the tag was at the very end of its operating range. We could observe a 28dB rise in the amount of power received on the adversary's antenna when the tag was placed on the desk. This shows that although the reader emits a signal which is much stronger than the tag's backscatter, we can effectively ignore it if we properly locate the adversary's antenna.

**4.2.3. Results.** This section shows our main result – a remote power analysis attack against Generation 1 tags.

4.2.3.1. *Differences Between the Reader Signal and the Tag's Backscatter.* This first result demonstrates that a tag modulates its backscatter even when it is supposed to be completely passive.

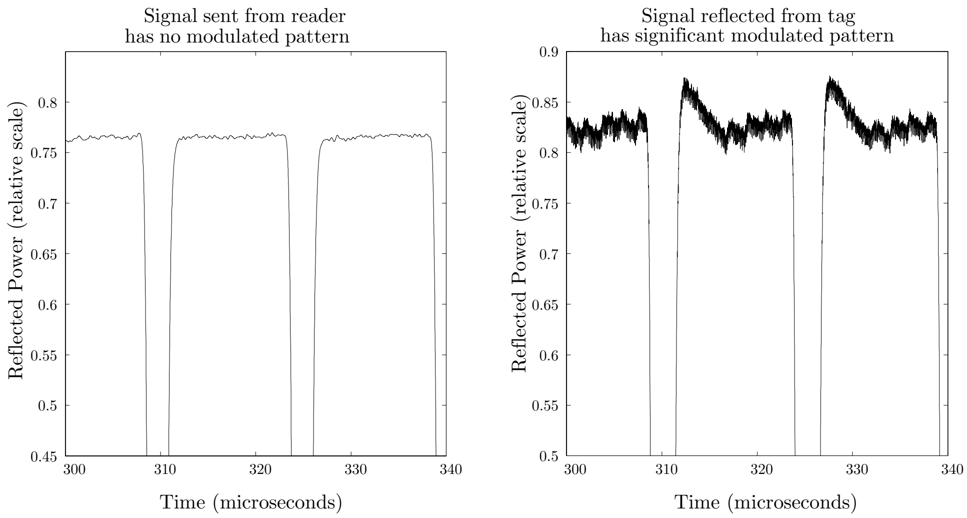


FIGURE 4.3. Signal reflected from Generation 1 tags has a significant modulated pattern

To generate Figure 4.3 on page 54, we measured both the signal sent to the tag and the one reflected from the tag as the tag was receiving a long string of zeros from the reader. In the subfigure shown on the left, the experiment was performed in the absence of a tag. In the subfigure shown on the right, the receive antenna was placed in close proximity to a tag but in the reader antenna's dead zone, giving the tag signal a 28dB advantage over the reader signal. Approximately 80 traces were averaged to create each subfigure. We stress that according to the Generation 1 protocol the tag is not supposed to be transmitting anything at this stage.

This experiment shows how the RF front end of the tag influences the signal scattered by the tag. The figures show the topmost part of the signal. The amplitude of the parasitic backscatter modulation in this case was about 22dB less than the overall peak-to-peak amplitude of the measured signal, or a bit less than 1%. As shown in Subsection 3.1.1.4, the peaks in reflected power correspond (perhaps counter-intuitively) to areas with higher power consumption.

Referring again to the various components of a tag, as described in Figure 3.1 on page 24, one can attribute the distinctive sawtooth shape to several different sources in the tag. We can quickly rule out the modulator as the source of this pattern, since the tag is not transmitting anything at the moment. The control circuit can also be ruled out, since the tag under discussion is known to perform calculations only once per rising edge, while the sawtooth pattern is finer (5 distinctive ridges for every rising edge). This leaves two possible culprits – the demodulator and the power extraction circuit.

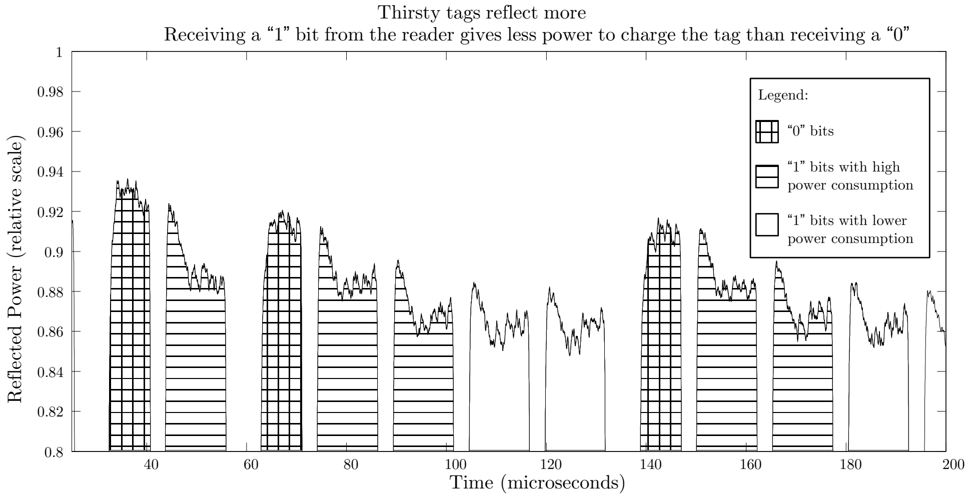


FIGURE 4.4. “Thirsty” tags reflect more power

It is probable that the demodulation circuit is based on a detector similar to the one shown in [**24**, Subsection III.C]. As stated in Subsection 3.1.1.4, the reader uses pulse width modulation, using a longer pulse to signify a “0” symbol and a shorter pulse to signify a “1”. This means the tag can be expected to measure the width of the incoming pulse to decide on the value of the incoming bit. Once a certain threshold width is passed, there is no need to continue measuring the width of the symbol, since the symbol is clearly a “0”, offering a possible way of explaining the lower power consumption. In our specific case the first large ridge seems to be too narrow to allow for this explanation. We could also see the same ridged pattern when observing a “1” bit, ruling out such a direct connection.

The other potential source for this pattern is the power extractor module. Without knowledge of the specific tag’s RF front end, it is hard for us to explain how the power extractor can create such a shape. However, we can claim that the first ridge in every pulse is taller since the power extractor has to compensate for the drain on the internal capacitor during the period of low transmit power.

4.2.3.2. *Effect of Power Consumption on Backscatter.* In this experiment we sent the Generation 1 tag a sequence of ones and zeros and measured the differences in reflected power levels.

Compared with a ’0’ bit (shown plain or with light horizontal hatching), a ’1’ bit (shown with cross-hatching) has a wider gap followed by a narrower pulse. Now, examine the wider gap before a ’1’ bit. As mentioned in Subsection 3.2.5, the tag’s internal power storage is depleted during these low-power gaps. At the end of the long gap which forms the beginning of the ’1’ bit, the tag’s power supply is relatively low. This makes it draw more power from the next pulse it receives.

As the tag consumes more power, it causes a stronger current to flow through its antenna. Because of this stronger current, the tag radiates a stronger reflected field, as shown by the cross-hatched pulses.

The '1' bit has more than a wider gap – it has a narrower pulse as well. This means the tag's power storage is not fully charged up even at the end of the cross-hatched pulse. As a result, the tag also draws more power from the next '0' bit, as indicated by the horizontal hatching. As the tag receives more '0' bits, it slowly charges up, reducing the current flowing through its antenna. This finally causes the tag to reflect less power, as witnessed in the plain areas.

While this result shows that the tag's increased power consumption is observable from the increase in the strength of its backscattered field, this in itself is not enough to show the possibility of power analysis attacks. To launch such an attack we need a sufficiently high signal-to-noise ratio in the reflected trace to provide insights about the control circuitry. However, the relatively large amount of noise that exists in the RF medium, together with the relatively simple circuitry running on the tag's control circuit (resulting in a weak signal), mean that there is still the chance that the power analysis traces embedded in the backscatter will have a SNR which is too low for practical use. We address this concern in the following subsection.

4.2.3.3. *Effect of Internal Tag Memory on Power Consumption.* In this experiment we sent the incorrect kill command to a tag which was programmed with various different IDs. As stated in Subsection 3.1.2.2, the Generation 1 kill command is simply a listing of the tag's entire internal memory. In this experiment we actually looked at a part of the internal tag memory (ITM) which was not related to the kill password. Instead, we modified parts of the payload.

The top subfigure of Figure 4.5 on page 57 was obtained while the tag was programmed with a payload of $10FF00000000_x$, while the subfigure on the bottom was obtained while the tag was programmed with $1000FF00000000_x$. The difference between the two traces is shown in the middle. Observing this difference, we can see that the traces behave in one way where the memory is identical (time 1.5–3.5*$10^{-4}$) and in another way in areas in which the tag memory is different between traces (time 3.5–5*$10^{-4}$).

This result shows the first link between the goings-on of the control circuits and the tag's backscatter. In the next subsection we show how the kill password can be examined by similar means.

4.2.3.4. *A Power Analysis Attack on the Kill Password.* Our final result against generation 1 tags shows a correlation between the kill password assigned to a tag and the tag's backscatter.

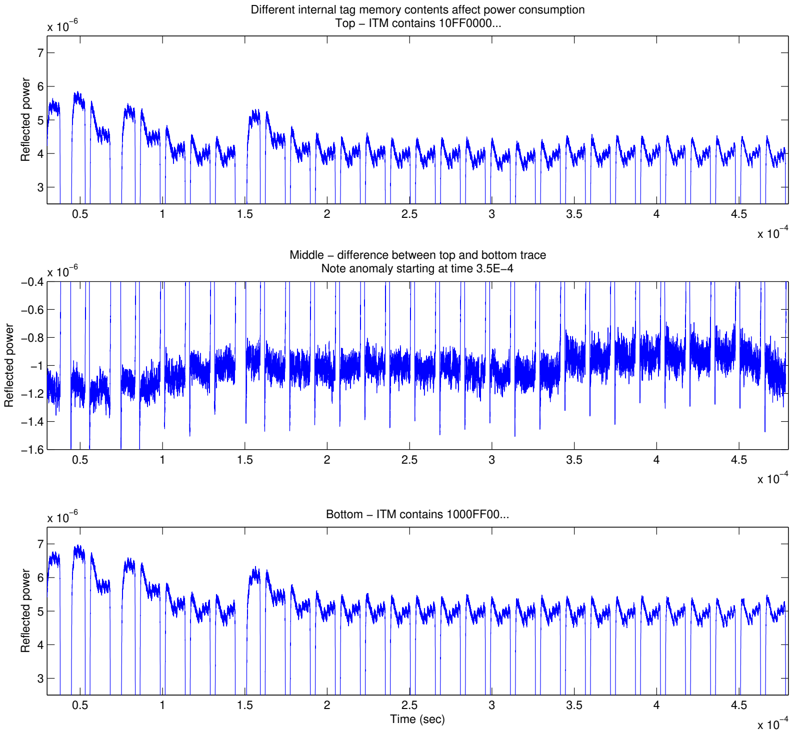


FIGURE 4.5. Internal tag memory has an effect on power consumption

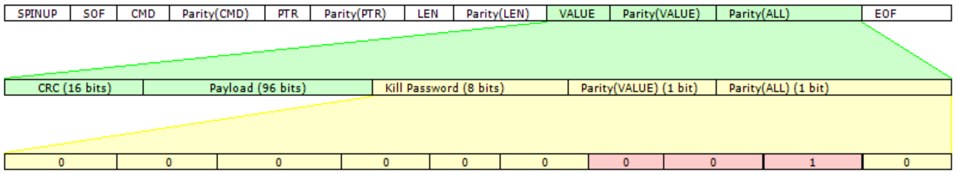


FIGURE 4.6. The location of the trace shown in Figure 4.7 on page 58

Figure 4.7 on page 58 shows a close-up view of the last 2 bits of a kill password being sent to a Generation 1 tag, followed by the first parity bit following them. Figure 4.6 on page 57 indicates in red the exact location of the trace we are about to see: the final bits of the kill password, right at the end of the VALUE parameter of the command.

The exact format of a generation 1 kill command is defined in [**5**, subsections 4.1 and 4.2.2] and described in this work in Subsection 3.1.2.2.

Let us now compare the two traces shown in Figure 4.7 on page 58. In the experiment shown on top, the tag expects a kill password of $FF_h = 11111111_b$

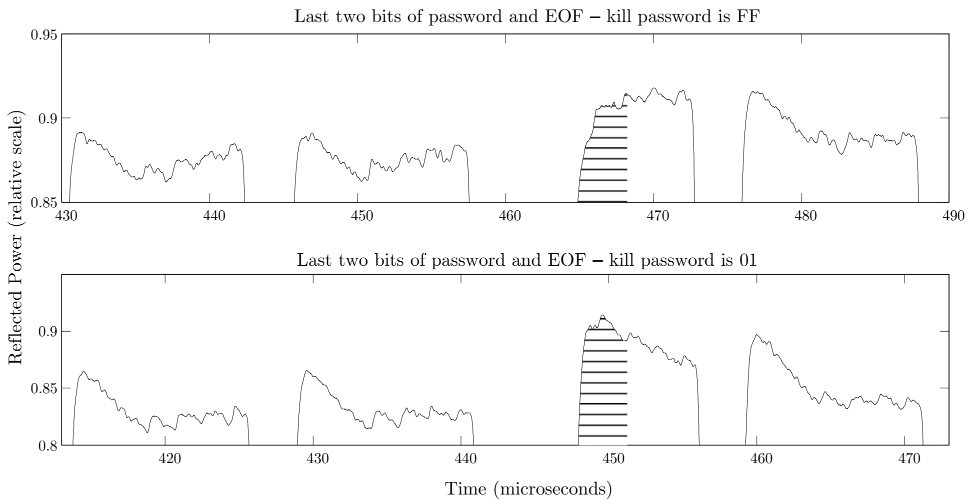


FIGURE 4.7. Recovering one bit of the kill password

(reading from left to right), while on the bottom it expects a password of $01_h = 00000001_b$. In both cases the kill password sent to the tag is $00_h = 00000000_b$.

Let us examine this situation in detail. Both experiments involved the same tag at the same location receiving the same data. Turning to the general desription of a tag, as discussed in Subsection 3.1.1.1, we note that there is no change to the inputs of the power extraction circuit or of the demodulator. Thus, it is safe to assume the differences in power consumption between the two experiments are solely the result of changes in the control logic, in the EEPROM, or in the interface between the two.

The next step requires some insight into the internal workings of the tag. Since we did not have the equipment required to reverse-engineer the tag under attack, we have to make some assumptions on the tag's design. In general, the tag should compare the password bits coming in to the stored password, and kill itself only if the entire set of 8 bits matches its stored value. There are two ways to go about this – the tag can either compare the bits coming in one by one, or accumulate the entire 8 bits in scratch memory before comparing them to the stored password in a single batch. As stated discussed in the previous chapter, tag vendors try to minimize the amount of internal storage they use. Thus, we assume that the tag compares incoming bits one by one[4].

[4]Storing batches of bits may still be a possible design choice if the extra memory was dictated by another tag feature and simply reused in the execution of the kill password. However, it is less likely that longer passwords (such as the 32-bit Generation 2 password or any future larger passwords) will be stored in their entirety. In the case that tags compare several bits at a time, power analysis can be used to discover the Hamming weight of this group of bits. As described in [**38**], this data can still be used to compromise the security of the device under attack.

We also assume that when the tag receives the first incorrect bit it will consume more power than does when it receives a correct bit. This observation is based on practical experience, and can be justified by noting that the tag has to react to an incorrect bit by transitioning to a special error state and performing other "housekeeping" tasks.

The tag in question evaluates the password bits coming in from left to right. This means that the tag shown in the top subfigure tag already knows the kill command will fail, having previously received many wrong bits. The bottom tag, however, only learns that the kill password is wrong after the falling edge identifying the last '0' bit. This "exciting" event causes the tag to carry out certain additional computations in the trough between this final "0" bit and the parity bit which follows, resulting in a larger power draw from the tag's internal capacitor. When the next rising edge arrives, the tag's power extraction circuit has to replenish the internal capacitor, requiring it to draw more power from the reader's field. The increased power consumption of the tag in the lower tag can be seen by the spike it exhibits as it starts receiving the parity bit, as compared to the gentler slope on the top subfigure, as indicated by the hatched area. This demonstrates how a single bit of kill password can be extracted via power analysis over the air.

## 4.3. An Attack on Generation 2 Tags

The objective of this attack is the same as the attack on Generation 1 tags – the recovery of the kill password using a linear time effort. There are several complications in attacking Generation 2 tags, as mentioned in Subsection 3.1.2.2. In a future work we hope to fully demonstrate how Generation 2 passwords can also be extracted by a somewhat more complicated version of the parasitic backscatter attack. What we show here is that the privacy of Generation 2 tag users can be compromised using the same attack.

The test was executed in the same way as the Generation 1 attack. However, the newer Generation 2 tags we used were more power efficient than the Generation 1 tags, requiring the tag to be repositioned to achieve borderline performance.

**4.3.1. Results.** Figure 4.8 on page 60 shows a trace similar to Figure 4.3 on page 54, comparing the signal transmitted by the reader and the signal reflected by the tag. The noticable addition of the cusp shows that the tag is modulating its reflected signal. It is also evident that tags A and B, each from a different vendor, have different RF signatures.

In our experiments we noted that a dead tag (i.e., a tag which has received a kill command with the correct kill password) presents essentially the same backscatter signature as a live tag. Dead tags do not participate in EPC inventory commands

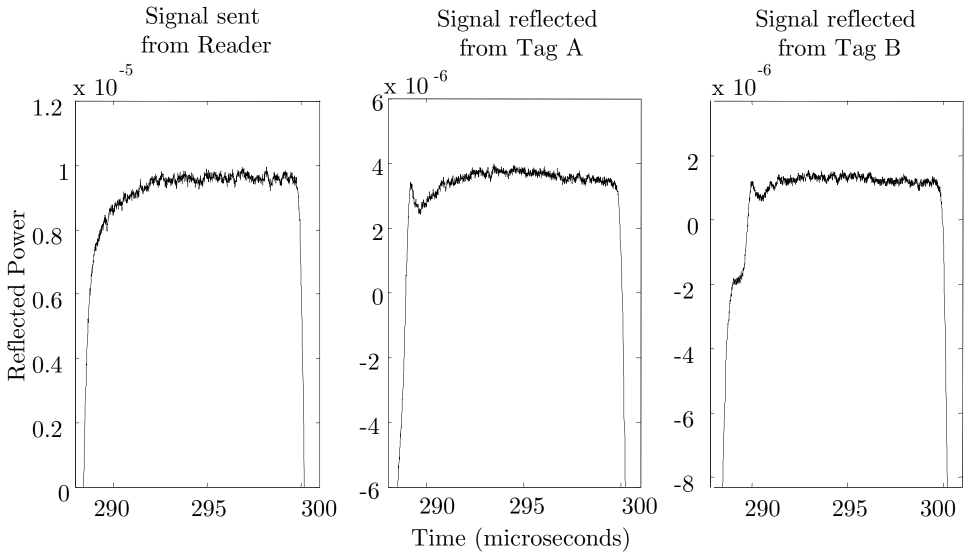


FIGURE 4.8. Signal reflected from Generation 2 tags has a significant modulated pattern, which differs between tag vendors

and, as such, are considered invisible in the conventional RFID security model. However, a killed tag's RF front end is still functional, and thus a dead tag modulates its reflected field in practically the same way it does when the tag is active. This means that the *existence* of a killed tag can be detected by an adversary using an attack technique similar to ours, even though the tag's payload has been erased as part of the kill command. The different design choices made by tag vendors in implementing their RF front ends cause each brand of tag to modulate the reader's signal in a slightly different way. Thus, not only is it possible to tell apart a dead (or privacy-enhanced) tag from a reflecting surface which does not modulate the incident signal, such as a short segment of wire, but it is even possible to discover the brand of a specific dead tag, simply by observing this tag's backscatter. By sweeping a directed beam with changing polarization over a person, an adversary can thus learn about the type and orientation of the various tags carried by this person, even if the tags are dead and cannot be interrogated. This calls into question the entire concept of application-layer privacy and gives credence to the opinion that only physically destroying a tag can truly silence it[**23**].

CHAPTER 5

# Discussion

This section will discuss the practical implications of our results, ways to avoid their detrimental effects and some future directions for research.

## 5.1. Practical Implications

As stated before, countermeasures against power analysis come at a price – they usually increase the cost and decrease the range of the tag. In addition, implementing any type of countermeasure requires vendors to commit resources to its design and its testing. Unless the attack described in this work is widely publicised and reproduced, there is little chance of vendors making this effort. The pressure on vendors to reduce the cost and increase the range of their devices forces them to compete for the lowest gate count. Any non-essential functionality of the device, such as the ability to withstand an esoteric attack on an old generation of tags, will be under fierce scrutiny to justify its existence.

This means that for all practical purposes, devices produced in the next year or two should be considered insecure. Tag integrators concerned about the security of their implementations may be forced to resort to other measures to prevent their tags from being killed or rewritten by rogue attackers. Fortunately, the EPC protocol specifies a stopgap measure called *permalock* that can protect today's systems. As shown in [**19**, subsections 6.3.2.9 and 6.3.2.10.3.5], tags can be programmed in such a way that they may never be killed or rewritten, regardless of the password supplied to them. Security-sensitive applications should make sure the kill command is disabled by use of this permalock mechanism. Obviously, this compromise does away with the privacy protecting capabilities of the tag – a permalocked tag can be killed only by physically destroying it. It should be noted that given the industry's track record in properly implementing security commands, integrators should always test to see whether the permalock feature works as designed in the particular brand of tags they use.

## 5.2. Countermeasures

This work concentrates on attacks rather than on defences. Nevertheless, we will review some common countermeasures and explain why they are problematic to implement by RFID chip designers. The interested reader is invited to look at the introduction to [**48**] or at [**38**] for a more detailed survey.

**5.2.1. Mitigation and Prevention Countermeasures.** In general, power analysis countermeasures fall into one of two categories: mitigation and prevention. Mitigation countermeasures try to reduce the signal-to-noise ratio of the secret information located in the power consumption trace, either by attenuating or by hiding it in noise. Prevention countermeasures try to completely remove secret information from the trace.

A common type of mitigation countermeasure involves the addition of *random noise* to the power consumption of a device[**29**]. Since power is supplied to tags by the reader, it sounds tempting to add this noise source to the reader's signal and not to the tag, thus saving a painful redesign of the tags and keeping their costs low. However, there are several reasons why this approach won't work. First a differential setup such as the one described in Subsection 3.2.4 will not be affected by this added noise. Second, the reader can only add very limited narrowband noise to the signal because of the strict regulatory constraints placed on its high-powered output.

An example of a prevention countermeasure is the introduction of *balanced logic* – a method of designing the circuit such that the same number of outputs switches between states every clock cycle [**30**]. The unintuitiveness of this requirement can be eased by using prefabricated hardware description language (HDL) components which encapsulate this behaviour (see for example [**49**]). The main drawback of this approach is in the price circuit designers have to pay – the added gate count raises the cost of every device, while the larger amount of transitions per clock cycle translates immediately into a higher power consumption and thus a lower read range. It may be tempting to isolate the circuit into secure and non-secure components and apply logic balancing only to the secure components. However, care must be taken when deciding which parts are secure and which are not. For example, a chip designer may try and protect against the password-sniffing attack by balancing only the one-bit register containing the result of the comparison of the current password bit and the received bit. However, if the tag's data bus is not balanced, it is still possible to detect individual bytes of the password as they are read from memory and learn about their Hamming weights.

**5.2.2. Double-Buffered Power Supply.** A feasible solution, which is perhaps the most compatible with the current RF front ends found on tags, would be the separation of the power supply from the power consumption by use of a *double-buffering power supply* mechanism as described in [**48**]. As shown in Figure 5.1 on page 63, this mechanism consists of a pair of capacitors switched by power

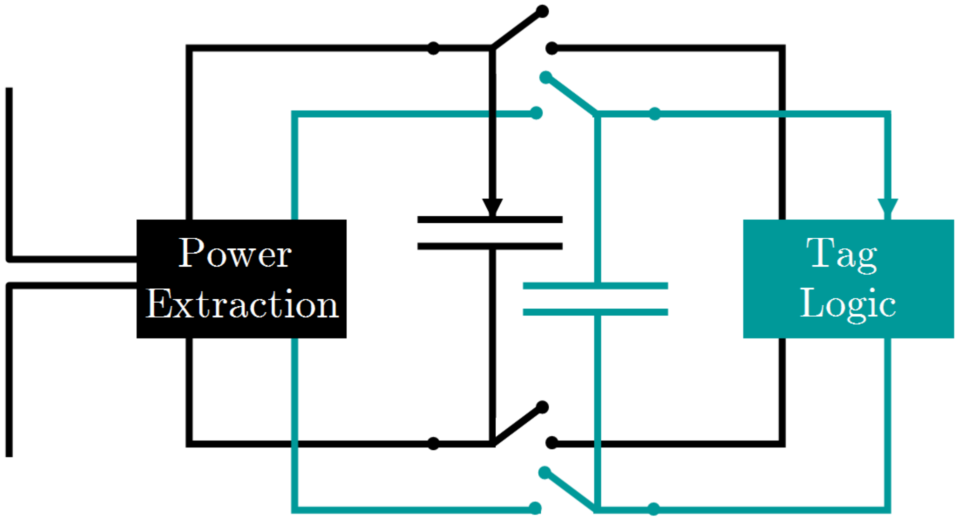


FIGURE 5.1. The double buffered power supply

transistors. At any stage in time, one capacitor is charged by the reader while the other is being discharged by the circuit. With proper design, this approach can almost eliminate the power consumption information. Moreover, it involves changes only to the RF front-end of the tag, making it the quickest to roll out. To make this countermeasure more effective, large flat capacitors can be attached to the plastic inlay next to the printed antenna. Tag vendors can easily produce two versions of their ICs – a protected version for secure applications and an insecure version for cost-conscious applications – while sharing the internal logic and only dropping in different RF front ends. To further reduce costs, vendors can create a single IC with redundant contact points. Such an IC will offer power analysis resistance when fixed to inlays with the extra capacitor and degrade to insecure operation when fixed to inlays without such a capacitor. Tags using this protective mechanism still have to take care that power consumption does not leak out through the intentional backscatter modulation mechanism, which has to come out of the circuit proper and connect to the antenna.

## 5.3. Improving the Current Attack

**5.3.1. Increasing the Sensitivity.** One issue involved in bringing this attack from the lab out into the open revolves around the issue of its maximal range. Since the parasitic tag modulation uses a similar mechanism to the tag's intentional backscatter, the maximum read range of a tag is an approximate upper bound on the range of our attack.

As of early 2006, tag vendors claim to have a read range of around 8 meters in ideal conditions. Tags have been shown to be detectable using a single antenna from as far away as 21 meters [**44**]. The attackers of [**44**] used a high-gain directional Yagi antenna to focus only on the signal reflected by the tag and ignore noise coming from other sources. Since the backscatter mechanism used for intentional tag responses is the same mechanism we use for our attack, this can be assumed as a reasonable bound on the range of our attack. However, one must remember that an increase in the directionality of a Yagi antenna usually comes with a reduction of its usable frequency bandwidth. The bandwidth of an intentional tag backscatter is on the order of 50 KHz, whereas our attack used a resolution bandwidth of 5 MHz.

The authors of [**31**] were faced with a similar detection scenario when trying to launch TEMPEST attacks against video displays. In [**31**, p. 95] they suggest using an array of Yagi antennas in a grid pattern to maximize directionality while keeping the usable bandwidth of the antenna array over 10% of the center frequency. According to their calculations, a 24-antenna $(4 \times 6)$ array of 4-element Yagi antennas tuned to 900 MHz will fit inside a 1m $\times$ 1m $\times$ 33cm panel and have an impressive directional gain of over 22dB while maintaining a signal bandwidth of 90MHz.

**5.3.2. Lowering the Cost.** Our attack was performed using lab equipment under lease. The rental cost of this equipment was under $1000 per month. There are several ways of reducing the cost even further.

The most expensive piece of equipment we used was the HP4405B-AYX spectrum analyzer. There are several ways of replacing this device with a lower-priced counterpart. First, TV and VCR devices contain a sensitive RF receiver. As discussed in Subsection 4.1.1.1, we can use this tuner box to construct a homebrew spectrum analyzer[**41**]. If we can keep the RF environment in which the tag is operating reasonably clean, one can also assume that the only signal of non-negligible amplitude incident on the attacker's antenna is the tag signal itself. This can be pretty close to the truth if the attacker uses a high-gain directional antenna pointed at the tag and having a narrow reception bandwidth centered around the reader's operating frequency. In this case, a trivial AM detector consisting of a Schottky diode and a lowpass filter can be used as a cheap replacement for the spectrum analyzer.

An interesting alternative for the entire lab setup is a cellphone with modified firmware. UHF tags and GSM cell phones have very similar operating frequencies. The cell phone's antenna has the right shape for talking to RFID tags. The transmitter is strong enough. The receiver is more than sensitive enough. The air interface protocol of modern cell phones is much more complicated than the RFID air interface . This means that with the appropriate firmware a cellphone can be modified to attack UHF tags. HF tags have different frequencies and antennas, but more and more vendors are adding HF reader circuitry to their phones. With some firmware modifications. these so-called "wallet phones" will be capable of attacking HF tags. We have not tested these ideas, but we see no insurmountable technical difficulties in carrying them out.

## 5.4. New Directions of Attack

**5.4.1. Attacking HF Tags.** The discussion so far was focused on UHF (EPC) tags, which operate in the 900 MHz frequency range. These tags have a higher read range which makes them easier to attack. Another common type of tag is HF tag (ISO/IEC 14443), which has an operating frequency of 13.56 MHz. These tags rely on slightly different principles to provide power to the tag. As opposed to the radiative form of tag-reader connection used in UHF tags, HF tags use inductive coupling. This form of coupling has a wavelength-related maximum read range imposed by the physics of the inductive coupling method. For standard ISO/IEC 14443 tags, this range is about 3.6 meters [**55**, p. 43], although the regulatory restrictions on the reader's output power typically cause the maximum read range to be much lower. In practical situations, the range of near-field tags is on the order of 10cm, and in nearly all cases no more than 50cm[**27**].

The theoretical power-analysis attack shown in Figure 5.2 on page 66 is designed to work in this short-range scenario. It contains a sensitive ammeter connected to a coil antenna which is to be sandwiched between the tag and the reader. Because of inductive coupling, current flowing through the tag's coil antenna will cause a proportional current to flow through the adversary's antenna and to be picked up by the ammeter. This signal will be superimposed on the powerful current caused by the reader's signal. By measuring the minute changes in the current flowing through the ammeter, the attacker can estimate the current flowing through the tag and, through that, measure its power consumption.

Many HF readers are located in semi-public places such as entranceways and subway stations. Considering the flat form factor of coil antennas, an adversary can design his attacking device in the form of a sticker containing a coil antenna, a basic signal-processing front end and some non-volatile memory. The reader under attack will even be generous enough to provide the attacker with a free power supply. The attacker can paste this device on top of the reader's surface, let it capture raw data about a day's worth of tags, then peel off the attack device and take it back to his secret hideout for offline processing. Given the proliferation

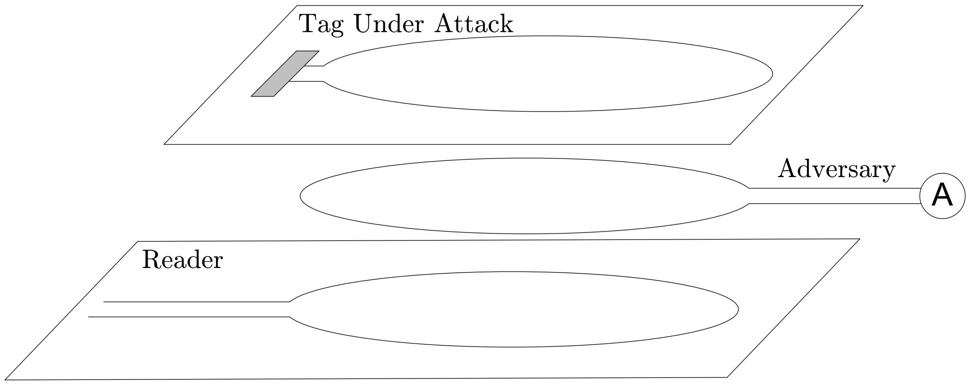


FIGURE 5.2. A theoretical remote power analysis attack on HF tags

of stickers already found in retail and public transportation locations, this extra sticker will be reasonably hard to detect[1].

**5.4.2. A Smart Fault Attack Based on Jamming.** The work of [**17**] shows how simple it is for a rogue device to corrupt the operation of readers in its vicinity. If we somewhat refine this result we can adapt it to form a smart fault attack. We do this by noting that a tag receiving invalid data will probably transition to some error state or even reset itself. By carefully controlling the point in time in which a reader's signal is jammed, the attacker can preempt another reader's incoming signal right in the middle of a calculation and force it to transition to the error state at any desired location in time.

Since forcing a device into the error state changes its registers to some predefined state, the power consumed by a tag when it is suddenly jammed is related to the Hamming weight of its internal registers. By forcing a reset by smart jamming and then measuring the power consumed by the tag the attacker can perform remote fault analysis of the tag and learn about its internal state, again without physical manipulation. This attack will be much more effective if the attacker can control the reader as well as the jamming device.

A more esoteric fault attack can be based on destructive interference. If the reader is in a fixed location and the attacker has time to prepare, he can create a device that generates "null zones" on demand. This fault-generating circuit, presented in Figure 5.4.2 on page 67, consists of a receive antenna, a delay line and amplifier, and finally a properly polarized transmit antenna. The device receives the reader's signal, changes its relative phase using the delay line and sends it back

[1]It is important to note that HF tags have more generous power budgets and, as such, are more easily modified to include power analysis countermeasures. This is especially the case with inductively coupled contactless smart cards, which typically inherit the power analysis countermeasures from their physically coupled siblings.

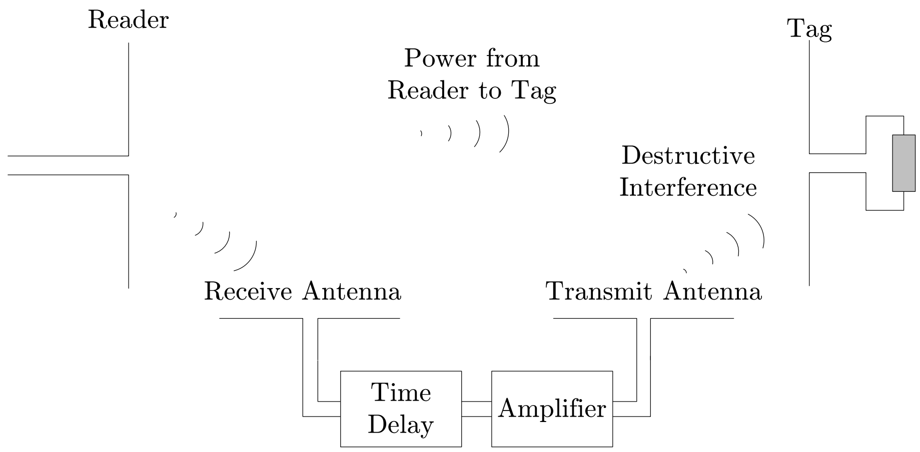


FIGURE 5.3. A theoretical setup for creating destructive interference, to be used for remote fault analysis

in the direction of the tag. By configuring the phase and the amplitude to match and oppose the reader's signal, the signal presented to the tag can be decreased arbitrarily. We noticed in our tests that we could see more pronounced power consumption information when tags were placed in a relatively weak field, so weakening the reader using this method is one practical way of remotely raising the signal to noise ratio of the attacker.

CHAPTER 6

# Closing Remarks

In this work, we have shown how power analysis, a well known cryptanalytic attack method which was extensively researched in the field of smartcards, can be applied to the field of passive RFID devices, even though they have no explicit connection to an external power source.

The big players in the field of smart cards aggressively promote their brand names. Smartcard vendors compete on the security of their products, and smart card customers (both end-users and system integrators) realize the benifits of a higher-security device and understand why they should pay more for it. RFID tags, on the other hand, are ubiquitous devices which play a very minor role in the life cycle of a product. RFID system integrators do not generally care which specific brand of tag they use, leaving tag vendors to compete mainly on price. This fact is exacerbated by the limited feature set and strict standardisation of tags, which leaves vendors little room for innovation.

The specifics of the RFID market mean that it will be difficult to convince tag vendors to offer a higher-security tag at a higher price and with reduced range, unless the advantage provided by this added security is clearly understood. Tag vendors have all the tools they need to create safer tags today, but they will only be created if the power analysis threat is well understood by system integrators, driving a demand for these safer tags. Achieving this requires the active participation of the research community, and there are many exciting directions in which these results can be extended.

# Bibliography


[1] Manfred Aigner and Elisabeth Oswald. Power analysis tutorial. Online, December 2000. http://www.iaik.tugraz.at/aboutus/people/oswald/papers/dpa_tutorial.pdf.

[2] Henri Barthel. UHF RFID regulations. Online, June 2006. http://www.epcglobalinc.org/standards_technology/uhf_rfid.html.

[3] Matt Blaze. Safecracking for the computer scientist. Online, December 2004. http://www.crypto.com/papers/safelocks.pdf.

[4] David Brumley and Dan Boneh. Remote timing attacks are practical. *Computer Networks*, 48, 2005. http://crypto.stanford.edu/~dabo/papers/ssl-timing.pdf.

[5] Auto-ID Center. 860MHz–930MHz class I radio frequency identification tag radio frequency & logical communication interface specification candidate recommendation, version 1.0.1. Online, November 2002. http://www.epcglobalinc.org/standards_technology/Secure/v1.0/UHF-class1.pdf.

[6] Anantha P. Chandrakasan, Samuel Sheng, and Robert W. Brodersen. Low-power CMOS digital design. *IEEE Journal of Solid-State Circuits*, 27(4):473–484, April 1992. http://ieeexplore.ieee.org/search/wrapper.jsp?arnumber=126534.

[7] LeCroy Corporation. Lecroy 9300C series digital oscilloscopes. Online, January 1998. http://www.lecroy.com/tm/library/manuals/9300Series/OperatorsManual/9300_OM_REVA.pdf.

[8] Verichip Corporation. Veriguard security suite. Online, 2006. http://www.verichipcorp.com/images/VeriGuard(web).pdf.

[9] Raghu Das. RFID tag sales in 2005 – how many and where. Online, December 2005. http://www.idtechex.com/products/en/articles/00000398.asp.

[10] Tim Dawson. Using .NET languages to make your application scriptable. Online, May 2003. http://www.divil.co.uk/net/articles/plugins/scripting.asp.

[11] Daniel Dobkin. The RF in RFID. Online, October 2005. http://www.enigmatic-consulting.com/Communications_articles/RFID/RF_in_RFID_index.html.

[12] Klaus Finkenzeller. *RFID Handbook : Fundamentals and Applications in Contactless Smart Cards and Identification*. John Wiley & Sons, 2003.

[13] International Organization for Standardization. ISO/IEC 14443-2 – identification cards – contactless integrated circuit(s) cards – proximity cards – part 2: Radio frequency power and signal interface, March 1999.

[14] Jovan Dj. Golić and Christophe Tymen. Multiplicative masking and power analysis of AES. In Burt Kaliski and Çetin Kaya Koç, editors, *Cryptographic Hardware and Embedded Systems - CHES 2002: 4th International Workshop, Lecture Notes in Computer Science*, volume 2523, pages 198–212. Springer-Verlag GmbH, August 2002. http://www.springerlink.com/openurl.asp?genre=article&issn=0302-9743&volume=2523&spage=198.

[15] Louis Goubin and Jacques Patarin. DES and differential power analysis (the "duplication" method). In Çetin Kaya Koç and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 1999: 1st International Workshop, Lecture Notes*

*in Computer Science*, volume 1717, pages 158–172. Springer-Verlag GmbH, July 1999. http://link.springer.de/link/service/series/0558/bibs/1717/17170158.htm.

[16] PCTEL Antenna Products Group. MaxRad 800/900 MHz directional panel antenna series. Online, May 2005. http://205.234.153.66/images_catalog_group/pdf_docs/PCTEL_EXW05.pdf.

[17] SCISSEC Security Research Group. RFID vulnerabilities uncovered in new UHF tags. Online, April 2006. http://scissec.scis.ecu.edu.au/wordpress/?p=39.

[18] EPCglobal Inc. The EPCglobal network: Overview of design, benefits, & security. Online, September 2004. http://www.epcglobalinc.org/news/EPCglobal_Network_Overview_10072004.pdf.

[19] EPCglobal Inc. EPC radio-frequency identity protocols class-1 generation-2 UHF RFID protocol for communications at 860 MHz – 960 MHz, version 1.0.9. Online, September 2005. http://www.epcglobalinc.org/standards_technology/EPCglobal2UHFRFIDProtocolV109122005.pdf.

[20] WJ Communications Inc. MPR series PC card UHF generation 2 class 1 application programmer's interface. Online, February 2006. http://www.wj.com/Support/ApplicationNotes.cfm.

[21] Dynamic Sciences International. R-110B wide range AM/FM receiver. Online, August 2003. http://www.dynamicsciences.com/R110.pdf.

[22] Ari Juels, Ronald L. Rivest, and Michael Szydlo. The blocker tag: selective blocking of rfid tags for consumer privacy. In *CCS '03: Proceedings of the 10th ACM conference on Computer and communications security*, pages 103–111, New York, NY, USA, 2003. ACM Press. http://doi.acm.org/10.1145/948109.948126.

[23] Günter Karjoth and Paul Moskowitz. Disabling RFID tags with visible confirmation: Clipped tags are silenced. In *Workshop on Privacy in the Electronic Society – WPES*, Alexandria, Virginia, USA, November 2005. ACM, ACM Press. http://domino.watson.ibm.com/library/cyberdig.nsf/papers/D25E54DB29DAA9AA8525707C00702C9F/$File/rc23710.pdf.

[24] Udo Karthus and Martin Fischer. Fully integrated passive UHF RFID transponder IC with 16.7-$\mu$W minimum RF input power. *IEEE Journal of Solid-state Circuits*, 39(10):1602–1608, October 2003. http://www.ee.washington.edu/research/seal/internal/files/RFID_Karthaus_Fischer.pdf.

[25] Matthai Philipose Kenneth P. Fishkin, Bing Jiang and Sumit Roy. I sense a disturbance in the force: Unobtrusive detection of interactions with RFID-tagged objects. In Itiro Siio Nigel Davies, Elizabeth Mynatt, editor, *UbiComp 2004: Ubiquitous Computing, 6th International Conference, Lecture Notes in Computer Science*, volume 3205, pages 268–282. Springer-Verlag GmbH, January 2004. http://www.intel-research.net/Publications/Seattle/062420041544_244.pdf.

[26] Ziv Kfir and Avishai Wool. Picking virtual pockets using relay attacks on contactless smartcard systems. Cryptology ePrint Archive, Report 2005/052, 2005. http://eprint.iacr.org/2005/052.

[27] Ilan Kirschenbaum and Avishai Wool. How to build a low-cost, extended-range RFID skimmer. Cryptology ePrint Archive, Report 2006/054, 2006. http://eprint.iacr.org/2006/054.

[28] Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. *Lecture Notes in Computer Science*, 1666:388–397, 1999. http://www.cryptography.com/resources/whitepapers/DPA.pdf.

[29] Paul Kocher, Joshua Jaffe, and Benjamin Jun. US patent 6,327,661: Using unpredictable information to minimize leakage from smartcards and other cryptosystems, 2001. http://www.cryptography.com/technology/dpa/Patent6327661.pdf.

[30] Paul Kocher, Joshua Jaffe, and Benjamin Jun. US patent 6,510,518: Balanced cryptographic computational method and apparatus for leak minimizational in smartcards and other cryptosystems, 2003. http://www.cryptography.com/technology/dpa/Patent6510518.pdf.

[31] Markus G. Kuhn. Compromising emanations: eavesdropping risks of computer displays. Technical Report 577, University of Cambridge Computer Laboratory, December 2003. http://www.cl.cam.ac.uk/TechReports/UCAM-CL-TR-577.pdf.

[32] Butler W. Lampson. Hints for computer system design. *Operating Systems Review*, 15(5):33–48, October 1983. http://research.microsoft.com/~lampson/33-Hints/WebPage.html.

[33] Roy W. Lewallen. EZNEC demo v. 4.0.28. Online, August 2006. http://www.eznec.com/.

[34] Panasonic Limited. Operating instructions for NV-7200. Online, October 1982. http://www.panasonic.co.uk/customer-Support/download-centre.asp?did=128753&fmt=pdf.

[35] RCS Maintainers. Official RCS homepage. Online, August 2005. http://www.cs.purdue.edu/homes/trinkle/RCS/.

[36] Stefan Mangard. Exploiting radiated emissions – EM attacks on cryptographic ICs. In *Proceedings of Austrochip 2003*, 2003. http://www.iaik.tu-graz.ac.at/research/sca-lab/publications/pdf/Mangard2003ExploitingRadiatedEmissions.pdf.

[37] Thomas S. Messerges, Ezzy A. Dabbish, and Robert H. Sloan. Power analysis attacks of modular exponentiation in smartcards. In Çetin Kaya Koç and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 1999: 1st International Workshop, Lecture Notes in Computer Science*, volume 1717, pages 158–172. Springer-Verlag GmbH, July 1999. http://www.springerlink.com/openurl.asp?genre=article&issn=0302-9743&volume=1717&spage=144.

[38] Thomas S. Messerges, Ezzy A. Dabbish, and Robert H. Sloan. Examining smart-card security under the threat of power analysis attacks. *IEEE Transactions on Computers*, 51(5):541–552, May 2002. http://dx.doi.org/10.1109/TC.2002.1004593.

[39] Tim "Minime" and Christopher "Mahajivana". RFID zapper. 22nd Chaos Communication Congress, December 2005. https://events.ccc.de/congress/2005/wiki/RFID-Zapper(EN).

[40] U.S. Government Printing Office. Code of federal regulations, title 47 – telecommunication, chapter I – federal communications commission, part 15 – radio frequency devices. Online, November 2005. http://www.access.gpo.gov/nara/cfr/waisidx_05/47cfr15_05.html.

[41] Green Bay Professional Packet Radio. Homebrew RF test equipment and software. Online, 2001. http://www.qsl.net/n9zia/wireless/appendixF.html#9.

[42] Patrick Rakers, Larry Connell, Tim Collins, and Dan Russell. Secure contactless smartcard ASIC with DPA protection. *IEEE Journal of Solid-State Circuits*, 36:559–565, March 2001. http://ieeexplore.ieee.org/search/wrapper.jsp?arnumber=910496.

[43] Fractus S.A. Fractus EZConnect Zigbee/RFID chip antenna. Online, May 2006. http://www.fractus.com/868-915mhz.htm.

[44] Greg Sandoval. Hackers' prowess on display at Defcon conference. *MIT Technology Review*, August 2005. http://www.technologyreview.com/articles/05/08/ap/ap_080405.asp.

[45] Sanjay Sarma and Daniel W. Engels. On the future of RFID tags and protocols. Technical Report MIT-AUTOID-TR018, Auto-ID Center, June 2003. http://www.autoidlabs.org/single-view/dir/article/6/176/page.html.

[46] Sanjay E. Sarma, Stephen A. Weis, and Daniel W. Engels. Security and privacy aspects of low-cost radio frequency identification systems. In *First International Conference on Security in Pervasive Computing*, 2003. http://theory.lcs.mit.edu/~sweis/pdfs/spc-rfid.pdf.

[47] J. Schroet. Gen 2 tag clock rate – what you need to know. Online, October 2005. http://www.impinj.com/page.cfm?ID=Document_Center.

[48] Adi Shamir. US patent 6,507,913: Protecting smart cards from power analysis with detachable power supplies, 2003. http://patft1.uspto.gov/netacgi/nph-Parser?patentnumber=6507913.

[49] Danil Sokolov, Julian Murphy, Alexander Bystrov, and Alex Yakovlev. Design and analysis of dual-rail circuits for security applications. *IEEE Transactions on Computers*, 54(4):449–460, April 2005. http://dx.doi.org/10.1109/TC.2005.61.

[50] STMicroelectronics. UHF, EPCglobal class 1b, contactless memory chip 96 bit ePC with inventory and kill function. Online, October 2005. http://www.st.com/stonline/products/literature/ds/11097/xra00.htm.

[51] Agilent Technologies. E4405B-STD ESA-E standard analyzer, 9 kHz to 13.2 GHz. Online, 2006. http://www.home.agilent.com/cgi-bin/pub/agilent/Product/cp_Product.jsp?NAV_ID=-536902958.536894443.00.

[52] International Telecommunications Union. Recommendation ITU-R BS.450-3 - transmission standards for FM sound broadcasting at VHF. Online, November 2001. http://www.itu.int/rec/R-REC-BS.450-3-200111-I.

[53] International Telecommunications Union. Recommendation ITU-R BT.1701-1 - characteristics of radiated signals of conventional analogue television systems. Online, February 2005. http://www.itu.int/rec/R-REC-BT.1701.

[54] Wim van Eck. Electromagnetic radiation from video display units: An eavesdropping risk? *Computers & Security*, 4:269–286, 1985. http://cryptome.org/emr.pdf.

[55] Roy Want. The magic of RFID – just how do those little things work anyway? *ACM Queue*, 2(7):40–48, October 2004. http://doi.acm.org/10.1145/1035594.1035619.

[56] Jonathan Westhues. Cloning a verichip. Online. http://cq.cx/verichip.pl.