Technology Evaluation ■

The Security Implications of VeriChip Cloning

John Halamka, MD, Ari Juels, Adam Stubblefield, MD, Jonathan Westhues, MD

Abstract The VeriChip™ is a 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 the authors explain, the VeriChip is vulnerable to simple, over-the-air *spoofing* attacks. In particular, an attacker capable of scanning a VeriChip, eavesdropping on its signal, or simply learning its serial number can create a spoof device whose radio appearance is indistinguishable from the original. We explore the practical implications of this security vulnerability. The authors 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 spoof; an attacker then has less incentive to coerce victims or extract VeriChips from victims' bodies.

■ **J Am Med Inform Assoc.** 2006;13:601–607. DOI 10.1197/jamia.M2143.

Introduction

The VeriChip is a commercially produced, human-implantable microchip.[24] It is designed to serve as an 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 health care 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.[25] 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[15,19] marketed under the product name VeriPay. It has even acquired a degree of chic among certain technophiles, who are exploring applications in daily life.[3]

Affiliations of the authors: Beth Israel-Deaconess Medical Center, Boston, MA (JH); RSA Laboratories, Bedford, MA (AJ); Johns Hopkins University, Baltimore, MD (AS, JW).

Correspondence and reprints: John Halamka, MD, 1135 Tremont, Boston, MA 02120; e-mail: <jhalamka@caregroup.harvard.edu>.

Received for review: 05/10/06; accepted for publication: 08/07/06

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.[4]

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 ever-present 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: The VeriChip is effectively a kind of license plate for people.

Privacy is not the only concern that the VeriChip raises. As we explain in this paper, the VeriChip is vulnerable to a straightforward *spoofing* 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 simulating device as the target VeriChip. In fact, in principle an attacker can simulate a VeriChip on the basis of its serial number alone. We emphasize that by spoofing, we mean emulation of device communication, not physical duplication. (As a VeriChip reader does not visually perceive a communicating chip, physical duplication is not necessary for spoofing attacks, and thus not relevant to most of our security discussion here.) For most security applications, the claim by VeriChip Corporation that the VeriChip "cannot be counterfeited" is effectively untrue.

Our main thesis is that use of the VeriChip for authentication, i.e., as a *proof* of identity, is inappropriate. As an implanted security device, the VeriChip heightens the risk to its bearers of physically coercive attack, i.e., to use of the VeriChip under duress. An attacker can readily seize and use a physically transferable authenticator, such as an ATM card, without seizing its owner; this is not true of the VeriChip. (To use a common distinction among authentication types, a card is "something you have," while a VeriChip is more closely akin to "something you are.") 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.

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 spoofing may apply to such devices.

Organization

The rest of this paper is organized as follows. First, we provide an overview of the VeriChip. In the next section we describe several healthcare applications that motivate the use of implantable RFID tags like the VeriChip. We present the results of our efforts at spoofing VeriChips in the section after that, and discuss the security and privacy implications in the following section. We then conclude in the last section.

Overview of the VeriChip

The VeriChip is an RFID tag. As noted above, RFID tags are proliferating into many domains of use. Already they play a considerable role in everyday activities. Contactless building-access cards—the type that operate when held in proximity to readers, rather than swiped—contain RFID tags; these devices are often called "proximity cards." Many major banks are incorporating RFID tags into credit cards to enable "touch-and-go" payment functionality. These deployments follow on the heels of ExxonMobil Speedpass, an RFID-based payment device in use for years. Similarly, public transit systems are increasingly offering contactless fare-payment devices that contain RFID tags.

As an implanted device, the VeriChip is unusual among RFID tags in its long life span. VeriChip Corporation specifies a life span of fifteen years, a figure based on the normal maximum age of animals with implanted RFID tags that are similar to the VeriChip.[20]

The VeriChip 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 synchronization 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, later in this paper.

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. We believe that healthcare 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 example of this type—or the chip can reside in an identification wristband 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. At the time of writing, even the least expensive tags cost more than 10 cents apiece in quantity.[17] RFID tags may have up to a 5% failure rate during manufacturing, resulting in a potentially unreadable wristband.[16] 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 identifier, (in the case of VeriChip, this identifier is a 16 digit number built into the chip) 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.[9]

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: medical records at BIDMC and Verichip contain the patient's gender, age, and other demographic information, all of which can serve as a 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 referenced below (VeriChip Corporation, 2006).[5]

Implanted chips, like any device inserted into the human body, may elicit an adverse tissue reaction, result in infection, and may migrate from the original site of insertion. Such side effects are rare, but were described during the FDA approval process. To mitigate the risk of migration, especially when the chip is subjected to MRI magnetic forces, the chip is coated with glass and a substance which encourages cellular adherence so that it becomes fixed subcutaneously.

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 sensitive 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 a patient beds. These active tags are about the size of a 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 of 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.

Spoofing the VeriChip

We now explain our spoofing experiments on the VeriChip. For these experiments we used the "proxmarkii" generalized RFID tag reader/spoofer. The proxmarkii is an RFID reading and simulation device developed by Westhues, who used an earlier version to demonstrate spoofing attacks against proximity cards.[26,27] 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.

We were able to study three different VeriChip tags (two unimplanted, one implanted). The unimplanted VeriChips were provided to us as free samples by the medical director of VeriChip, while the implanted VeriChip belonged to a human volunteer.

We first discovered that 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 134 kHz, a carrier clock cycle lasts about 7.46 μs.) 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). In other words, each bit is encoded as either the transition from a high to a low amplitude or a low to high amplitude.

We identified only 32 bits of the 128-bit transmitted value that appear to vary among the three tags we studied. 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 (base 10) code that the legitimate reader reports. It is possible (although unlikely in our view) 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 spoofing, however, does not require a deep look into the structure of the tags' IDs. Since the VeriChip always transmits exactly the same information, spoofing a VeriChip is just a matter of determining the signal that the tag transmits and building a device that mimics 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 spoofing 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* spoofing, 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. We were thus able to record the response of a legitimate VeriChip and later use our encoding of the response to send that same response to a legitimate reader.

Viewed another way, we performed a *replay* attack against the VeriChip, meaning that we simply captured a signal from a VeriChip 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 spoofing, i.e., the harvested signal may be replayed indefinitely while appearing valid to a reader.

Implications of Spoofing

As we see, the practicality of our spoofing attack is determined not by any cryptographic factors, but simply by the distance from which the tag can be read. 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 spoofing 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.[12] 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 spoof *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.

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 adopt a false identity. For example, a drug addict might attempt to spoof 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. 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 spoof or perhaps modify an existing VeriChip to output a false ID. Such an attack could be mitigated by confirming the patient's identity via other information stored in the associated records.

Existential Spoofing

In addition to the risk of spoofing practiced through surreptitious scanning and replay of VeriChip signals, there is also, as mentioned above, a threat of existential spoofing. The IDs in the three VeriChips we obtained appeared to 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 a 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. Our possession of only three VeriChips somewhat constrains our ability 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). 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.

There is in principle a possibility that the checksum is keyed, i.e., it depends upon a secret key shared among VeriChip readers. We consider this possibility remote, in part given comments on our work by VeriChip Corporation as cited below reference 20[20], but also because the use of global secret keys would render reader engineering more difficult and would require rather more specialized RFID tags (as opposed to, e.g, redeployment of the RFID devices currently implanted in house-pets). If the VeriChip system did depend upon global keying, however, 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. An obvious upper bound on the checksum is 30 bits (as at least 2 bits are required to render 3 IDs distinct), so even in the very worst case the "brute force" attack of the third method would likely be possible.

If an attacker can mount an existential spoofing attack, the implications are serious. Consider a corporation that uses the VeriChip to control access to a secured physical area. If IDs are indeed 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 spoofed, revoking access privileges for that employee would be insufficient: the attacker could simulate other valid IDs in the system. Moreover, the challenging 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? 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 spoofing attacks enabled by the current system—at least until they undergo surgical replacement of their implanted chips.

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

Privacy Implications

There are well known cryptographic tools, like challenge-response protocols, that can defend against over-the-air spoofing attacks like the one we have demonstrated. Paradoxically, though, there is a compelling reason to ensure that an implantable RFID tag is in fact spoofable: an adversary then has little incentive to perform a physical attack against the chip and its bearer. 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.[14] An attacker has a similar incentive to obtain physical possession of a VeriChip or to coerce its bearer if the chip is: 1) Used to secure access to valuable resources; and 2) Hard to spoof.

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 operational setting—i.e., condition 2) above—is a matter largely beyond the control of its bearer. It seems imprudent to rely on the policies of system architects or administrators to avoid implantable-RFID authentication. The following example illustrates this point.

Example: XYZ Pharmacy creates a system in which patrons can obtain their prescription medications from automated dispensing stations. Originally, this system authenticates pharmacy patrons based upon an iris scan and PIN. At this time, XYZ Pharmacy only allows its patrons to identify themselves to dispensing stations and other pharmacy services using their implantable RFID tags. Several years later, however, XYZ Pharmacy changes its policy and offers a convenient new service whereby patrons can authenticate to dispensing stations using just their implantable RFID tags.

Many prescription medications have high street value. Thus, in this example, the pharmacy has created an incentive for physical attack against its patrons' implanted RFID tags. Not only has the pharmacy created this incentive unilaterally, but some of its patrons might 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, like the VeriChip, be easy to spoof by design. In particular, it should not contain cryptographic protections against spoofing, like challenge-response schemes.

What, however, are the implications for user privacy of this requirement?

An RFID tag that emits a static identifier, like the VeriChips in use today, poses a threat to the privacy of its bearer. As explained above, a VeriChip transmits to any off-the-shelf reader what is essentially a unique identifier. That identifer permits physical tracking of a VeriChip: A reader or network of readers can uniquely identify a VeriChip within range. Furthermore, an association can be made between the tag identifier and a real-world identity of the bearer. For example, when a consumer pays for an item in a shop using a credit card, if the shop can scan the consumer's implanted VeriChip, it can establish a binding between the name of the consumer and the VeriChip itself. In principle, then, the shop or any of its affilitates can detect the entry of that consumer into its shops. The short read range of the VeriChip provides some assurance against such attacks, but installations with large antennas, e.g., shops with portal readers, high-power readers, and improvements in RFID technology can erode such protection. See, e.g., Juels et al., 2005; Garfinkel and Rosenberg, 2005 for overviews of the topic of RFID privacy.[11,7]

Researchers have proposed a range of techniques for protecting the privacy of RFID-tag bearers.[10] Most such proposals involve tags emitting identifiers that change over time in a secret, cryptographically determined manner, i.e., such that the outputs of a given tag are unlinkable and unpredictable. At first glance, it may therefore seem that privacy—in the sense of protection from clandestine tracking—cannot co-exist with spoofability. How is it possible to spoof a tag without knowing its internal secrets or what future outputs it will emit?

Somewhat surprisingly, it is possible by using appropriate cryptographic tools to construct a device that achieves the apparent contradiction of *simultaneous* privacy and spoofability. In the paper's online Appendix, available at www.jamia.org, we briefly describe such a device, which we refer to as an *iChip*. Employing a special form of public-key cryptography, the iChip is an illustration of how the characteristic of spoofability need not imply neglect of bearer privacy.

Conclusion

We have highlighted and discussed the vulnerability of the VeriChip to simple spoofing 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.

Somewhat paradoxically, though, we maintain that a VeriChip *should* be vulnerable to spoofing 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.

In reviewing the work presented here, a representative of the VeriChip Corporation has provided basic confirmation of our technical results, stating that "The observations presented by the authors are valid based on the ISO open architecture standards utilized the VeriMedTM implantable microtransponder in broadcasting a unique sixteen digit identification number over its rated distance of 2.5 inches." This representative objected, however, that our exposition "contains numerous conjectures and assumptions regarding possible VeriChip applications and successful attacks on them. These are stated as possibilities and there was no attempt to prove the conjectures or assumptions accurate by putting the proposed attacks to tests in real world situations."[20]

We offer no categorical judgment as to whether or not VeriChip 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,[9] and is effectively serving as an experimental subject. This may be the only good way to explore the pros and cons of such devices. We encourage the technological community, however, to reflect on the security and privacy features of implantable RFID tags as carefully and as early as possible.[6,8,13,18,23]

References ■

1. Second meeting of the American Health Information Community, 29 November 2005. Available at: http://www.hhs.gov/healthit/m20051129.html.
2. Albrecht K, McIntyre L. The Spychips Threat: Why Christians Should Resist RFID and Computer Tracking. Nelson Current, 2006.
3. Bahney A. High tech, under the skin. New York Times, 2 February 2006. Available at: http://www.nytimes.com/2006/02/02/fashion/thursdaystyles/02tags.html. Accessed February 2, 2006.
4. Bono S, Green M, Stubblefield A, Juels A, Rubin A, Szydlo M. Security analysis of a cryptographically-enabled rfid device. In

14th USENIX Security Symposium, pp. 1–16, Baltimore, Maryland, USA, July-August 2005. USENIX Videos and other information available at www.rfidanalysis.org.
5. VeriChip Corporation. Procedure for VeriMedTM system use, 2006. Available at: www.verimedinfo.com/files/VeriMed%20ER%20Protocol.pdf.
6. El Gamal T. A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Trans Inform Theory. 1985; 31:469–72.
7. RFID applications, security, and privacy. Addison-Wesley, 2005.
8. Golle P, Jakobsson M, Juels A, Syverson P. Universal re-encryption for mixnets. In Okamoto, editor, T. RSA Conference—Cryptographers' Track (CT-RSA). Springer-Verlag, 2004.
9. Halamka J. Straight from the shoulder. New Engl J Med. 2005;353:331–3.
10. Juels A. RFID security and privacy: a research survey. IEEE J Select Areas Comm. 2006;24(2):381–95.
11. Juels A, Garfinkel S, Pappu R. RFID privacy: an overview of problems and proposed solutions. IEEE Security and Privacy. 2005;3(3):34–43.
12. Juels A, Molnar D, Wagner D. Security and privacy issues in e-passports. In D. Gollman, G. Li, and G. Tsudik, editors, IEEE/CreateNet SecureComm, 2005. Available at: http://www.cs.berkeley.edu/dmolnar/papers/papers.html.
13. Juels A, Pappu R. Squealing Euros: Privacy protection in RFID-enabled banknotes. In Wright, editor, R. Financial Cryptography—FC '03. Springer-Verlag, 2003. LNCS no. 2742, pp. 103–21.
14. Kent J. Malaysia car thieves steal finger. BBC News, 31 March 2005. Available at: http://news.bbc.co.uk/2/hi/asia-pacific/4396831.stm.
15. Leyden J. Barcelona nightclub chips customers. The Register, 19 May 2004. Available at: at http://www.theregister.co.uk/2004/05/19/veripay/.
16. Malone R. RFID—it's more than price! Forbes, 12 December 2006. Available at: www.forbes.com/technology/2005/12/12/rfid-reliability-data-cx rm 1212rfid.html. Accessed December 12, 2006.
17. OConnor MC. Alien drops tag price to 12.9 cents. RFID Journal, 15 September 2005. Available at http://www.rfidjournal.com/article/articleview/1870/1/1/. Accessed September 15, 2006.
18. Paillier P. Public-key cryptosystems based on composite degree residuosity classes. In Stern, editor, J. EUROCRYPT 99. Springer-Verlag, 1999. LNCS no. 1592, pp. 223–38.
19. Purohit C. Technology gets under clubbers skin. *CNN*, 9 June 2004. Available at: http://edition.cnn.com/2004/WORLD/europe/06/09/spain.club/.
20. Seelig R. Vice President for Medical Applications, VeriChip Corporation. Personal communication via e-mail, 12 July 2006.
21. Scheeres J. Tracking junior with a microchip. Wired News, 10 October 2003. Available at: http://www.wired.com/news/technology/0,1282,60771,00.html. Accessed October 10, 2006.
22. Scheeres J. Implantable chip, on sale now. Wired News, 25 October 2002. Available at: 2006 at http://www.wired.com/news/politics/0,55999-0.html. Accessed October 25, 2006.
23. Schneier B. Applied Cryptography: Protocols, Algorithms, and Source Code in C, Second Edition Wiley, 1995.
24. Verichip corporation web site, 2006. Available at: http://www.verichipcorp.com/.
25. Weissert W. Mexican attorney general personally goes high-tech for security. USA Today, 14 July 2004. Associated Press. Available at: http://www.usatoday.com/tech/news/2004-07-14-mexsecurity-implant x.htm. Accessed July 14, 2006.
26. Westhues J. Hacking the prox card. In Garfinkel and S. Rosenberg, editors, B. RFID: Applications, Security, and Privacy. Addison-Wesley, 2005, pp. 291–300.
27. Westhues J. Proxmarkii description, 2006. Web site. Available at: http://cq.cx/proxmarkii.pl.