

The Security Implications of VeriChipTM Cloning

John Halamka¹ and Ari Juels² and Adam Stubblefield³ and Jonathan Westhues⁴

¹ Beth Israel-Deaconess Medical Center
Boston, MA, USA

jhalamka@caregroup.harvard.edu

² RSA Laboratories, Bedford, MA, USA

ajuels@rsasecurity.com

³ Johns Hopkins University, Baltimore, MD, USA

astubble@cs.jhu.edu

⁴ jwesthues@ccq.cx

10 March 2006

Abstract. The VeriChipTM is an Radio-Frequency Identification (RFID) tag produced commercially for implantation in human beings. Its proposed uses include identification of medical patients, physical access control, contactless retail payment, and even the tracing of kidnapping victims.

As we explain, the VeriChip is vulnerable to simple *cloning* attacks. An attacker capable of scanning a VeriChip, eavesdropping on its signal, or simply learning its serial number can create a clone device whose radio appearance is indistinguishable from the original. We explore the practical security implications of this vulnerability to cloning. We argue that:

1. The VeriChip should serve exclusively for *identification*, and not *authentication* or access control.
2. Paradoxically, for bearer safety a VeriChip should be easy to clone; an attacker then has less incentive to coerce victims or extract VeriChips from victims' bodies.

Given the privacy concerns that arise from the possibility of physically tracking VeriChip bearers, however, cryptographic protection of their identifiers is desirable. We propose an alternative to the VeriChip, called the iChip, that offers resistance to tracking, but expressly permits cloning.

Keywords: medical identification, RFID, privacy, security, VeriChip

1 Introduction

The VeriChip is a commercially produced, human-implantable microchip [24]. It is designed to serve as a identification device, effectively a kind of wireless barcode or dog tag for people. About the size of a grain of rice, the VeriChip is surgically implanted under the skin of its bearer, typically on the back of the arm. When interrogated by a nearby reading device, it communicates a unique serial number over the air. This serial number may be referenced in a database to identify its bearer.

VeriChip Corporation, the manufacturer of the device, asserts that the VeriChip “cannot be lost, stolen, misplaced, or counterfeited,” and advocates a range of applications for the device [24]. In healthcare settings, the VeriChip can help identify a “Jane Doe” or “John Doe,” that is, an incapacitated or disoriented patient whose identity is difficult to establish. In private facilities, the VeriChip can enhance physical access control, as it permits automated identification of individuals and tracking of their movements in buildings. For example, the Attorney General of Mexico and members of his staff underwent surgical implantation of VeriChips as a measure to control access to a federal anti-crime information center [26]. A few years ago, a Mexican distributor announced plans to create an anti-kidnapping system for children using the VeriChip [21]. The VeriChip has also seen limited deployment as a payment device, essentially a credit-card replacement [17, 20] marketed under the product name VeriPay. It has even acquired a degree of chic among certain technophiles, who are exploring applications in daily life [5].

The VeriChip lies at the confluence of several technological trends. About fifty million house pets around the world already bear implanted wireless microchips similar in form and function to the VeriChip. These chips help shelters and veterinarians identify lost animals. For human beings, biometric authentication is becoming widespread as a tool for both physical and logical access control. Popular forms include fingerprint and iris scanning, voice identification, and face recognition. The VeriChip may be viewed as a kind of “prosthetic biometric”: like a finger, it cannot be misplaced. At the same time, the VeriChip offers a convenient digital interface and circumvents the poor reliability of natural biometrics. As a broad technology, Radio-Frequency Identification (RFID) is proliferating into many applications, including tracking of crates and pallets in industrial and military supply chains, contactless payment devices, and anti-theft systems for automobiles [6].

The spread of RFID has provoked a backlash from privacy advocates concerned about the increasing presence of tags in the possession of consumers. Because RFID tags respond silently and automatically to interrogation by readers, they permit some degree of clandestine tracking of their bearers. (Certain types of RFID tags also convey information about the types of items they are attached to, e.g., medications, and can thus facilitate invasive inventorying of personal items.) As a permanent and everpresent device, the VeriChip has proven a lightning rod for RFID privacy concerns, particularly since its approval for human implantation in 2002 by the United States Food and Drug Administration (FDA) [22]. Religious groups have gone so far as to claim that the VeriChip may be a realization of the Mark of the Beast as described in the New Testament [2].

Basic RFID tags like the VeriChip are *passive*. They do not contain an internal source of power, but instead receive transmission power from an interrogating reader. As such, they have short read ranges. Some tags can be scanned at distances up to tens of feet. Under ordinary circumstances, the effective read range of the VeriChip is on the order of several inches. As we discuss, however, an attacker can potentially capture VeriChip signals from a longer range. The short read range of the VeriChip diminishes but does not negate privacy concerns.

Privacy is not the only concern that the VeriChip raises. As we explain in this paper, the VeriChip is vulnerable to a straightforward *cloning* attack. By this we mean that an attacker that scans a VeriChip — or eavesdrops while it is scanned — can program a separate device to emit an undistinguishable simulation of the VeriChip signal that appears valid at all future times. Such an attacker can then easily spoof a reader into accepting the clone device as the target VeriChip. In fact, in principle an attacker can simulate a VeriChip on the basis of its serial number alone.

For most security applications, the claim by VeriChip Corporation that the VeriChip “cannot be ... counterfeited” is effectively untrue.

In this paper we consider the implications of cloning to the actual and envisioned applications of the VeriChip. We argue that use of the VeriChip for authentication, i.e., as a *proof* of identity, is inappropriate and dangerous. As a security device, the VeriChip exposes its bearer to coercive attacks, i.e., to use of a VeriChip under duress. Worse still, an attacker may be tempted to extract the VeriChip from the body of a victim. If suitable for implantation (which it may or may not be), the VeriChip should only serve for identification, i.e., as a convenient automated label, not for security.

Even as a mere identification tool, we assert that the VeriChip should ideally be designed to address the privacy concern of clandestine tracking. Toward this end, we propose an alternative implantable RFID device that we call an *iChip*. The iChip implements a very simple cryptographic system that permits only an authorized reader to determine its identifier and thus renders arbitrary surveillance infeasible. While cryptographic in nature, however, the iChip system *does not resist cloning attacks*. In fact, it intentionally renders cloning attacks easy, thereby negating the value of physical attacks.

In addition to their implantable product, VeriChip Corporation sells RFID tags for human identification that are wearable (and detachable), as well as theft-prevention tags for physical assets. In this paper we focus mainly on the implantable VeriChip. We have not examined the security features of other VeriChip RFID devices. Nonetheless, some of our observations regarding cloning may apply to such devices.

1.1 Organization

The rest of this paper is organized as follows. In section 2, we provide an overview of the VeriChip. In section 3 we describe several healthcare applications that motivate the use of implantable RFID tags like the VeriChip. We present the results of our efforts at cloning VeriChips in section 4. We describe our proposal for the clonable, privacy-enhancing iChip in section 5, and then conclude in section 6. The appendix contains VeriChip Corporation instructions for use of their VeriMedTM patient identification system.

2 Overview of the VeriChip

The VeriChip is an RFID tag. It operates at 134 kHz: when the tag is excited by a sufficiently strong magnetic field at that frequency, the circuitry on the chip powers up and transmits a unique identifier over the air. Communication is unidirectional, from the tag to the reader. The tag does not receive any acknowledgment from the reader that its ID has been successfully received. It therefore transmits its ID repeatedly, whenever it is powered. In this sense it is identical in concept to most of the ‘first generation’ RFID tags and proximity cards (for example, Indala’s FlexPass, or HID’s Prox Card II). The VeriChip differs from tags that communicate bidirectionally, like ExxonMobil Speedpass, which executes a challenge-response protocol, or the widely used ISO 14443 tags, which accept reader input aimed at preventing radio-signal collisions among nearby tags.

The VeriChip’s ID comprises 128 bits. In theory this means that there could exist 2^{128} VeriChips, each with a unique ID. In practice there must be fewer. First, because the ID is “looped,” the reader

knows the tag's ID only up to a cyclic shift: there is no designated first or last bit in the bit stream that the VeriChip emits. It is thus necessary to assign some bits as a synchronisation marker or to resolve this ambiguity through some other coding method. Second, it is likely that some of the bits in the VeriChip emission represent a checksum or some other error-detecting or -correcting code. Due to our limited access to VeriChip devices, we have been unable to determine the exact format of the ID at present.

We present more details of the ID's structure, however, in section 4.

3 RFIDs as Identifiers in Healthcare

In this section we examine the utility of the VeriChip and human-identification RFID more generally in the healthcare industry. Healthcare, as we explain, is a particularly attractive environment for VeriChip deployment. Medical applications for the VeriChip are also particularly interesting because, in contrast to access-control scenarios, simple unauthenticated identification can be a useful goal. A VeriChip or equivalent device that provides identification but not authentication is suited to a variety of tasks.

Passive, or battery-less, RFIDs are available in two main form factors for use in tracking humans in healthcare settings. Either the chip can be implanted into the body—the VeriChip being the leading exemplar of this type—or the chip can reside in an identification wrist-band worn by patients. Both of these form factors provide significant advantages over the printed barcodes that they are designed to replace.

Unlike barcodes, RFID tags do not require line of sight reading. Hence an RFID reader can read the tags of sleeping patients or of swaddled babies in intensive care units without repositioning their bodies. Moreover, RFID tags are better suited than barcodes for a variety of environmental conditions, as they are resistant to moisture, crushing, and tearing. Unfortunately, current RFID tags are more expensive than simple printed bar codes. RFID tags may have up to a 5% failure rate during manufacturing, resulting in a potentially unreadable wrist band [18]. RFID tags are also much harder to read if any sort of metal barrier exists between the reader and the tag.

Current implantable tags emit a simple medical record identifier which can be used by a patient's physician to access the corresponding database records through an access-controlled Web-based interface. For the most part, human use has been limited, although passive RFID tags currently serve two applications at Beth Israel Deaconess medical center in Boston [10].

The Beth Israel Deaconess Emergency Department is outfitted with passive RFID scanners to read implanted chips. If an unconscious, confused, or non-responsive patient arrives for care, he or she is scanned. If an implanted RFID with a medical record identifier is present, it can be used to retrieve the patient's medical history from the medical database. The RFID identifiers in this system need not serve as definite authenticators: the first line of each medical record contains the patient's gender, age, and (generally) race, all of which can serve as quick check to ensure that the identification is correct. Additionally, each record contains the social and medical history that the patient has elected to share with clinicians, which may also help confirm the patient's identity. The instructions furnished by VeriChip Corporation for their VeriMed system, which supports scanning of implanted VeriChips in patients, may be found in the appendix or referenced at [7].

In the Beth Israel Deaconess Neonatal Intensive Care Unit (NICU), babies are outfitted with RFID wristbands. These RFID tags serve two main purposes. First, to ensure accurate matching of

mother's milk and babies, each mother's milk is tagged upon storage in NICU refrigerators. When a nurse feeds a baby, she first scans the milk, then scans the baby. A software application ensures that the right infant receives the right milk and automatically creates an audit trail. Additionally, RFID scanners are implanted in door frames to detect babies passing in and out of the NICU. In both these cases, tags serve as identifiers, not authenticators. The hospital threat model does not regard nurses (or babies) as adversaries, and physical controls restrict unauthorized access by other parties.

One can imagine several future uses of implanted RFID tags in healthcare:

Automated registration: As patients arrive for care in outpatient, inpatient, or emergency room settings, they can be scanned and automatically registered, bypassing the "clip board" which patients generally fill out with demographics, insurance and medical information. Eliminating the clip board is one of the most important problems in healthcare IT: the Secretary of Health and Human Services recently named it as one of the three most important healthcare IT goals in 2006 [1]. Implanted chips offer one potential solution for identifying patients without imperfect identifiers such as names or overloaded (and private) identifiers such as Social Security numbers.

Patient safety: Currently, blood samples are taken from patients and medications are given to patients without confirmation of patient identity. Many hospitals use a system of stickers with warnings like "Name Check" when several patients with similar names are admitted concurrently. This problem is exacerbated further if multiple patients with exactly the same name are admitted. Blood tests and medications could be easily confused between two John Smiths, causing potential medical error and patient harm. If each patient is scanned as a blood sample is drawn, the sample can be tagged with accurate patient identifiers. Similarly, scanning patients prior to the delivery of medications can eliminate errors of identification. Of course, RFID wristbands could support these same operations, but implantable tags prevent errors that might result from inaccurate wrist-banding.

Patient tracking: As patients move from location to location in the hospital, they could be scanned with door-frame scanners or hand-held devices. Patient location information would empower workflow enhancement. When a patient arrives in the operating room, the surgeon and anesthesiologist could be automatically paged. When a patient leaves the Emergency Department and goes for an X-ray in radiology, the emergency room physician could see the patient's location on a dashboard, preventing loss of time to searches for the patient.

Active RFID tags (those with a battery) are already used to track medical personnel and equipment such as patient beds. These active tags are about the size of the pager, require battery replacement every 6 months and cost \$50 each. As with many new technologies, their size is decreasing, their battery life is lengthening, and the cost per tag is dropping significantly. These active RFID transmitters are generally of two types—based on either WiFi (802.11b at 2.4 GHz) or a proprietary protocol (at 488 MHz). The advantage WiFi is that the existing hospital wireless network can read tag locations. Active RFID over WiFi can be rapidly and cost effectively deployed for uses that require room-level tag location. Proprietary systems can provide location to the level of the square meter, but do require the installation of a dedicated RFID-reader network. Beth Israel Deaconess is currently using active tags to track equipment such as ventilators, IV pumps,

and EKG devices in the emergency department. The search times for such tracked devices have dropped to nearly zero.

4 Cloning the VeriChip

We now explain our cloning experiments on the VeriChip. For these experiments we used the “proxmarkii” generalized RFID tag reader/cloner. The proxmarkii is an RFID reading and simulation device developed by Westhues, who used an earlier version to demonstrate cloning attacks against proximity cards [27, 28]. Given its design for research applications, proxmarkii is capable of dealing with a large variety of formats for the signal over the air. It is also capable of simulating any kind of low-frequency RFID tag, and thus of replaying stored VeriChip IDs to readers.

Because the VeriChip transmits its ID repeatedly, its signal is periodic. By examining its autocorrelation, we determined that the tag ID is emitted over a period of 4096 carrier clock cycles. (When the tag transmits at its nominal operating frequency of 134kHz, a carrier clock cycle lasts about $7.46\mu\text{s}$.) By looking at a graph of the signal received from the tag, we were able to determine that each bit is emitted over an interval of 32 clock cycles; this led us to determine that the full length of the ID is $4096/32 = 128$ bits. The ID appears to be transmitted using Manchester-coded Amplitude-Shift Keying (ASK).

We obtained three different VeriChip tags (two unimplanted, one implanted), giving us three IDs to study. We identified only 32 bits of the 128-bit transmitted value that appear to vary among tags. These 32 bits are separated into two 16-bit sections surrounded by bit patterns that most probably synchronize the reader. It is possible that some of the other bits in the signal also transmit ID data, but the 128-bit tag IDs we observed contained mostly 0’s. It is also likely that some bits are a checksum. Given our limited sample size, we did not make more than a first-order attempt to determine the mapping between the 128-bit string and the sixteen-digit code that the legitimate reader reports. It is possible that VeriChip Corporation has implemented cryptographic techniques to make this mapping harder to determine; we have not determined whether this is indeed the case.

Basic cloning, however, does not require a deep look into the structure of the tags’ IDs. Since the VeriChip always transmits exactly the same information, cloning a VeriChip is just a matter of determining the signal that the tag transmits and building a device that mimicks that signal. There is no need to know the meaning or encoding of the signal. It is helpful to know a little bit about the structure of that signal—whether we have read a valid signal or one corrupted by noise, for example—but not a fundamental requirement.

All of the operations described above, in which the tag is energized, and measurements are made on the signal received over the air, are identical to the operations performed by a legitimate RFID tag reader. If the specifications for the VeriChip were known, then it would be possible to perform the “read” portion of the cloning using a commercial off-the-shelf reader. We could then take the ID that that reader provides, and map it back on to a signal over the air, according to the specification. Indeed we could perform *existential* cloning, meaning that we could create a simulated VeriChip with an ID whose signal we have never actually observed.

Not knowing the mapping from reader-displayed IDs to radio signals, we employed our own reader and devised our own (arbitrary) format in which to store tag IDs for later mapping back to signals over the air.

Viewed another way, we performed *replay* attack against the VeriChip, meaning that we simply captured a signal from the target device and re-transmitted it to a reader. The complexity of our attack results only from the engineering details of the communications link over the air. Because the VeriChip emits only a static identifier, a replay attack is equivalent to full-blown cloning, i.e., the harvested signal may be replayed indefinitely while appearing valid to a reader.

Replay attacks: A VeriChip could in principle be designed to prevent replay attacks, or to render replay attacks less effective than full-blown cloning. Such design would require that the VeriChip modify its emitted ID over time and would therefore necessitate additional resources in the tag. A VeriChip that transmits unidirectionally, i.e., an output-only device, cannot prevent replay attacks. Indeed, if a tag is stateless, i.e., has no clock or storage, then its output values are subject to replay at any future time. Provided that an output harvested by an attacker has not yet been replayed to a legitimate reader, it will appear to be legitimate. On the other hand, if a unidirectional tag maintains state—either through memory or an onboard, powered clock—it can minimize the impact of replay attacks. When scanned by a legitimate reader, such a tag can transmit state information, e.g., a counter value or timestamp, that invalidates any previously harvested outputs. Tags that execute bidirectional protocols such as challenge-response algorithms can defend against replay attacks. (No logical-layer protocols, though, can protect an RFID tag against *relay* attacks, as described in, e.g., [16].)

4.1 Implications of cloning

As we see, the practicality of our cloning attack is determined not by any cryptographic factors, but simply by the read range of the reader used to clone the tag. Consequently, the VeriChip’s small size is its biggest security feature. The antenna inside the VeriChip is very small, and therefore inefficient. Only a powerful carrier can excite the tag, and the information-bearing signal that the tag returns is weak.

To achieve a longer read range, it is necessary to use a physically large read antenna, or to deliver high power to the antenna. It would be difficult to achieve a read range of more than a few inches with a portable, battery-operated reader. The execution range of a cloning attack is therefore limited, but not impractically so. For example, where the VeriChip is deployed for access control, authenticating its bearer to unlock a door, it is easy to imagine an attacker following victims from their workplaces and stealing their IDs on crowded subways.

Furthermore, an attacker can harvest a VeriChip ID via an eavesdropping attack. Rather than reading a VeriChip directly, the attacker can intercept the signal emitted by a VeriChip as it is scanned by a legitimate reader. Because the attacker does not in this case power the target VeriChip directly, eavesdropping is feasible at a considerably longer range than direct reading – possibly from some tens of feet away, as experiments with RFID-enabled passports suggest [11].⁵

If the VeriChip came to be widely used for payments, then even less specific attacks would be practical; it would be beneficial to an attacker to clone any stranger’s ID, because it would be possible to make purchases with it. An attacker could push clumsily through any sort of crowd, gathering IDs along the way, or eavesdrop near a payment-system reader.

⁵ RFID-enabled passports operate at 13.56MHz, however, and may therefore have a longer eavesdropping range than VeriChips.

The risks associated with healthcare applications are less obvious, because in that case, the VeriChip does not grant access to anything with immediate financial value. Still, an attacker who could read a patient’s VeriChip and had access to the associated database could obtain the patient’s medical records. This attack is fairly obvious, but its practicality depends greatly on external factors, relating to how access to the database is controlled.

Depending on how the VeriChip came to be used, it might be advantageous for an attacker to appear to a physician to be another person. For example, a drug addict might attempt to clone the tag of a patient with a disease treated with narcotics. This attack is complicated by the fact that the tag ID would be read not by an unattended machine, but by a physician, who would presumably notice if instead of presenting his shoulder, the patient presented a hand-held electronic device.⁶ Moreover, as mentioned above, patient records contain an abundance of information that serves to confirm a patient’s identity.

4.2 Existential cloning

In addition to the risk of cloning practiced through surreptitious scanning and replay of VeriChip signals, there is also, as mentioned above, a threat of existential cloning. The IDs in the three VeriChips we obtained appeared very likely come from a small identifier space. Setting aside what appears to be a fixed header value (‘1022’ in decimal), all three decimal IDs that we observed were integers less than 50,000. (To protect the anonymity of the owners, we do not reveal the specific ID values.) Indeed, it is conceivable that VeriChips emerge from production process that assigns sequential or otherwise non-random serial numbers to chips.

That said, we have not yet attempted to determine the mapping from sixteen-digit IDs to over-the-air signals. It is unclear how much effort would be required to do so. As explained above, we observed 32 bits whose values varied among the over-the-air signals of our three tags. Our educated guess is that 16 to 24 of these bits encode ID values while the remaining 8 to 16 bits encode a checksum of some kind, e.g., a cyclic redundancy code (CRC).⁷ If the checksum is unkeyed, i.e., if it depends on the ID alone, then we believe that with some additional work, it would be relatively easy to perform existential forgery. On the other hand, if the checksum is keyed, i.e., if it depends upon a secret key shared among VeriChip readers, then existential forgery would be more difficult. To compute the correct checksum for a given ID, an attacker would need to: (1) Extract the secret key from a reader by means of reverse engineering or tampering; (2) Determine the secret key by means of cryptanalysis; or (3) Guess random checksums and test them against a valid reader or reader component.

If an attacker can mount an existential cloning attack, the implications are serious. Consider a corporation that uses the VeriChip to control access to a secured physical area. If IDs are indeed

⁶ A clever attacker could sidestep this problem by building an “active VeriChip,” powered not by the reader signal, but a small battery. This device would have much longer range than a legitimate VeriChip. The attacker could conceal this device on his person, without implanting it. When the physician scanned the patient’s shoulder (or wherever the VeriChip was supposed to be implanted), the “active VeriChip” would report its ID. A truly determined attacker could build an implantable VeriChip clone or perhaps modify an existing VeriChip to output a false ID.

⁷ An obvious upper bound on the checksum is 30 bits – as at least two bits are required to render three IDs distinct.

assigned sequentially in production, for instance, then an attacker that observes the ID of one employee in a given corporation can probably guess the IDs of other employees, which are likely to be nearby decimal values. Thus even if the corporation discovers that the VeriChip of one of its employees has been cloned, revoking access privileges for that employee would be insufficient: the attacker could simulate other valid IDs in the system.⁸ In other words, it appears that as they are currently assigned, the IDs themselves in the VeriChip system cannot reasonably be regarded as secret.

The assignment of random VeriChip IDs over a large enough space would in principle minimize the risks of existential forgery. The possibility of existential attacks, however, illustrates yet one more potential pitfall in use of the VeriChip for authentication. Moreover, until a new system of ID assignment is created, all of those who have VeriChips implanted will be indefinitely vulnerable to any existential cloning attacks enabled by the current system—at least until they undergo corrective surgery.

In summary, given the risks of basic cloning and existential cloning, the VeriChip as designed is perhaps appropriate for *identification* of its bearers, but its vulnerability to cloning renders it inappropriate for *authentication*.

5 Clonability, Privacy, and the iChip

There are well known cryptographic tools, like challenge-response protocols, that can defend against over-the-air, logical-layer cloning attacks like the one we have demonstrated. Paradoxically, though, there is a compelling reason to ensure that an implantable RFID tag is in fact clonable: an adversary then has little incentive to perform a physical attack against the chip. As a “prosthetic biometric,” a VeriChip carries the same dangers as a real biometric, such as a fingerprint. For example, in 2005, thieves severed a man’s finger in order to steal his Mercedes, which had fingerprint-based access control [15]. An attacker has a similar incentive to obtain physical possession of a VeriChip or to coerce its bearer if the chip is: (1) Hard to clone and (2) Used to secure access to valuable resources.

For this reason, we extend our claim about the VeriChip to a larger principle. We maintain that *no matter how they are designed*, implantable RFID tags should be used only for identification, and not authentication. In most situations, sacrificing authentication functionality in a VeriChip or similar device is well worth the elimination of incentives for adversaries to mount physical attacks against bearers.

Whether or not and how an implantable RFID tag serves for access control in a given system – i.e., condition (2) above – is a matter largely beyond the control of its bearer. Thus, it seems imprudent to rely on avoidance of implantable-RFID authentication at the system level. The following example illustrates this point.

⁸ Moreover, the thorny question arises of how to re-establish access rights for compromised devices. How is a surgical implant revoked? How would an employee react to a request for chip removal and re-implantation?

Example: IronClad Bank offers a convenient system in which patrons can identify themselves to ATMs using their implantable RFID tags. In this system, authentication relies exclusively upon an iris scan and PIN, not upon the RFID tags. Several years later, however, IronClad bank implements a system in which the implantable RFIDs of its enrolled patrons control access to their safety deposit boxes.

In this example, the bank creates an incentive for physical attack against its patrons' implanted RFID tags. Not only does the bank has create this incentive unilaterally, but some its patrons may not even know of its existence. To discourage such applications — or at least protect bearers from their consequences — it is important that an implantable RFID chip should be easy to clone by design. In particular, it should not contain cryptographic protections against cloning, like challenge-response schemes.

What, then, are the implications for user privacy?

5.1 The iChip

At first glance, it may seem that privacy — in the sense of protection from clandestine tracking — cannot co-exist with clonability. Certainly, an RFID tag that emits a static identifier, like the VeriChips in use today, does not afford privacy. We now describe our iChip proposal. It employs a very simple scheme that achieves both privacy and clonability simultaneously by outputting randomized encrypted values. Our exposition assumes a basic familiarity with public-key encryption; for an introduction, see, e.g., [23].

Our scheme is as follows. Let (SK, PK) represent a private/public key pair for valid reader in the system. Each iChip stores its unique serial number s in encrypted form, that is, as a ciphertext C under PK , and also stores the public key PK itself. When queried by a reader, the iChip randomly *re-encrypts* C under PK to yield a fresh ciphertext C' , which it then outputs. The iChip need not store C' , and can therefore be stateless.

For this scheme, we require the use of a *homomorphic* public-key cryptosystem, whose such as the El Gamal cryptosystem [8]. A homomorphic cryptosystem possesses us two essential properties:

1. *Re-encryption:* With knowledge of the public key PK alone, any entity can re-encrypt a ciphertext C to yield new ciphertext C' . Both C and C' decrypt under secret key SK to the same plaintext s . Furthermore, C' itself can be re-encrypted.
2. *Untraceability:* The ciphertext C and re-encryption C' are unlinkable without knowledge of SK . More precisely, given a ciphertext C and public key PK , it is infeasible to determine whether or not a second value C' represents a valid re-encryption of C . (In other words, it is infeasible to tell whether C' represents a re-encryption of C or some other ciphertext — or if it is just a random value.)

The property of re-encryption ensures that an attacker can readily clone an iChip. To clone a target iChip, the attacker merely has to obtain an output ciphertext (re-encryption) C' . She can then create a clone that computes and outputs re-encryptions of C' . We do not provide details here. For any natural implementation of our idea, though, the output values of the clone will have a statistical distribution identical with that of the target iChip. In other words, the clone will be

perfectly indistinguishable from the original – even to a reader with knowledge of the private key SK .

The untraceability property of the cryptosystem ensures that despite being able to clone iChips, an attacker without knowledge of or access to SK cannot feasibly correlate the outputs of iChips in the system. In other words, an attacker that scans iChip clandestinely still cannot track them, as illustrated in the following example.

Example: ExCon, a retailer, would like to improve its customer service. When high-spending customers enter its shops, ExCon would like to identify them automatically so that they can be approached by its sales staff – and would like conversely to disregard low spenders. By embedding reader bays in the impulse-item stands, ExCon’s plan is to scan the iChips of patrons clandestinely as they pay for items at checkout counters and then link their iChip outputs with their expenditures in a database.

The ExCon system described in this example is not workable. It is infeasible for ExCon to determine an association between the ciphertext C'' output by a patron’s iChip when she enters a shop and the corresponding ciphertext C' registered in the ExCon database.

On the other hand, an attacker with access to a privileged system that leaks information about iChip serial numbers could perform clandestine scanning of tags. For example, if ExCon could penetrate a medical system that decrypts ciphertexts, looks up identifiers a database, and outputs the name of iChip owners, then its plan would succeed.

Strong privacy, therefore, demands that a system be designed such that readers do not promiscuously reveal information about iChip owners or serial numbers. For example, in a medical setting, a technician who scans the iChip of a patient should be required to authenticate herself in order to access the record of the patient. (Of course, an attacker can still potentially learn s through compromise of a reader, social engineering, etc. But such measures minimize the risk of information leakage of this kind.) In brief, our system helps alleviate the threat of tracking, but cannot eliminate it.

It is to be expected that multiple iChip domains may co-exist in the same environment. Each domain D_i would naturally have its own corresponding key pair (PK_i, SK_i) ; each iChip would emit a ciphertext under the public key for its associated domain (or possibly for multiple ones). In order for an iChip to be clonable, it would seem at first glance that it would have to reveal either its corresponding public key PK_j . If a iChip emits PK_j , however, then it betrays privacy-sensitive information – namely its corresponding domain.

The *universal re-encryption* scheme of Golle et al. [9] furnishes a solution to this conundrum. That scheme extends ciphertexts in such a way that they may be re-encrypted *without knowledge of their corresponding public keys*. In other words, the Golle et al. system eliminates the need for an iChip to reveal domain-specific information. A system proposed by Ateniese et al. [3] permits a further extension of functionality by enabling a reader to verify that a ciphertext has been digitally signed by an authorized party – with no sacrifice of privacy.

Public-key cryptography is well beyond the computational resources of most RFID devices. It would probably raise the cost of an iChip considerably above that of a VeriChip, and very likely reduce the effective read range. Given that a iChip is a personal medical device, we believe that

higher cost would not significantly impair its commercial viability. Reduced range might actually be beneficial to privacy.

Related work: Several papers have suggested the use of public-key-based re-encryption for RFID tags, e.g., [3, 9, 12]. In these cases, however, the aim was for an external device, i.e., a reader, to perform the re-encryption on behalf of the RFID tag. The advantage to such external re-encryption is that it can support privacy in very lightweight RFID tags, namely ones that do not themselves perform cryptographic operations. For a broad survey of security and privacy in RFID, see [4, 14].

Our idea here of intentionally enabling cloning is similar in spirit to previous cryptographic systems that explicitly suppress security features in the interest of privacy. Group signatures (and a lightweight variant called ring signatures) are perhaps the best examples [25]. A group signature is one in which any of a pre-established group of players can apply a valid digital signature to a document. An ordinary party that verifies a signature, however, cannot determine exactly *which* player signed the document. In other words, to support privacy, a group-signature scheme essentially suppresses the usual security property of individual accountability or repudiation present in traditional digital signature systems.

5.2 Private-key compromise

A serious problem arises if the private key SK in an iChip system is compromised, i.e., obtained by attackers. Short of surgical removal, there is no obvious mechanism for deactivating iChips. And indeed, even surgical removal can be complicated: because no good mechanism exists today for a surgeon to pinpoint the exact location of an implanted VeriChip, removal reportedly requires physical exploration under the skin [19]. There are a few possible approaches to incorporating mechanisms to address key compromise:

1. **Self-deactivation:** iChips might be constructed to disable itself permanently upon receipt of an authenticated “kill” command from a reader. “Wireless barcode” RFID tags, known as EPC (Electronic Product Code) tags, include a “kill” function aimed at privacy protection.
2. **Re-keying:** iChips might be constructed so as to permit authenticated re-keying, i.e., replacement of their programmed ciphertexts and public keys.
3. **Blocking:** If iChips do not include either of the above mechanisms, and an iChip bearer wishes to deactivate her implant without the risk of surgical removal, she might instead have a blocker tag [13] implanted.⁹ A blocker tag is itself a passive RFID device that can interfere with reading activity in its vicinity. Although initially designed to interfere with anti-collision protocols, it could be easily modified for this setting.

Note that in some sense, the problem of deactivation exists even for implanted RFID tags that do not implement our privacy-enhancing system, e.g., in today’s generation of VeriChips. Some bearers of VeriChips may simply come to the decision that they no longer wish to assume the privacy risks or other side-effects associated with their implants. Indeed, if a VeriChip bearer discovers that someone has cloned her device for abusive purposes, a good line of recourse is essential.

⁹ One can, however, imagine a chain of untoward consequences, as in the song about the old lady who swallowed the fly.

5.3 Subliminal channels in the iChip

The iChip permits easy cloning, but only if it is implemented as proposed. In principle, a variant protocol is possible in which a iChip communicates authentication information to a reader via a *subliminal channel*, i.e., one that is undetectable by an attacker. Such a channel could lead an attacker to construct a clone whose re-encryptions are identifiable by a reader as falsified. Even a stateless iChip can implement such a channel: the iChip merely needs to emit ciphertexts restricted to a subset R^* of the full space R of possible re-encryptions.¹⁰ Provided that a reader can recognize ciphertexts in R^* , it can distinguish between the re-encryptions of the original iChip and those of a clone. We now give an example of a simple channel of this kind.

Example: Suppose an iChip shares a secret key x with a legitimate reader. Let $X[a]$ for $0 \leq a \leq n$ denote the a^{th} bit of an n -bit string X and let $X[a, b]$ for $a < b$ denote the sequence of bits from the a^{th} to the b^{th} . The iChip might only emit values C' such that $C'[k + 1] = h(x, C'[1, k])[0]$. In other words, the iChip would emit ciphertexts that contain a cryptographic *check-bit* based on x . The output of such an iChip would be indistinguishable to an attacker from that of an iChip performing ordinary re-encryption. But a reader with knowledge of x could verify that the check-bit was correct. Thus an attacker could unknowingly produce a clone that outputs re-encryptions without the necessary checkbit. A valid reader would identify such a clone with high probability (about 1/2) as invalid.

We know of no way to create an iChip that can provide good privacy and also provides proof of the non-existence of a subliminal channel. In practice, open design specifications and software are probably the most practical way to provide assurance of the clonability property. It is also worth recalling, from our discussion above, that a stateless iChip *is* subject to simple replay attacks. An attacker can successfully replay a harvested value (once) at any future time irrespective of the existence of a subliminal channel. (An iChip that maintains state, on the other hand—a counter, for instance—could transmit its state information over a subliminal channel and thereby render old, harvested outputs such that they would no longer appear valid during replay.)

We leave as an open problem the design of techniques to help demonstrate the non-existence of a subliminal channel in an iChip, as well as proofs on the limitations of such techniques.

6 Conclusion

We have highlighted and discussed the vulnerability of the VeriChip to simple cloning attacks. For security systems that rely on VeriChips for authentication – like payment systems and physical access-control systems – the consequences are serious. With little sophistication and at little expense, an attacker can undermine system security by surreptitiously capturing and replaying VeriChip signals.

¹⁰ Alternatively, an iChip could randomize its identifier, i.e., output ciphertexts on plaintext (s, r) , where r is a random value. A reader could then detect a cloned VeriChip on the basis of multiple uses of the same randomized string r . Such a reader would, however, have to have access to a history of presented ciphertexts.

Somewhat paradoxically, though, we maintain that a VeriChip *should* be vulnerable to cloning by design, to discourage physical attacks on VeriChip bearers. We maintain that VeriChips should consequently serve only to *identify* their bearers, not to *authenticate* them. We propose a design for an implantable RFID tag that we call an iChip. An iChip emits an identifier through a simple cryptographic scheme that helps protect privacy but at the same time expressly enables straightforward cloning.

We offer no categorical judgment as to whether or not VeriChip (or iChip) implantation is beneficial on balance and no prognostication as to whether or not it will become popular. One author of this paper himself bears an implanted VeriChip [10], and is effectively serving as an experimental subject. This may be the only good way to explore the pros and cons of such devices. It certainly behooves the technological community, however, to reflect on the security and privacy features of iChips as carefully and as early as possible.

References

1. Second meeting of the American Health Information Community, 29 November 2005. Online materials referenced 2006 at <http://www.hhs.gov/healthit/m20051129.html>.
2. K. Albrecht and L. McIntyre. *The Spychips Threat: Why Christians Should Resist RFID and Computer Tracking*. Nelson Current, 2006.
3. G. Ateniese, J. Camenisch, and B. de Madeiros. Untraceable RFID tags via insubvertible encryption. In *ACM Conference on Computer and Communications Security (CCS)*, 2005.
4. G. Avoine. Security and privacy in RFID systems. <http://lasecwww.epfl.ch/~gavoine/rfid/>, 2006. Online bibliography.
5. A. Bahney. High tech, under the skin. *New York Times*, 2 February 2006. Referenced 2006 at <http://www.nytimes.com/2006/02/02/fashion/thursdaystyles/02tags.html>.
6. S. Bono, M. Green, A. Stubblefield, A. Juels, A. Rubin, and M. Szydlo. Security analysis of a cryptographically-enabled rfid device. In *14th USENIX Security Symposium*, pages 1–16, Baltimore, Maryland, USA, July-August 2005. USENIX. Videos and other information available at www.rfidanalysis.org.
7. VeriChip Corporation. Procedure for verimedTM system use, 2006. Referenced 2006 at www.verimedinfo.com/files/VeriMed%20ER%20Protocol.pdf.
8. T. El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31:469–472, 1985.
9. P. Golle, M. Jakobsson, A. Juels, and P. Syverson. Universal re-encryption for mixnets. In T. Okamoto, editor, *RSA Conference - Cryptographers' Track (CT-RSA)*. Springer-Verlag, 2004.
10. J. Halamka. Straight from the shoulder. *The New England Journal of Medicine*, 353:331–333, 28 July 2005.
11. A. Juels, D. Molnar, and D. Wagner. Security and privacy issues in e-passports. In D. Gollman, G. Li, and G. Tsudik, editors, *IEEE/CreateNet SecureComm*, 2005. Referenced 2006 at <http://www.cs.berkeley.edu/~dmolnar/papers/papers.html>.
12. A. Juels and R. Pappu. Squealing Euros: Privacy protection in RFID-enabled banknotes. In R. Wright, editor, *Financial Cryptography – FC '03*, pages 103–121. Springer-Verlag, 2003. LNCS no. 2742.
13. A. Juels, R.L. Rivest, and M. Szydlo. The blocker tag: Selective blocking of RFID tags for consumer privacy. In V. Atluri, editor, *ACM Conference on Computer and Communications Security (CCS)*, pages 103–111. ACM Press, 2003.
14. Ari Juels. RFID security and privacy: A research survey. *IEEE Journal on Selected Areas in Communication*, 24(2):381–395, February 2006.

15. J. Kent. Malaysia car thieves steal finger. *BBC News*, 31 March 2005. Referenced 2006 at <http://news.bbc.co.uk/2/hi/asia-pacific/4396831.stm>.
16. Z. Kfir and A. Wool. Picking virtual pockets using relay attacks on contactless smartcard systems. In *SecureComm'05*, 2005. Referenced 2006 at <http://eprint.iacr.org/2005/052>.
17. J. Leyden. Barcelona nightclub chips customers. *The Register*, 19 May 2004. Referenced 2006 at <http://www.theregister.co.uk/2004/05/19/veripay/>.
18. R. Malone. RFID — its more than price! *Forbes*, 12 December 2006. Referenced 2006 at www.forbes.com/technology/2005/12/12/rfid-reliability-data-cx_rm_1212rfid.html.
19. A. Newitz. Personal communication, 2006. Annalee Newitz, a correspondent for *Wired News*, underwent a VeriChip implant.
20. C. Purohit. Technology gets under clubbers skin. *CNN*, 9 June 2004. Referenced 2006 at <http://edition.cnn.com/2004/WORLD/europe/06/09/spain.club/>.
21. J. Scheeres. Tracking junior with a microchip. *Wired News*, 10 October 2003. Referenced 2006 at <http://www.wired.com/news/technology/0,1282,60771,00.html>.
22. J. Scheeres. Implantable chip, on sale now. *Wired News*, 25 October 2002. Referenced 2006 at <http://www.wired.com/news/politics/0,55999-0.html>.
23. B. Schneier. *Applied Cryptography: Protocols, Algorithms, and Source Code in C, Second Edition*. Wiley, 1995.
24. Verichip corporation web site, 2006. <http://www.verichipcorp.com/>.
25. G. Wang. Bibliography on group signatures, 2006. Online bibliography. Referenced 2006 at www.i2r.a-star.edu.sg/icsd/staff/guilin/bible/group-sign.htm.
26. W. Weissert. Mexican attorney general personally goes high-tech for security. *USA Today*, 14 July 2004. Associated Press. Referenced 2006 at http://www.usatoday.com/tech/news/2004-07-14-mex-security-implant_x.htm.
27. J. Westhues. Hacking the prox card. In S. Garfinkel and B. Rosenberg, editors, *RFID: Applications, Security, and Privacy*, pages 291–300. Addison-Wesley, 2005.
28. J. Westhues. Proxmarkii description, 2006. Web site. Referenced 2006 at <http://cq.cx/proxmarkii.pl>.

A Appendix: VeriMedTM Scanning Instructions

Procedure for VeriMedTM System Use Patient Scanning—Information Retrieval

Preparation

- Place the VeriMed Reader on or adjacent to a vital signs assessment cart or in patient registration area.
- Log on to your facility’s patient information system (or ED Information system if present).

Scanning Procedure Integrated with Patient Assessment

- Any patient presenting confused or unconscious should be assumed to have a subdermal RFID Microchip (VeriMed Microchip).
- Be prepared to scan patient for a subdermal microchip while the automated vital sign devices are obtaining data (i.e., BP insulation cuff, oximeter).
- VeriMed Microchip location: The common Site for VeriMed Microchip insertion is the posterior right upper arm, midway between the shoulder and elbow. The site may vary based on preexisting medical conditions or patient preference.

Scanning Procedure

- For more detailed information refer to the instruction manual packaged with the VeriMed Reader.
- Activate the reader by pressing down on the large button located below the display screen.
- The reader will go through several status checks.
- When the word “Ready” or “Searching” displays, depress button and hold down to initiate the scanning process. “Searching” will be displayed on the screen.
- With the button continuously depressed, slowly pass the reader over the area of the patient to be scanned. The VeriMed Reader should be no more than three Inches from skin level for optimal read accuracy. It is not necessary to remove clothing to read a VeriMed subdermal microchip.
- When a VeriMed subdermal RFID microtransponder is encountered, a beep will sound and the screen will briefly display “VeriChip” followed by the 16-digit VeriMed ID number.
- The number will be displayed for at least 60 seconds, after which the unit will power down to conserve battery life. Record the ID number.
- Repeat the above steps to reacquire the number or to learn another patient’s ID number.
- The numbers are not stored within the VeriMed Reader and thus cannot be recalled.

Accessing Patient Information

Follow your institution’s information System protocol for entering the VeriMed ID number to obtain associated information.

Should a VeriMed ID number be obtained which is not accessible within your institution’s information System or your facility does not utilize an Electronic Medical Record system (EMR or SHR), a VeriChip Corporation credentialed physician can obtain patient-provided information furnished at the time of microchip insertion. This procedure is initiated by logging on to a password-protected section of www.verimedinfo.com.