

Authenticating Pervasive Devices with Human Protocols

Ari Juels¹ and Stephen A. Weis²

¹ RSA Laboratories, Bedford, MA, USA
ajuels@rsasecurity.com

² Massachusetts Institute of Technology, Cambridge, MA, USA
sweis@mit.edu

Abstract. Forgery and counterfeiting are emerging as serious security risks in low-cost pervasive computing devices. These devices lack the computational, storage, power, and communication resources necessary for most cryptographic authentication schemes. Surprisingly, low-cost pervasive devices like Radio Frequency Identification (RFID) tags share similar capabilities with another weak computing device: people.

These similarities motivate the adoption of techniques from human-computer security to the pervasive computing setting. This paper analyzes a particular human-to-computer authentication protocol designed by Hopper and Blum (HB), and shows it to be practical for low-cost pervasive devices. We offer an improved, concrete proof of security for the HB protocol against passive adversaries.

This paper also offers a new, augmented version of the HB protocol, named HB⁺, that is secure against active adversaries. The HB⁺ protocol is a novel, symmetric authentication protocol with a simple, low-cost implementation. We prove the security of the HB⁺ protocol against active adversaries based on the hardness of the Learning Parity with Noise (LPN) problem.

Keywords: authentication, HumanAut, Learning Parity with Noise (LPN), pervasive computing, RFID

1 Introduction

As low-cost computing devices become more pervasive, counterfeiting may become a more serious security threat. For example, the security of access control or payment systems will rely on the authenticity of low-cost devices. Yet in many settings, low-cost pervasive devices lack the resources to implement standard cryptographic authentication protocols. Low-cost Radio Frequency Identification (RFID) tags exemplify such resource-constrained devices. Viewing them as possible beneficiaries of our work, we use RFID tags as a basis for our discussions of the issues surrounding low-cost authentication.

Low-cost RFID tags in the form of Electronic Product Codes (EPC) are poised to become the most pervasive device in history [10]. Already, there are

billions of RFID tags on the market, used for applications like supply-chain management, inventory monitoring, access control, and payment systems. Proposed as a replacement for the Universal Product Code (UPC) (the barcode found on most consumer items), EPC tags are likely one day to be affixed to everyday consumer products.

Today’s generation of basic EPC tags lack the computational resources for strong cryptographic authentication. These tags may only devote hundreds of gates to security operations. Typically, EPC tags will passively harvest power from radio signals emitted by tag readers. This means they have no internal clock, nor can perform any operations independent of a reader.

In principle, standard cryptographic algorithms – asymmetric or symmetric – can support authentication protocols. But implementing an asymmetric cryptosystem like RSA in EPC tags is entirely infeasible. RSA implementations require tens of thousands of gate equivalents. Even the storage for RSA keys would dwarf the memory available on most EPC tags.

Standard symmetric encryption algorithms, like DES or AES, are also too costly for EPC tags. While current EPC tags may have at most 2,000 gate equivalents available for security (and generally much less), common DES implementations require tens of thousands of gates. Although recent light-weight AES implementations require approximately 5,000 gates [11], this is still too expensive for today’s EPC tags.

It is easy to brush aside consideration of these resource constraints. One might assume that Moore’s Law will eventually enable RFID tags and similar devices to implement standard cryptographic primitives like AES. But there is a countervailing force: Many in the RFID industry believe that pricing pressure and the spread of RFID tags into ever more cost-competitive domains will mean little effective change in tag resources for some time to come, and thus a pressing need for new lightweight primitives.

Contribution. This paper’s contribution is a novel, lightweight, symmetric-key authentication protocol that we call HB^+ . HB^+ may be appropriate for use in today’s generation of EPC tags or other low-cost pervasive devices. We prove the security of this protocol against both passive eavesdroppers and adversaries able to adaptively query legitimate tags. We also offer an improved, concrete security reduction of a prior authentication protocol HB that is based on the same underlying hardness problem.

Organization. In Section 2, we describe the basic “human authentication” or “HumanAut” protocol, due to Hopper and Blum (HB), from which we build an authentication protocol appropriate for RFID tags that is secure against passive eavesdroppers. We discuss the underlying hardness assumption, the “Learning Parity with Noise” (LPN) problem, in Section 3. Section 4 offers our new, enhanced variant of the HB protocol, HB^+ , that is secure against adversaries able to query legitimate tags *actively*. Section 5 presents an improved, concrete reduction of the LPN problem to the security of the HB protocol, and shows a concrete

reduction of security from the HB protocol to the HB⁺ protocol. Finally, Section 6 states several open problems related to this work.

In this preliminary version of the full paper, we relegate many details to the appendix, most notably our security proofs and discussion of our security model.

1.1 The Problem of Authentication

It seems inevitable that many applications will come to rely on basic RFID tags or other low-cost devices as authenticators. For example, the United States Food and Drug Administration (FDA) proposed attaching RFID tags to prescription drug containers in an attempt to combat counterfeiting and theft [13].

Other RFID early-adopters include public transit systems and casinos. Several cities around the world use RFID bus and subway fare cards, and casinos are beginning to deploy RFID-tagged gambling chips and integrated gaming tables. Some people have even had basic RFID tags with static identifiers implanted in their bodies as payment devices or medical-record locators [40].

Most RFID devices today promiscuously broadcast a static identifier with no explicit authentication procedure. This allows an attacker to surreptitiously scan identifying data in what is called a *skimming* attack. Besides the implicit threat to privacy, skimmed data may be used to produce cloned tags, exposing several lines of attack.

For example, in a *swapping* attack, a thief skims valid RFID tags attached to products inside a sealed container. The thief then manufactures cloned tags, seals them inside a decoy container (containing, e.g., fraudulent pharmaceuticals), and swaps the decoy container with the original. Thanks to the ability to clone a tag and prepare the decoy in advance, the thief can execute the physical swap very quickly. In the past, corrupt officials have sought to rig elections by conducting this type of attack against sealed ballot boxes [37].

Clones also create denial-of-service issues. If multiple, valid-looking clones appear in a system like a casino, must they be honored as legitimate? Or must they all be rejected as frauds? Cloned tags could be intentionally designed to corrupt supply-chain databases or to interfere with retail shopping systems. Denial of service is especially critical in RFID-based military logistics systems.

Researchers have recently remonstrated practical cloning attacks against real-world RFID devices. Mandel, Roach, and Winstein demonstrated how to read access control proximity card data from a range of several feet and produce low-cost clones [27] (despite the fact that these particular proximity cards only had a legitimate read range of several inches). A team of researchers from Johns Hopkins University and RSA Laboratories recently elaborated attacks against a cryptographically-enabled RFID transponder that is present in millions of payment and automobile immobilization systems [6]. Their attacks involved extraction of secret keys and simulation of target transponders; they demonstrated an existing risk of automobile theft or payment fraud from compromise of RFID systems.

Example EPC Specifications	
Storage:	128-512 bits of read-only storage.
Memory:	32-128 bits of volatile read-write memory.
Gate Count:	1000-10000 gates.
Security Gate Count Budget:	200-2000 gates.
Operating Frequency:	868-956 MHz (UHF).
Scanning Range:	3 meters.
Performance:	100 read operations per second.
Clock Cycles per Read:	10,000 clock cycles.
Tag Power Source:	Passively powered by Reader via RF signal.
Power Consumption:	10 microwatts.
Features:	Anti-Collision Protocol Support Random Number Generator

Fig. 1. Example specification for a 5-10 low-cost RFID tag.

1.2 Previous Work on RFID Security

As explained above, securing RFID tags is challenging because of their limited resources and small physical form. Figure 1 offers specifications that might be realistic for a current EPC tag. Such limited power, storage, and circuitry, make it difficult to implement traditional authentication protocols. This problem has been the topic of a growing body of literature.

A number of proposals for authentication protocols in RFID tags rely on the use of symmetric-key primitives. The authors often resort to a hope for enhanced RFID tag functionality in the future, and do not propose use of any particular primitive. We do not survey this literature in any detail here, but refer the reader to, e.g., [17, 32, 35, 36, 39, 43].

Other authors have sought to enforce privacy or authentication in RFID systems while avoiding the need for implementing standard cryptographic primitives on tags, e.g., [12, 21, 22, 24, 23, 32].

Feldhofer, Dominikus, and Wolkerstorfer [11] propose a low-cost AES implementation, potentially useful for higher-cost RFID tags, but still out of reach for basic tags in the foreseeable future.

1.3 Humans vs. RFID Tags

Low-cost RFID tags and other pervasive devices share many limitations with another weak computing device: human beings. The target cost for a EPC-type RFID tag is in the US 0.05-0.10 (5-10) range. The limitations imposed at these costs are approximated in Figure 1. We will see that in many ways, the computational capacities of people are similar.

Like people, tags can neither remember long passwords nor keep long calculations in their working memory. Tags may only be able to store a short secret

of perhaps 32-128 bits, and be able to persistently store 128-512 bits overall. A working capacity of 32-128 bits of volatile memory is plausible in a low-cost tag, similar to how most human beings can maintain about seven decimal digits in their immediate memory [31].

Neither tags nor humans can efficiently perform lengthy computations. A basic RFID tag may have a total of anywhere from 1000-10000 gates, with only 200-2000 budgeted specifically for security. (Low-cost tags achieve only the lower range of these figures.) As explained above, performing modular arithmetic over large fields or evaluating standardized cryptographic functions like AES is currently not feasible in a low-cost device or for many human beings.

Tags and people each have comparative advantages and disadvantages. Tags are better at performing logical operations like ANDs, ORs and XORs. Tags are also better at picking random values than people – a key property we rely on for the protocols presented here. However, tag secrets can be completely revealed through physical attacks, such as electron microscope probing [1]. In contrast, physically attacking people tends to yield unreliable results.

Because of their similar sets of capabilities, this paper considers adopting human authentication protocols in low-cost pervasive computing devices. The motivating human-computer authentication protocols we consider were designed to allow a person to log onto an untrusted terminal while someone spies over his/her shoulder, without the use of any scratch paper or computational devices. Clearly, a simple password would be immediately revealed to an eavesdropper.

Such protocols are the subject of Carnegie Mellon University’s HumanAut project. Earlier work by Matsumoto and Imai [29] and Matsumoto [28] propose human authentication protocols that are good for a small number of authentications [41]. Naor and Pinkas describe a human authentication scheme based on “visual cryptography” [33]. However, this paper will focus primarily on the human authentication protocols of Hopper and Blum [18, 19].

2 The HB Protocol

We begin by reviewing Hopper and Blum’s secure human authentication protocol [18, 19], which we will refer to as the HB protocol. We then place it in the RFID setting. The HB protocol is only secure against passive eavesdroppers – not active attackers. In Section 4, we augment the HB protocol against active adversaries that may initiate their own tag queries.

Suppose Alice and a computing device C share an k -bit secret \mathbf{x} , and Alice would like to authenticate herself to C . C selects a random challenge $\mathbf{a} \in \{0, 1\}^k$ and sends it to Alice. Alice computes the binary inner-product $\mathbf{a} \cdot \mathbf{x}$, then sends the result back to C . C computes $\mathbf{a} \cdot \mathbf{x}$, and accepts if Alice’s parity bit is correct.

In a single round, someone imitating Alice who does not know the secret \mathbf{x} will guess the correct value $\mathbf{a} \cdot \mathbf{x}$ half the time. By repeating for r rounds, Alice can lower the probability of naïvely guessing the correct parity bits for all r rounds to 2^{-r} .

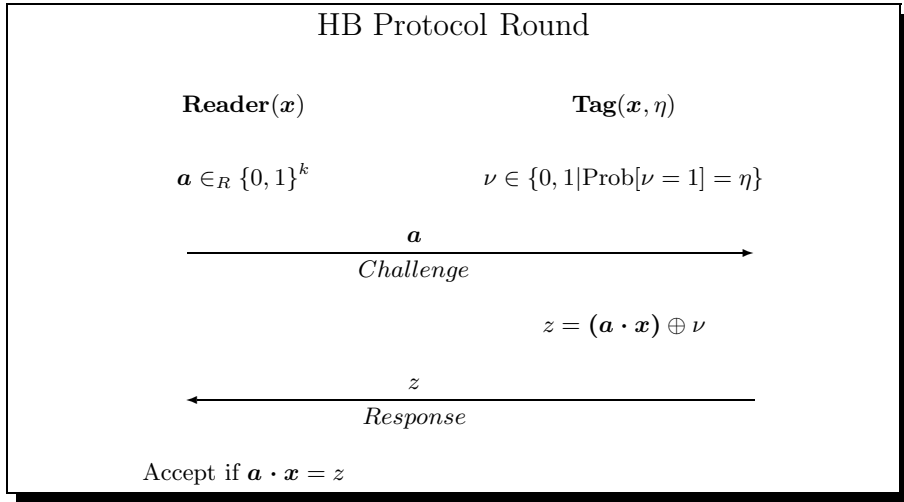


Fig. 2. A single round of the HB authentication protocol.

Of course, an eavesdropper capturing $O(k)$ valid challenge-response pairs between Alice and C can quickly calculate the value of \mathbf{x} through Gaussian elimination. To prevent revealing \mathbf{x} to passive eavesdroppers, Alice can inject noise into her response. Alice intentionally sends the wrong response with constant probability $\eta \in (0, \frac{1}{2})$. C then authenticates Alice's identity if fewer than ηr of her responses are incorrect.

Figure 2 illustrates a round of the HB protocol in the RFID setting. Here, the tag plays the role of the prover (Alice) and the reader of the authenticating device C . Each authentication consists of r rounds, where r is a security parameter.

The HB protocol is very simple to implement in hardware. Computing the binary inner product $\mathbf{a} \cdot \mathbf{x}$ only requires bitwise AND and XOR operations that can be computed on the fly as each bit of \mathbf{a} is received. There is no need to buffer the entire value \mathbf{a} . The noise bit ν can be cheaply generated from physical properties like thermal noise, shot noise, diode breakdown noise, metastability, oscillation jitter, or any of a slew of other methods. Only a single random bit value is needed in each round. This can help avoid localized correlation in the random bit stream, as occurs in chaos-based or diode breakdown random number generators.

Remark: The HB protocol can be also deployed as a *privacy-preserving* identification scheme. A reader may initiate queries to a tag without actually knowing whom that tag belongs to. Based on the responses, a reader can check its database of known tag values and see if there are any likely matches. This preserves the privacy of a tag's identity, since an eavesdropper only captures an instance of the LPN problem, which is discussed in the Section 3.

3 Learning Parity in the Presence of Noise

Suppose that an eavesdropper, i.e., a passive adversary, captures q rounds of the HB protocol over several authentications and wishes to impersonate Alice. Consider each challenge \mathbf{a} as a row in a matrix \mathbf{A} ; similarly, let us view Alice's set of responses as a vector \mathbf{z} . Given the challenge set \mathbf{A} sent to Alice, a natural attack for the adversary is to try to find a vector \mathbf{x}' that is functionally close to Alice's secret \mathbf{x} . In other words, the adversary might try to compute a \mathbf{x}' which, given challenge set \mathbf{A} in the HB protocol, yields a set of responses that is close to \mathbf{z} . (Ideally, the adversary would like to figure out \mathbf{x} itself.)

The goal of the adversary in this case is akin to the core problem on which we base our investigations in this paper. This problem is known as the *Learning Parity in the Presence of Noise*, or LPN problem. The LPN problem involves finding a vector \mathbf{x}' such that $|(\mathbf{A} \cdot \mathbf{x}') \oplus \mathbf{z}| \leq \eta q$, where $|\mathbf{v}|$ represents the Hamming weight of vector \mathbf{v} . Formally, it is as follows:

Definition 1 (LPN Problem). *Let \mathbf{A} be a random $q \times k$ binary matrix, let \mathbf{x} be a random k -bit vector, let $\eta \in (0, \frac{1}{2})$ be a constant noise parameter, and let \mathbf{v} be a random q -bit vector such that $|\mathbf{v}| \leq \eta q$. Given \mathbf{A} , η , and $\mathbf{z} = (\mathbf{A} \cdot \mathbf{x}) \oplus \mathbf{v}$, find a k -bit vector \mathbf{x}' such that $|(\mathbf{A} \cdot \mathbf{x}') \oplus \mathbf{z}| \leq \eta q$.*

The LPN problem may also be formulated and referred to as as the *Minimum Disagreement Problem* [9], or the problem of finding the closest vector to a random linear error-correcting code; also known as the syndrome decoding problem [2, 26]. Syndrome decoding is the basis of the McEliece public-key cryptosystem [30] and other cryptosystems, e.g., [8, 34]. Algebraic coding theory is also central to Stern's public-key identification scheme [38]. Chabaud offers attacks that, although infeasible, help to establish practical security parameters for error-correcting-code based cryptosystems [7].

The LPN problem is known to be NP-Hard [2], and is hard even within an approximation ratio of two [16]. A longstanding open question is whether this problem is difficult for random instances. A result by Kearns proves that the LPN is not efficiently solvable in the statistical query model [25]. An earlier result by Blum, Furst, Kearns, and Lipton [3] shows that given a random k -bit vector \mathbf{a} , an adversary who could weakly predict the value $\mathbf{a} \cdot \mathbf{x}$ with advantage $\frac{1}{k^c}$ could solve the LPN problem. Hopper and Blum [18, 19] show that the LPN problem is both pseudo-randomizable and log-uniform.

The best known algorithm to solve random LPN instances is due to Blum, Kalai, and Wasserman, and has a subexponential runtime of $2^{O(\frac{k}{\log k})}$ [4]. Based on a concrete analysis of this algorithm, we discuss estimates for lower-bounds on key sizes for the HB and HB⁺ protocols in the full version of the paper.

As mentioned above, the basic HB protocol is only secure against passive eavesdroppers. It is not secure against an active adversary with the ability to query tags. If the same challenge \mathbf{a} is repeated $\Omega((1 - 2\eta)^{-2})$ times, an adversary can learn the error-free value of $\mathbf{a} \cdot \mathbf{x}$ with very high probability. Given $\Omega(k)$ error-free values, an adversary can quickly compute \mathbf{x} through Gaussian elimination.

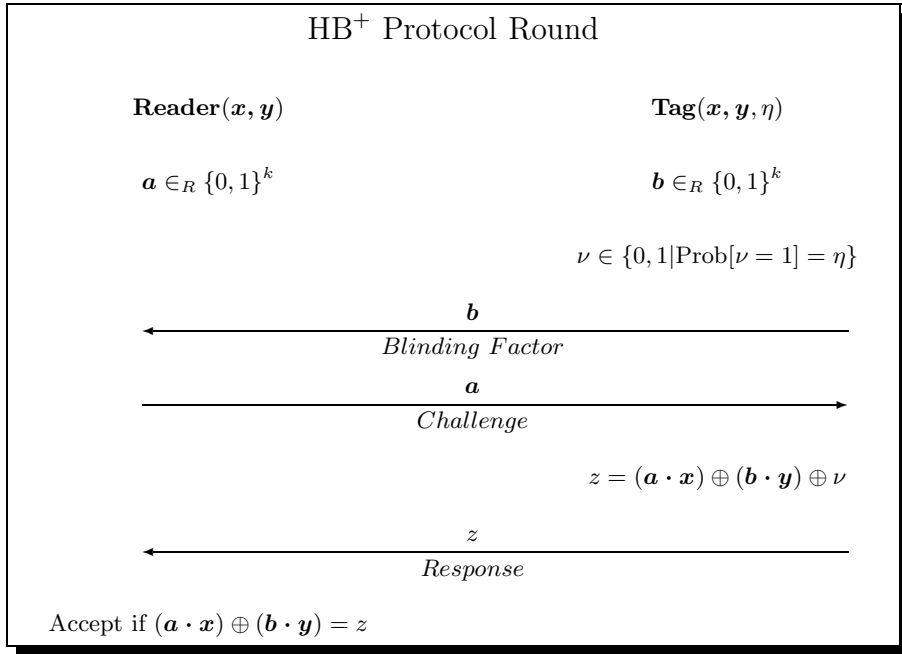


Fig. 3. A single round of the HB⁺ protocol.

4 Authentication Against Active Adversaries

In this section, we show how to strengthen the HB protocol against active adversaries. We refer to the improved protocol as HB⁺. HB⁺ prevents corrupt readers from extracting tag secrets through adaptive (non-random) challenges, and thus prevents counterfeit tags from successfully authenticating themselves. Happily, HB⁺ requires marginally more resources than the “passive” HB protocol in the previous section.

4.1 Defending Against Active Attacks: The HB⁺ Protocol

The HB⁺ protocol is quite simple, and shares a familiar “commit, challenge, respond” format with classic protocols like Fiat-Shamir identification. Rather than sharing a single k -bit random secret \mathbf{x} , the tag and reader now share an additional k -bit random secret \mathbf{y} .

Unlike the case in the HB protocol, the tag in the HB⁺ protocol first generates random k -bit “blinding” vector \mathbf{b} and sends it to the reader. As before, the reader challenges the tag with an k -bit random vector \mathbf{a} .

The tag then computes $z = (\mathbf{a} \cdot \mathbf{x}) \oplus (\mathbf{b} \cdot \mathbf{y}) \oplus \nu$, and sends the response z to the reader. The reader accepts the round if $z = (\mathbf{a} \cdot \mathbf{x}) \oplus (\mathbf{b} \cdot \mathbf{y})$. As before, the reader authenticates a tag after r rounds if the tag’s response is incorrect in less than ηr rounds. This protocol is illustrated in Figure 3.

One reason that Hopper and Blum may not have originally proposed this protocol improvement is that it is inappropriate for use by humans. It requires the tag (playing the role of the human), to generate a random k -bit string \mathbf{b} on each query. If the tag (or human) does not generate uniformly distributed \mathbf{b} values, it may be possible to extract information on \mathbf{x} or \mathbf{y} .

To convert HB^+ into a two-round protocol, an intuitive idea would be to have the tag transmit its \mathbf{b} vector along with its response bit z . Being able to choose \mathbf{b} after receiving \mathbf{a} , however, may give too much power to an adversarial tag. In particular, our security reduction in Section 5.4 relies on the tag transmitting its \mathbf{b} value first. It's an open question whether there exists a secure two-round version of HB^+ . Another open question is whether security is preserved if \mathbf{a} and \mathbf{b} are transmitted simultaneously on a duplex channel.

Beyond the requirements for the HB protocol, HB^+ only requires the generation of k random bits for \mathbf{b} and additional storage for an k -bit secret \mathbf{y} . As before, computations can be performed bitwise; there is no need for the tag to store the entire vectors \mathbf{a} or \mathbf{b} . Overall, this protocol is still quite efficient to implement in hardware, software, or perhaps even by a human being with a decent randomness source.

4.2 Security Intuition

As explained above, an active adversary can defeat the basic HB protocol and extract \mathbf{x} by making adaptive, non-random \mathbf{a} challenges to the tag. In our augmented protocol HB^+ , an adversary can still, of course, select \mathbf{a} challenges to mount an active attack.

By selecting its own random blinding factor \mathbf{b} , however, the tag effectively prevents an adversary from actively extracting \mathbf{x} or \mathbf{y} with non-random \mathbf{a} challenges. Since the secret \mathbf{y} is independent of \mathbf{x} , we may think of the tag as initiating an independent, interleaved HB protocol with the roles of the participants reversed. In other words, an adversary observing \mathbf{b} and $(\mathbf{b} \cdot \mathbf{y}) \oplus \nu$ should not be able to extract significant information on \mathbf{y} .

Recall that the value $(\mathbf{b} \cdot \mathbf{y}) \oplus \nu$ is XORed with the the output of the original, reader-initiated HB protocol, $\mathbf{a} \cdot \mathbf{x}$. This prevents an adversary from extracting information through non-random \mathbf{a} challenges. Thus, the value $(\mathbf{b} \cdot \mathbf{y}) \oplus \nu$ effectively “blinds” the value $\mathbf{a} \cdot \mathbf{x}$ from both passive and active adversaries.

This observation underlies our proof strategy for the security of HB^+ . We argue that an adversary able to efficiently learn \mathbf{y} can efficiently solve the LPN problem. In particular, an adversary that does not know \mathbf{y} cannot guess $\mathbf{b} \cdot \mathbf{y}$, and therefore cannot learn information about \mathbf{x} from a tag response z .

The blinding therefore protects against leaking the secret \mathbf{x} in the face of active attacks. Without knowledge of \mathbf{x} or \mathbf{y} , an adversary cannot create a fake tag that will respond correctly to a challenge \mathbf{a} . In other words, cloning will be infeasible. In Section 5, we will present a concrete reduction from the LPN problem to the security of the HB^+ protocol. In other words, an adversary with some significant advantage of impersonating a tag in the HB^+ protocol can be used to solve the LPN problem with some significant advantage.

5 Security Proofs

We will first present concrete security notation in Section 5.1. Section 5.2 reviews key aspects of the Blum et al. proof strategy that reduces the LPN problem to the security of the HB protocol [3]. We offer a more thorough and concrete version of the Blum et al. reduction in Section 5.3. In Section 5.4, we present a concrete reduction from the HB protocol to the HB⁺ protocol. Finally, in Section 5.5, we combine these results to offer a concrete reduction of the LPN problem to the security of the HB⁺ protocol.

5.1 Notation and Definitions

We define a tag-authentication system in terms of a pair of probabilistic functions $(\mathcal{R}, \mathcal{T})$, namely a reader function \mathcal{R} and a tag function \mathcal{T} .

The tag function \mathcal{T} is defined in terms of a noise parameter η , a k -bit secret \mathbf{x} , and a set of q random k -bit vectors $\{\mathbf{a}^{(i)}\}_{i=1}^q$ that we view for convenience as a matrix \mathbf{A} . Additionally, \mathcal{T} includes a k -bit secret \mathbf{y} for protocol HB⁺. We let q be the maximum number of protocol invocations on \mathcal{T} in our experiments.

For protocol HB, we denote the fully parameterized tag function by $\mathcal{T}_{x,\mathbf{A},\eta}$. On the i th invocation of this protocol, \mathcal{T} is presumed to output $(\mathbf{a}^{(i)}, (\mathbf{a}^{(i)} \cdot \mathbf{x}) \oplus \nu)$. Here ν is a bit of noise parameterized by η . This models a passive eavesdropper observing a round of the HB protocol. Note that the oracle $\mathcal{T}_{x,\mathbf{A},\eta}$ takes no input and essentially acts as an interface to a flat transcript. For this protocol, the reader \mathcal{R}_x takes as input a pair (\mathbf{a}, z) . It outputs either “accept” or “reject”.

For protocol HB⁺, we denote a fully parameterized tag function as $\mathcal{T}_{x,\mathbf{y},\eta}$. This oracle internally generates random blinding vectors \mathbf{b} . On the i th invocation of \mathcal{T} for this protocol, the tag outputs some random $\mathbf{b}^{(i)}$, takes a challenge vector $\mathbf{a}^{(i)}$ (that could depend on $\mathbf{b}^{(i)}$) as input, and outputs $z = (\mathbf{a}^{(i)} \cdot \mathbf{x}) \oplus (\mathbf{b}^{(i)} \cdot \mathbf{y}) \oplus \nu$. This models an active adversary querying a tag in a round of the HB⁺ protocol. For this protocol, the reader $\mathcal{R}_{x,\mathbf{y}}$ takes as input a triple $(\mathbf{a}, \mathbf{b}, z)$ and outputs either “accept” or “reject”.

For both protocols HB and HB⁺, we consider a two-phase attack model involving an adversary comprising a pair of functions $\mathcal{A} = (\mathcal{A}_{query}, \mathcal{A}_{clone})$, a reader \mathcal{R} , and a tag \mathcal{T} . In the first, “query” phase, the adversarial function \mathcal{A}_{query} has oracle access to \mathcal{T} and outputs some state σ .

The second, “cloning” phase involves the adversarial function \mathcal{A}_{clone} . The function \mathcal{A}_{clone} takes as input a state value σ . In HB⁺, it outputs a blinding factor \mathbf{b}' (when given the input command “initiate”). In both HB and HB⁺, when given the input command “guess”, \mathcal{A}_{clone} takes the full experimental state as input, and outputs a response bit z' .

We presume that a protocol invocation takes some fixed amount of time (as would be the case, for example, in an RFID system). We characterize the total protocol time by three parameters: the number of queries to a \mathcal{T} oracle, q ; the computational runtime t_1 of \mathcal{A}_{query} ; and the computational runtime t_2 of \mathcal{A}_{clone} . Let D be some distribution of $q \times k$ matrices. We let $\stackrel{R}{\leftarrow}$ denote uniform random assignment. Other notation should be clear from context.

<p>Experiment $\mathbf{Exp}_{\mathcal{A},D}^{HB-attack}[k, \eta, q]$</p> <p>$x \xleftarrow{R} \{0, 1\}^k;$</p> <p>$\mathbf{A} \xleftarrow{R} D$</p> <p>$\sigma \leftarrow \mathcal{A}_{query}^{T_{x,\mathbf{A}},\eta};$</p> <p>$\mathbf{a}' \xleftarrow{R} \{0, 1\}^k;$</p> <p>$z' \leftarrow \mathcal{A}_{clone}(\sigma, \mathbf{a}', \text{“guess”});$</p> <p>Output $\mathcal{R}_x(\mathbf{a}', z')$.</p>	<p>Experiment $\mathbf{Exp}_{\mathcal{A}}^{HB^+-attack}[k, \eta, q]$</p> <p>$x, y \xleftarrow{R} \{0, 1\}^k;$</p> <p>$\sigma \leftarrow \mathcal{A}_{query}^{T_{x,y},\eta};$</p> <p>$\mathbf{b}' \leftarrow \mathcal{A}_{clone}(\sigma, \text{“initiate”});$</p> <p>$\mathbf{a}' \xleftarrow{R} \{0, 1\}^k;$</p> <p>$z' \leftarrow \mathcal{A}_{clone}(\sigma, \mathbf{a}', \mathbf{b}', \text{“guess”});$</p> <p>Output $\mathcal{R}_{x,y}(\mathbf{a}', \mathbf{b}', z')$.</p>
---	--

Consider \mathcal{A} 's advantage for key-length k , noise parameter η , over q rounds. In the case of the HB-attack experiment, this advantage will be over matrices \mathbf{A} drawn from the distribution D :

$$\text{Adv}_{\mathcal{A},D}^{HB-attack}(k, \eta, q) = \left| \Pr \left[\mathbf{Exp}_{\mathcal{A},D}^{HB-attack}[k, \eta, q] = \text{“accept”} \right] - \frac{1}{2} \right|$$

Let $\text{Time}(t_1, t_2)$ represent the set of all adversaries \mathcal{A} with runtimes t_1 and t_2 , respectively. Denote the maximum advantage over $\text{Time}(t_1, t_2)$:

$$\text{Adv}_D^{HB-attack}(k, \eta, q, t_1, t_2) = \max_{\mathcal{A} \in \text{Time}(t_1, t_2)} \{ \text{Adv}_{\mathcal{A},D}^{HB-attack}(k, \eta, q) \}$$

The definitions for Adv are exactly analogous for HB^+ -attack, except that there is no input distribution D , as adversarial queries are active.

Note on the model: It is important to point out that this adversarial model is not the strongest possible, as the adversary lacks oracle access to the reader during the query phase. Rather, our experiments specify a “detection model” of anti-counterfeiting. The goal of the adversary in our experiment is to insert a counterfeit tag into the system without detection by the reader. (In contrast, in a “prevention” model, an adversary could not create a counterfeit tag under any circumstances.) We discuss this model and its implications in detail in appendix A.

5.2 Blum et al. Proof Strategy Outline

Given an adversary \mathcal{A} that achieves the advantage $\text{Adv}_{\mathcal{A},U}^{HB-attack}(k, q, \eta, t_1, t_2) = \epsilon$, Blum et al. [3] offer a proof strategy to extract bits of x , and thus solve the LPN problem. If ϵ is a non-negligible function of k , then x can be extracted by their reduction in polynomial time.

To extract the i th bit of the secret x , the Blum et al. reduction works as follows. The reduction takes a given LPN instance (\mathbf{A}, \mathbf{z}) and randomly modifies it to produce a new instance $(\mathbf{A}', \mathbf{z}')$.

The modification involves two steps. First, a vector \mathbf{x}' is chosen uniformly at random and $\mathbf{z}' = (\mathbf{z} \oplus \mathbf{A}) \cdot \mathbf{x}' = (\mathbf{A} \cdot (\mathbf{x} \oplus \mathbf{x}')) \oplus \nu$ is computed. Note

that thanks to the random selection of x' , the vector $(\mathbf{x} \oplus \mathbf{x}')$ is uniformly distributed. Second, the i th column of \mathbf{A} is replaced with random bits. To view this another way, denote the subspace of matrices obtained by uniformly randomizing the i th column of \mathbf{A} as $R_i^{\mathbf{A}}$. The second step of the modification involves setting $\mathbf{A}' \stackrel{R}{\leftarrow} R_i^{\mathbf{A}}$. Once computed as described, the modified problem instance $(\mathbf{A}', \mathbf{z}')$ is fed to an HB adversary \mathcal{A}_{query} .

Suppose that the i th bit of $(\mathbf{x} \oplus \mathbf{x}')$, which we denote $(x \oplus x')_i$, is a binary ‘1’. In this case, since \mathbf{A} is a randomly distributed matrix (because HB challenges are random), and the secret \mathbf{x} is also randomly distributed, the bits of \mathbf{z}' are random. In other words, thanks to the ‘1’ bit, the randomized i th row of \mathbf{A}' “counts” in the computation of \mathbf{z}' , which therefore comes out random. Hence \mathbf{z}' contains no information about the correct value of $\mathbf{A} \cdot (\mathbf{x} \oplus \mathbf{x}')$ or about the secret x . Since \mathcal{A}_{query} cannot pass any meaningful information in σ to \mathcal{A}_{clone} in this case, \mathcal{A}_{clone} can do no better than random guessing of parity bits, and enjoys no advantage.

In contrast, suppose that $(x \oplus x')_i$ is a binary ‘0’. In this case, the i th row of \mathbf{A}' does not “count” in the computation of \mathbf{z}' , and does not have a randomizing effect. Hence \mathbf{z}' may contain meaningful information about the secret \mathbf{x} in this case. As a result, when \mathcal{A}_{clone} shows an advantage over modified problem instances $(\mathbf{A}', \mathbf{z}')$ for a particular fixed choice of \mathbf{x}' , it is clear for those instances that $(x \oplus x')_i = 0$, i.e. $x_i = x'_i$.

In summary then, the Blum et al. reduction involves presentation of suitably modified problem instances $(\mathbf{A}', \mathbf{z}')$ to HB adversary \mathcal{A} . By noting choices of \mathbf{x}' for which \mathcal{A} demonstrates an advantage, it is possible in principle to learn individual bits of the secret \mathbf{x} . With presentation of enough modified problem instances to \mathcal{A} , it is possible to learn \mathbf{x} completely with high probability.

5.3 Reduction from LPN to HB-attack

We will show a concrete reduction from the LPN problem to the HB-attack experiment. This is essentially a concrete version of Blum et al.’s asymptotic reduction strategy from [3] and is an important step in proving Theorem 1.

Unfortunately, the original Blum et al. proof strategy does not account for the fact that while \mathcal{A} ’s advantage may be non-negligible over random matrices, it may actually be negligible over modified $(\mathbf{A}', \mathbf{z}')$ values, i.e., over the distribution $R_i^{\mathbf{A}}$. Matrices are not independent over this distribution: Any two sample matrices are identical in all but one column. Thus, it is possible in principle that \mathcal{A} loses its advantage over this distribution of matrices and that the reduction fails to work. This is a problem that we must remedy here.

We address the problem by modifying a given sample matrix only once. A modified matrix \mathbf{A}' in our reduction is uniformly distributed. This is because it is chosen uniformly from a random $R_i^{\mathbf{A}}$ subspace associated with a random matrix \mathbf{A} . Additionally, since we use a fresh sample for each trial, our modified matrices are necessarily independent of each other. The trade-off is that kL times as many sample matrices are needed for our reduction, where L is the number of trials per bit.

This is an inefficient solution in terms of samples. It is entirely possible that the adversary’s advantage is preserved when, for each column j , samples are drawn from the $R_i^{\mathbf{A}_j}$ subspace for a matrix \mathbf{A}_j . It might even be possible to devise a rigorous reduction that uses a single matrix \mathbf{A} for all columns. We leave these as open questions.

Lemma 1. *Let $\text{Adv}_U^{\text{HB-Attack}}(k, \eta, q, t_1, t_2) = \epsilon$, where U is a uniform distribution over binary matrices $\mathbb{Z}_2^{q \times k}$, and let \mathcal{A} be an adversary that achieves this ϵ -advantage. Then there is an algorithm \mathcal{A}' with running time $t'_1 \leq kLt_1$ and $t'_2 \leq kLt_2$, where $L = \frac{8(\ln k - \ln \ln k)}{(1-2\eta)^2} \left(\frac{1+\epsilon}{\epsilon}\right)^2$, that makes $q' \leq kLq + 1$ queries that can correctly extract all k bits of x with probability $\epsilon' \geq \frac{1}{k}$.*

We provide the proof of this lemma in appendix B.

5.4 Reduction from HB-attack to HB⁺-attack

We show that an HB⁺-attack adversary with ζ -advantage can be used to build an HB-attack adversary with advantage $\frac{\zeta^3(k-2)-2}{4k}$. We provide concrete costs of this reduction that will be used for Theorem 1.

Lemma 3. *If $\text{Adv}_U^{\text{HB}^+-\text{Attack}}(k, \eta, q, t_1, t_2) = \zeta$, then*

$$\text{Adv}_U^{\text{HB-Attack}}(k, \eta, q', t'_1, t'_2) \geq \frac{\zeta^3(k-2)-2}{4k}$$

where $q' \leq q(2 + \log_2 q)$, $t'_1 \leq kq't_1$, $t'_2 \leq 2kt_2$, and $k \geq 9$.

(Lemma 2, a technical lemma, is skipped here. We give the lemma and its proof in appendix C.)

Lemma 3 is the main technical core of the paper. It is worth briefly explaining the proof intuition. The proof naturally involves a simulation where the HB-attack adversary \mathcal{A} makes calls to the furnished HB⁺-attack adversary, which we call \mathcal{A}^+ . In other words, \mathcal{A} simulates the environment for $\text{Exp}_{\mathcal{A}^+}^{\text{HB}^+-\text{attack}}$. The goal of \mathcal{A} is to use \mathcal{A}^+ to compute a correct target response w to an HB challenge vector \mathbf{a} that \mathcal{A} itself receives in an experiment $\text{Exp}_{\mathcal{A}}^{\text{HB-attack}}$.

\mathcal{A} makes its calls to \mathcal{A}^+ in a special way: It “cooks” transcripts obtained from its own HB oracle before passing them to \mathcal{A}^+ during its simulation of the query phase of $\text{Exp}_{\mathcal{A}^+}^{\text{HB}^+-\text{attack}}$. The “cooked” transcripts are such that the target value w is embedded implicitly in a secret bit of the simulated HB⁺ oracle.

In its simulation of the cloning phase of $\text{Exp}_{\mathcal{A}^+}^{\text{HB}^+-\text{attack}}$, the adversary \mathcal{A} extracts the embedded secret bit using a standard cryptographic trick. After \mathcal{A}^+ has committed a blinding value \mathbf{b} , \mathcal{A} rewinds \mathcal{A}^+ to as to make two different challenges $\mathbf{a}^{(0)}$ and $\mathbf{a}^{(1)}$ relative to \mathbf{b} . By looking at the difference in the responses, \mathcal{A} can extract the embedded secret bit and compute its own target response w .

There are two main technical challenges in the proof. The first is finding the right embedding of w in a secret bit of the simulated HB^+ -oracle. Indeed, our approach is somewhat surprising. One might intuitively expect \mathcal{A} instead to cause \mathcal{A}^+ to emit a *response* equal to w during the simulation; after all, w itself is intended to be a tag response furnished by \mathcal{A} , rather than a secret bit. (We could not determine a good way to have w returned as a response.) The second challenge comes in the rewinding and extraction. There is the possibility of a non-uniformity in the responses of \mathcal{A}^+ . An important technical lemma (Lemma 2) is necessary to bound this non-uniformity.

We give proofs for Lemma 3 (and the technical Lemma 2) in appendix 3.

5.5 Reduction of LPN to HB^+ -attack

By combining Lemmas 1 and 3, we obtain a concrete reduction of the LPN problem to the HB^+ -attack experiment. Given an adversary that has an ϵ -advantage against the HB^+ -attack experiment within a specific amount of time and queries, we can construct an adversary that solves the LPN problem within a concrete upper bound of time and queries. The following theorem follows directly from Lemmas 1 and 3.

Theorem 1. *Let $\text{Adv}^{\text{HB}^+ - \text{Attack}}(k, \eta, q, t_1, t_2) = \zeta$, where U is a uniform distribution over binary matrices $\mathbb{Z}_2^{q \times k}$, and let \mathcal{A} be an adversary that achieves this ζ -advantage. Then there is an algorithm that can solve a random $q' \times k$ instance of the LPN problem in time (t'_1, t'_2) with probability $\frac{1}{k}$, where $t'_1 \leq k^2 Lq(2 + \log_2 q)t_1$, $t'_2 \leq 2k^2 Lt_2$, $q' \leq kLq(2 + \log_2 q)$, $\epsilon = \frac{\zeta^{3(k-2)-2}}{4k}$, and $L = \frac{8(\ln k - \ln \ln k)}{(1-2\eta)^2} \left(\frac{1+\epsilon}{\epsilon}\right)^2$.*

To put this in asymptotic terms, the LPN problem may be solved by an adversary where $\text{Adv}^{\text{HB}^+ - \text{Attack}}(k, \eta, q, t_1, t_2) = \zeta$ in time $O\left(\frac{(k^5 \log k)(q \log q) t}{(1-2\eta)^2 \zeta^6}\right)$, where $t = t_1 + t_2$.

6 Conclusion and Open Questions

In summary, this paper presents a new authentication protocol named HB^+ that is appropriate for low-cost pervasive computing devices. The HB^+ protocol is secure in the presence of both passive and active adversaries and should be implementable within the tight resource constraints of today's EPC-type RFID tags.

A number of essential open questions remain, however, before the HB^+ can see practical realization.

Above all, the security of the HB^+ protocol is based on the LPN problem, whose hardness over random instances remains an open question.

The security of concurrent executions of the HB^+ protocol is also unknown. Our security proof uses a rewinding technique that would take time exponential in the number of concurrent rounds. This open question is a vital one,

as we feel it has a bearing on practical realization of our ideas. (It is also open question whether the two-round variant of HB^+ briefly mentioned in Section 4.1 is secure.)

As we have explained, our security model is a “detection model.” Whether our protocols or techniques can be extended to achieve security in stronger adversarial models is an essential line of future work.

Our results here do not offer direct practical guidance for parameterization in real RFID tags, something essential for real-world implementation. It would be desirable to see a much tighter concrete reduction than we give here. One avenue might be improvement to the Blum et al. reduction. As mentioned in Section 5.3, the efficiency of the modified concrete version of Blum et al. reduction [3] may be improved. Our version uses sample values only once. It may be possible to use a single sample to generate several trials per column, or perhaps to generate trials for every column. This lowers the concrete query costs. It is unclear, however, whether the reduction holds over non-uniform input distributions.

Finally, there is second human authentication protocol by Hopper and Blum, based on the “Sum of k Mins” problem and error-correcting challenges [5, 19]. Unlike the HB protocol, this protocol is already supposed to be secure against active adversaries. However, the hardness of the “Sum of k Mins” has not been studied as much as the LPN problem, nor is it clear whether this protocol can efficiently be adapted for low-cost devices. These remain open avenues of research.

Acknowledgments

Thanks to Mike Szydlo for contributing ideas and references.

References

- [1] ANDERSON, R., AND KUHN, M. Low Cost Attacks on Tamper Resistant Devices. In *International Workshop on Security Protocols* (1997), vol. 1361 of *Lecture Notes in Computer Science*, pp. 125–136.
- [2] BERLEKAMP, E. R., MCELIECE, R. J., AND TILBORG, V. On the Inherent Intractability of Certain Coding Problems. *IEEE Transactions on Information Theory* 24 (1978), 384–386.
- [3] BLUM, A., FURST, M., KEARNS, M., AND LIPTON, R. J. Cryptographic Primitives Based on Hard Learning Problems. In *Advances in Cryptology – CRYPTO’93* (1994), vol. 773 of *Lecture Notes in Computer Science*, pp. 278–291.
- [4] BLUM, A., KALAI, A., AND WASSERMAN, H. Noise-Tolerant Learning, the Parity Problem, and the Statistical Query Model. *Journal of the ACM* 50, 4 (July 2003), 506–519.
- [5] BLUM, M., LUBY, M., AND RUBINFELD, R. Self-Testing/Correcting with Applications to Numerical Problems. In *Symposium on Theory of Computation* (1990), pp. 73–83.
- [6] BONO, S., GREEN, M., STUBBLEFIELD, A., JUELS, A., RUBIN, A., AND SZYDLO, M. Security Analysis of a Cryptographically-Enabled RFID Device. In *USENIX Security* (2005). To appear. Available at <http://rfidanalysis.org/>.

- [7] CHABAUD, F. On the Security of Some Cryptosystems Based on Error-Correcting Codes. In *Advances in Cryptology - EUROCRYPT* (1995), vol. 950 of *Lecture Notes in Computer Science*, pp. 131–139.
- [8] COURTOIS, N., FINIASZ, M., AND SENDRIER, N. How to Achieve a McEliece-based Digital Signature Scheme. In *Advances in Cryptology - ASIACRYPT* (2001), vol. 2248 of *Lecture Notes in Computer Science*, pp. 157–174.
- [9] CRAWFORD, J. M., KEARNS, M. J., AND SHAPIRE, R. E. The Minimal Disagreement Parity Problem as a Hard Satisfiability Problem. Tech. rep., Computational Intelligence Research Laboratory and AT&T Bell Labs, February 1994.
- [10] EPCGLOBAL. Website. <http://www.epcglobalinc.org/>, 2005.
- [11] FELDHOFFER, M., DOMINIKUS, S., AND WOLKERSTORFER, J. Strong Authentication for RFID Systems using the AES Algorithm. In *Cryptographic Hardware in Embedded Systems (CHES)* (2004).
- [12] FLOERKEMEIER, C., AND LAMPE, M. Issues with RFID Usage in Ubiquitous Computing Applications. In *Pervasive Computing (PERVASIVE)* (2004), vol. 3001 of *Lecture Notes in Computer Science*, pp. 188–193.
- [13] FOOD AND DRUG ADMINISTRATION. Combating counterfeit drugs. Tech. rep., US Department of Health and Human Services, Rockville, Maryland, February 2004.
- [14] GILBERT, H., SIBERT, H., AND ROBshaw, M. An Active Attack Against a Provably Secure Lightweight Authentication Protocol. Preliminary Version, 2005.
- [15] GREENTECH COMPUTING. GT6 Algorithm Solves the Extended DIMACS 32-bit Parity Problem. Tech. rep., Greentech Computing Limited, London, England, 1998.
- [16] HÅSTAD, J. Some Optimal Inapproximability Results. In *Symposium on Theory of Computing* (1997), pp. 1–10.
- [17] HENRICI, D., AND MÜLLER, P. Hash-based Enhancement of Location Privacy for Radio-Frequency Identification Devices using Varying Identifiers. In *Pervasive Computing and Communications (PerCom)* (2004), IEEE Computer Society, pp. 149–153.
- [18] HOPPER, N., AND BLUM, M. A Secure Human-Computer Authentication Scheme. Tech. Rep. CMU-CS-00-139, Carnegie Mellon University, 2000.
- [19] HOPPER, N. J., AND BLUM, M. Secure Human Identification Protocols. In *Advances in Cryptology - ASIACRYPT* (2001), vol. 2248 of *Lecture Notes in Computer Science*, pp. 52–66.
- [20] JOHNSON, D. S., AND TRICK, M. A., Eds. *Cliques, Coloring, and Satisfiability: Second Dimacs Implementation Challenge* (1993), vol. 26, American Mathematical Society.
- [21] JUELS, A. Minimalist Cryptography for RFID Tags. In *Security in Communication Networks* (2004), C. Blundo and S. Cimato, Eds., vol. 3352 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 149–164.
- [22] JUELS, A. "Yoking Proofs" for RFID Tags. In *Pervasive Computing and Communications Workshop* (2004), IEEE Press.
- [23] JUELS, A., AND PAPPU, R. Squealing Euros: Privacy Protection in RFID-Enabled Banknotes. In *Financial Cryptography* (2003), vol. 2742 of *Lecture Notes in Computer Science*, pp. 103–121.

- [24] JUELS, A., RIVEST, R. L., AND SZYDLO, M. The blocker tag: selective blocking of RFID tags for consumer privacy. In *Proceedings of the 10th ACM conference on Computer and communication security* (2003), ACM Press, pp. 103–111.
- [25] KEARNS, M. Efficient Noise-Tolerant Learning from Statistical Queries. *Journal of the ACM* 45, 6 (November 1998), 983–1006.
- [26] MACWILLIAMS, F., AND SLOANE, N. *The Theory of Error-Correcting Codes*. North-Holland, Amsterdam, 1977.
- [27] MANDEL, J., ROACH, A., AND WINSTEIN, K. MIT Proximity Card Vulnerabilities. Tech. rep., Massachusetts Institute of Technology, March 2004.
- [28] MATSUMOTO, T. Human-computer Cryptography: An Attempt. In *Computer and Communications Security* (1996), ACM Press, pp. 68–75.
- [29] MATSUMOTO, T., AND IMAI, H. Human Identification through Insecure Channel. In *Advances in Cryptology - EUROCRYPT* (1991), vol. 547 of *Lecture Notes in Computer Science*, pp. 409–421.
- [30] McELIECE, R. J. DSN Progress Report. Tech. Rep. 42–44, JPL-Caltech, 1978.
- [31] MILLER, G. A. The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. *Psychological Review* 63 (1956), 81–97.
- [32] MOLNAR, D., AND WAGNER, D. Privacy and Security in Library RFID : Issues, Practices, and Architectures. In *Computer and Communications Security* (2004), B. Pfizmann and P. McDaniel, Eds., ACM, pp. 210 – 219.
- [33] NAOR, M., AND PINKAS, B. Visual Authentication and Identification. In *Advances in Cryptology - CRYPTO* (1997), vol. 1294 of *Lecture Notes in Computer Science*, pp. 322–336.
- [34] NIEDERREITER, H. Knapsack-Type Cryptosystems and Algebraic Coding Theory. *Problems of Control and Information Theory* 15, 2 (1986), 159–166.
- [35] OHKUBO, M., SUZUKI, K., AND KINOSHITA, S. Efficient Hash-Chain Based RFID Privacy Protection Scheme. In *Ubiquitous Computing (UBICOMP)* (September 2004).
- [36] SARMA, S. E., WEIS, S. A., AND ENGELS, D. W. RFID Systems and Security and Privacy Implications. In *Workshop on Cryptographic Hardware and Embedded Systems* (2002), vol. 2523, *Lecture Notes in Computer Science*, pp. 454–470.
- [37] SHAMOS, M. I. Paper v. Electronic Voting Records - An Assessment, 2004. Paper written to accompany panel presentation at Computers, Freedom, and Privacy Conference '04. Available at <http://euro.ecom.cmu.edu/people/faculty/mshamos/paper.htm>.
- [38] STERN, J. A New Paradigm for Public Key Identification. *IEEE Transactions on Information Theory* 42, 6 (1996), 1757–1768.
- [39] VAJDA, I., AND BUTTYAN, L. Lightweight Authentication Protocols for Low-Cost RFID Tags. In *Ubiquitous Computing (UBICOMP)* (2003).
- [40] VERICHIP. Website. <http://www.4verichip.com/>, 2005.
- [41] WANG, C.-H., HWANG, T., AND TSAI, J.-J. On the Matsumoto and Imai’s Human Identification Scheme. In *EuroCrypt '95* (1995), vol. 921 of *Lecture Notes in Computer Science*, pp. 382–392.
- [42] WARNERS, J. P., AND VAN MAAREN, H. A Two Phase Algorithm for Solving a Class of Hard Satisfiability Problems. *Operations Research Letters* 23, 3–5 (1999), 81–88.
- [43] WEIS, S. A., SARMA, S. E., RIVEST, R. L., AND ENGELS, D. W. Security and Privacy Aspects of Low-Cost Radio Frequency Identification Systems. In

Security in Pervasive Computing (2004), vol. 2802 of *Lecture Notes in Computer Science*, pp. 201–212.

A Discussion of Security Model

The security model we consider in this paper for HB^+ is not the strongest possible; we circumscribe the power of the adversary. It will be noted that during the query phase of the experiment HB^+ -attack, the adversary lacks oracle access to the reader \mathcal{R} .

This weakened model merits some explanation. It may be viewed as defining a *detection-based* authentication system in which the adversary is presumed to have a particular aim: The adversary seeks to insert a bogus tag into the system without detection. In other words, if an authentication session fails, and \mathcal{R} thus detects an ostensibly counterfeit tag, we consider the adversary to have been unsuccessful. (This is equivalent to saying that during the query phase of HB^+ -attack, the adversary can only initiate or observe successful authentication sessions. Such sessions reveal only information that the adversary can learn directly from the tag, i.e., that access to \mathcal{R} furnishes no additional information.)

It is instructive to compare a detection-based model against a *prevention-based* model where the adversary has unfettered access to oracles for \mathcal{T} and \mathcal{R} . The aim in a prevention-based model is to ensure against tag cloning irrespective of whether or not an adversary is detected in a counterfeiting attempt against a tag.

We consider a detection-based model to be natural and useful in centralized RFID systems, such as those that might be employed with RFID-tagged casino chips or proximity cards, or in tightly integrated supply chains. In such environments, a failed authentication attempt – such as an attempt at counterfeiting – would naturally trigger an alert. RFID tags have an important physical dimension, namely that an attacker must have some physical presence or proxy to mount an attack. Thus detection has a value in RFID systems not present in general communication networks where an attacker may operate remotely.

To ensure against leakage of tag secrets, a detection-based authentication system can employ throttling or a lockout in the face of multiple failed authentication attempts. (Of course, any such policy must be constructed carefully to account for the possibility of denial-of-service attacks.) Each authentication attempt can in principle leak up to at most a single bit of information about the secret contained in a tag, as a reader either accepts or rejects at the conclusion of a session.

Gilbert, Robshaw, and Sibert [14] have recently observed that our detection-based model in this work is vulnerable to active attack, i.e., it does not achieve the full strength of a prevention-based model. They demonstrate a simple, linear-time man-in-the-middle attack against HB^+ .

Of course, a prevention-based model is more desirable than a detection-based model. Our aim here, however, is to propose a lightweight system for pragmatic security in very low-cost wireless devices. Low-cost RFID tags, as we have explained, for instance, possess extremely too little computational power to execute canonical symmetric-key cryptographic primitives. Much of the literature on low-cost RFID seeks to achieve the pragmatic aim of security in models that involve some form of weakening, e.g., [21].

We speculate that the LPN-based techniques we employ here cannot achieve security in a prevention-based model without a very large increase in parameter sizes. Whether or not there exists an some efficient, prevention-based HB^{++} protocol is an open question.

B Proof of LPN to HB-attack Reduction

Lemma 1. *Let $\text{Adv}_U^{HB-Attack}(k, \eta, q, t_1, t_2) = \epsilon$, where U is a uniform distribution over binary matrices $\mathbb{Z}_2^{q \times k}$, and let \mathcal{A} be an adversary that achieves this ϵ -advantage. Then there is an algorithm \mathcal{A}' with running time $t'_1 \leq kLt_1$ and $t'_2 \leq kLt_2$, where $L = \frac{8(\ln k - \ln \ln k)(1+\epsilon)^2}{(1-2\eta)^2 \epsilon^2}$, that makes $q' \leq kLq + 1$ queries that can correctly extract all k bits of x with probability $\epsilon' \geq \frac{1}{k}$.*

Given an adversary \mathcal{A} such that $\text{Adv}_{\mathcal{A}, U}^{HB-attack}(k, q, \eta, t_1, t_2) = \epsilon$, we will show how to construct a simulator \mathcal{S} to extract all bits of an HB secret \mathbf{x} with high probability. Thus, \mathcal{S} will be able solve the LPN problem.

Consider a particular set of samples $(\hat{\mathbf{A}}, \hat{\mathbf{z}})$. The simulator \mathcal{S} will first select one sample row from $(\hat{\mathbf{A}}, \hat{\mathbf{z}})$ at random and denote it $(\hat{\mathbf{a}}, \hat{z})$. \mathcal{S} splits the remaining samples into kL sets of size q (L will be defined later). The simulator will then replace a random i th column of L different samples with random bits and randomized the associated z value as described in Section 5.2. Denote these samples as $(\mathbf{A}', \mathbf{z}')$, respectively.

\mathcal{S} will then input each $(\mathbf{A}', \mathbf{z}')$ sample to \mathcal{A}_{query} . In the cloning phase, \mathcal{S} replaces the i th bit of $\hat{\mathbf{a}}$ with a random bit and challenges \mathcal{A}_{clone} for the result. The simulator \mathcal{S} knows the noisy sample \hat{z} , thus can verify whether \mathcal{A}_{clone} 's result matches.

If $(x \oplus x')_i = 0$, then replacing the i th column of \mathbf{A}' does not affect \mathbf{z}' and $(\mathbf{A}', \mathbf{z}')$ is a valid LPN instance. Our hope is that \mathcal{A} would maintain its ϵ -advantage over this distribution of samples. However, it is conceivable that the adversary \mathcal{A} 's advantage over this is less than ϵ over the distribution of samples whose i th secret bit is necessarily zero.

However, since the samples are drawn from a valid LPN distribution, the adversary must still maintain an ϵ -advantage over all secrets. Thus, any under-performance over the distribution over '0'-valued i th bits is made up over the distribution of '1'-valued i th bits.

Denote the event that the i th secret bit is zero as Z , and when it is one as \bar{Z} . Suppose $Pr[\mathcal{A} \text{ succeeds} \mid Z] = (\epsilon - \delta)$ and $Pr[\mathcal{A} \text{ succeeds} \mid \bar{Z}] = (\epsilon + \delta)$. If δ is significantly large enough, then \mathcal{S} can simply run \mathcal{A} on the original, unaltered samples $\hat{\mathbf{A}}$ and observe its performance. An adversary that achieves advantage less than ϵ would indicate Z , while advantage greater than ϵ would indicate \bar{Z} . Note that a significant δ can be detected by generating random LPN instances and measuring \mathcal{A} 's performance conditioned on Z and \bar{Z} events.

Using this naïve approach, \mathcal{A} correctly outputs the i th bit when either both \hat{z} and \mathcal{A} are correct, or they are both wrong (recall that \hat{z} is a noisy sample).

This occurs with probability $(1 - \eta)(\frac{1}{2} + \frac{\delta}{2\epsilon}) + \eta(\frac{1}{2} - \frac{\delta}{2\epsilon})$. Thus \mathcal{A} would guess x_i with expected probability $\frac{1}{2} + \frac{1}{2}(1 - 2\eta)\frac{\delta}{\epsilon}$.

However, if we use the modified \mathbf{A}' samples, the \mathcal{A} will correctly guess x_i with expected probability $\frac{1}{2} + \frac{1}{2}(1 - 2\eta)(\epsilon - \delta)$. Thus, if $\frac{\delta}{\epsilon} \geq (\epsilon - \delta)$, it would be advantageous for \mathcal{S} to use the naïve method. This occurs if $\delta \geq \frac{\epsilon^2}{1+\epsilon}$.

If this is the case, then there exists a simulator that simply permutes columns of the original challenges $\widehat{\mathbf{A}}$ and maintains alignment with the corresponding rows of $\widehat{\mathbf{z}}$. Then a simulator can run the naïve attack to determine each key bit with advantage $\frac{1}{2} + \frac{1}{2}(1 - 2\eta)\frac{\delta}{\epsilon}$ per trial. The simulator can permute columns such that each key bit is assigned to the i th column exactly L times. This assumes the original samples are drawn from a random distribution, which is closed under permutation.

Thus, in the worst case we have that $\delta = \frac{\epsilon^2}{1+\epsilon}$ and $(\epsilon - \delta) = \frac{\epsilon}{1+\epsilon}$. For convenience, denote \mathcal{A} 's advantage at guessing a bit per trial as $\widehat{\epsilon} = \frac{1}{2}(1 - 2\eta)(\frac{\epsilon}{1+\epsilon})$.

Consider repeating L randomly modified trials per k bits of \mathbf{x} and taking the majority of the outcome for each bit. By a Chernoff bound, after L trials each guessed bit will be correct with probability $p = (1 - \exp(\frac{-L\widehat{\epsilon}^2}{1+2\widehat{\epsilon}})) \leq (1 - \exp(\frac{-L\widehat{\epsilon}^2}{2}))$. Thus, all k bits will be correct with probability p^k .

Thus, if \mathcal{A} requires q samples and runs in (t_1, t_2) time, this reduction will extract all bits of x using $q' = kLq + 1$ samples and running in time $t' = tLk$ with probability:

$$(1 - e^{-\frac{L\widehat{\epsilon}^2}{2}})^k \approx \exp\left(\frac{-k}{\exp(\frac{L\widehat{\epsilon}^2}{2})}\right)$$

Let $L = \frac{2(\ln k - \ln \ln k)}{\widehat{\epsilon}^2}$. Plugging this into the above formula gives us a success probability of $\frac{1}{k}$. Substituting in for $\widehat{\epsilon}$, we get that $L = 8(\ln k - \ln \ln k) \left(\frac{1+\epsilon}{(1-2\eta)\epsilon}\right)^2$. Thus, we can express $t'_1 = t_1 Lk$ and $t'_2 = t_2 Lk$ concretely in terms of k and η . \square

C HB-attack to HB⁺-attack Reduction

The following lemma is a technical one. It bounds the ability of an adversary to cause failures in our simulation in Lemma 3 in a step where we provide challenge vectors that differ in a random bit position. Here we let $v[j]$ denote the j^{th} bit of a vector v , and let \in_U denote uniform random selection from a set.

Lemma 2. *Consider an experiment that takes as input a matrix \mathbf{A} and a k -bit vector, where $k \geq 9$. The experiment yields either a '0' or a '1' output. Let $p_{\mathbf{A}}$ denote the probability of a '1' output over random vectors of k bits for a matrix \mathbf{A} . Suppose a pair of random k -bit vectors v_0 and v_1 is selected such that $v_0[j] = 0$ and $v_1[j] = 1$ for random $j \in \{1, \dots, k\}$. Let $q_{\mathbf{A}}$ be the probability that for vectors thus selected, both yield a '0' or both vectors yield a '1'. If $p = \sum_{\mathbf{A}} p_{\mathbf{A}} \geq 1/2 + \epsilon$, then $q = \sum_{\mathbf{A}} q_{\mathbf{A}} \geq 1/2 + \epsilon'$ for $\epsilon' = \epsilon^3/2 - (\epsilon^3 + 1)/k$.*

Proof. Suppose that v_0 and v_1 are selected as in the statement of the Lemma 2, i.e., with a ‘0’ and ‘1’ fixed respectively at a random position j . We observe that for a set S of k -bit vectors such that $|S| = 2^{k-d}$, $\Pr[v_0, v_1 \in S]$ is minimized when S consists of vectors whose first d bits are equal. Consequently, for $|S| > 2^{k-d}$, we have

$$\Pr[u_0[j] = u_1[j] \mid u_0, u_1 \in_U S, j \in_U \{1, \dots, k\}] \geq (k-d)/k. \quad (1)$$

Here \in_U denotes uniform random selection.

For a particular matrix \mathbf{A} , there is a set $S_{\mathbf{A}}$ of vectors for which the experiment outputs ‘1’, where $|S_{\mathbf{A}}| = p_{\mathbf{A}}2^k$. For clarity, we drop the subscript and denote this set by S . Let \bar{S} denote the complementary set. We shall also assume $j \in_U \{1, \dots, k\}$ as appropriate in what follows.

By Bayes’s rule and eq. 1, we have $\Pr[v_0, v_1 \in S \mid v_0[j] = 0, v_1[j] = 1] = \Pr[(v_0, v_1 \in S) \wedge (v_0[j] = 0, v_1[j] = 1)] / \Pr[v_0[j] = 0, v_1[j] = 1] \geq p_{\mathbf{A}}^2 \binom{k - \lceil \log_2 p_{\mathbf{A}} \rceil}{k}$.

Now consider two cases for a particular matrix \mathbf{A} .

Case 1, $p_{\mathbf{A}} \leq 1/4$: In this case, $q_{\mathbf{A}} \geq \Pr[v_0, v_1 \in \bar{S} \mid v_0[j] = 0, v_1[j] = 1] > (3/4)^2(k-1/k)$. As $k \geq 9$, it follows that $q_{\mathbf{A}} \geq 1/2$.

Case 2, $p_{\mathbf{A}} > 1/4$: Here, $q_{\mathbf{A}} = \Pr[v_0, v_1 \in S \mid v_0[j] = 0, v_1[j] = 1] + \Pr[v_0, v_1 \in \bar{S} \mid v_0[j] = 0, v_1[j] = 1] > p_{\mathbf{A}}^2 + (1-p_{\mathbf{A}})^2(k-2/k) = (2p_{\mathbf{A}}^2 - 2p_{\mathbf{A}} + 1)(k-2/k)$. Note that this last term is minimized for $p_{\mathbf{A}} = 1/2$, in which case $q_{\mathbf{A}} = (k-2)/2k$.

Since $p = 1/2 + \epsilon$, it is straightforward to show that $p_{\mathbf{A}} \geq 1/2 + \epsilon/2$ for greater than an ϵ -fraction of matrices \mathbf{A} . For such matrices, $q_{\mathbf{A}} > (2p_{\mathbf{A}}^2 - 2p_{\mathbf{A}} + 1)(k-2/k) = (1/2 + \epsilon^2/2) \binom{k-2}{k}$.

Thus, $q = \sum_{\mathbf{A}} q_{\mathbf{A}} > (1-\epsilon) \binom{k-2}{2k} + \epsilon(1/2 + \epsilon^2/2) \binom{k-2}{2k} = \binom{\epsilon^3+1}{2} \binom{k-2}{k}$. The lemma then follows by straightforward algebraic manipulation. \square

The next lemma is our main result in this paper, namely a reduction of the security of our HB^+ protocol to the security of the HB protocol. We note that the lemma is only meaningful for relatively large advantage values ζ . Small advantages, however, can be boosted through the standard technique of taking majority output from multiple (polynomial) adversarial executions.

Lemma 3. *If $\text{Adv}_U^{HB^+ - \text{Attack}}(k, \eta, q, t_1, t_2) = \zeta$, then*

$$\text{Adv}_U^{HB - \text{Attack}}(k, \eta, q', t'_1, t'_2) \geq \frac{\zeta^3(k-2) - 2}{4k}$$

where $q' \leq q(2 + \log_2 q)$, $t'_1 \leq kq't_1$, $t'_2 \leq 2kt_2$, and $k \geq 9$.

Suppose we are given an HB^+ adversary \mathcal{A}^+ that achieves advantage $\text{Adv}_{\mathcal{A}^+, U}^{HB^+ - \text{attack}}(k, \eta, q, t_1, t_2) = \zeta$, where ζ is non-negligible in k . We shall use this adversary to construct an HB adversary \mathcal{A} . As an HB adversary, \mathcal{A} queries a

tag oracle $\mathcal{T}_{x,\mathcal{A},\eta}$ in the query phase of $\mathbf{Exp}_{\mathcal{A}}^{HB-attack}$. We denote the challenge-response pairs from this phase by $(\mathbf{A}, w) = \{a^{(i)}, w^{(i)}\}_{i=1}^{q+rq}$. (Note that we assume a “pool” here of rq extra challenge-response pairs.) In the cloning phase of $\mathbf{Exp}_{\mathcal{A}}^{HB-attack}$, \mathcal{A} takes a challenge vector \mathbf{a} and aims to output a correct response $w = \mathbf{a} \cdot \mathbf{x}$. To accomplish this goal and determine the target value w , \mathcal{A} makes specially formulated calls to the adversary \mathcal{A}^+ between its experimental phases, as we now explain.

In its calls to \mathcal{A}^+ during the simulated query phase, the adversary \mathcal{A} simulates responses for an HB^+ tag oracle $\mathcal{T}_{x^+,y^+,\eta}$. To do so, it takes responses from its own tag oracle $\mathcal{T}_{x,\mathcal{A},\eta}$ and “folds in” its own k -bit secret \mathbf{s} before passing them to \mathcal{A}^+ . In effect, \mathcal{A} uses \mathbf{s} as the tag oracle secret \mathbf{x}^+ : Since the challenges $\{\mathbf{a}^{+(i)}\}$ that are XORed with \mathbf{x}^+ are selected actively by \mathcal{A}^+ , the adversary \mathcal{A} must have knowledge of \mathbf{x}^+ in order to perform the simulation successfully. In contrast, \mathcal{A} itself chooses the blinding factors $\{\mathbf{b}^{+(i)}\}$ that are XORed with \mathbf{y}^+ in the query phase. Therefore, \mathcal{A} is able to perform its simulation with $\mathbf{y}^+ = \mathbf{x}$, i.e., since it controls the challenges, it can incorporate the data (\mathbf{A}, w) here that it harvested (passively) in $\mathbf{Exp}_{\mathcal{A}}^{HB-attack}$.

Let $s[i]$ denote the i^{th} bit of \mathbf{s} . The adversary \mathcal{A} selects all bits of the secret \mathbf{s} at random *except* $s[j]$. It reserves the bit $s[j]$ as a special unknown one; in its simulation with \mathcal{A}^+ , it implicitly embeds the target value w in $s[j]$, as we shall explain.

Let us now describe how \mathcal{A} executes the query and cloning phases for \mathcal{A}^+ in its simulation of $\mathbf{Exp}_{\mathcal{A}^+}^{HB^+-attack}$.

Query phase: Recall that in this phase, \mathcal{A}_{query}^+ queries an HB^+ tag oracle $\mathcal{T}_{x^+,y^+,\eta}$. Let us consider the m^{th} query made by \mathcal{A}_{query}^+ , which we denote by $\mathbf{a}^{+(m)}$. Before the query is made, \mathcal{A} selects a random bit $g^{(m)}$. This is \mathcal{A} 's guess at the query bit $a^{+(m)}[j]$. If $g^{(m)} = 0$, then \mathcal{A} sets $\mathbf{b}^{+(m)} = \mathbf{a}^{(m)}$. If $g^{(m)} = 1$, it sets $\mathbf{b}^{+(m)} = \mathbf{a}^{(m)} \oplus \mathbf{a}$. \mathcal{A} passes the blinding factor $\mathbf{b}^{+(m)}$ to \mathcal{A}^+ as the first protocol flow.

If \mathcal{A} 's guess $g^{(m)}$ is incorrect, i.e., $g^{(m)} \neq a^{+(m)}[j]$, then \mathcal{A} rewinds to the beginning of the m^{th} query. It discards the pair $(\mathbf{a}^{(m)}, w^{(m)})$ from (\mathbf{A}, w) and replaces it with the next challenge-response pair. In effect, \mathcal{A} draws from the “pool” of extra challenge-response pairs in (\mathbf{A}, w) . It halts and outputs a random guess at w if the “pool” is exhausted. \mathcal{A} then repeats its simulation for the m^{th} query with a new guess $g^{(i)}$ and the new challenge-response pair.

If \mathcal{A} 's guess $g^{(i)}$ is correct, then \mathcal{A} computes its response bit as $z^{+(m)} = \bigoplus_{i \neq j} (a^{+(m)}[i]s[i]) \oplus w^{(m)}$. If $g^{(m)} = a^{+(m)}[j] = 0$, then observe that there is an omitted term $u = a^{+(m)}[j]s[j]$ in this response bit; since $a^{+(m)}[j] = 0$, this omitted value $u = 0$, so the response $z^{+(m)}$ is still correct. If $g^{(m)} = a^{+(m)}[j] = 1$, then the omitted term $u = a^{+(m)}[j]s[j] \oplus w = s[j] \oplus w$. In other words, the response is correct if and only if $s[j] = w$. This is how \mathcal{A} embeds the target value w in the secret bit $s[j]$ (without knowing w).

\mathcal{A} noises its response according to probability value η before transmitting it to \mathcal{A}^+ .

Cloning phase: In the cloning phase, the goal of \mathcal{A} is to extract the target value $s[j] = w$ from \mathcal{A}^+ . In this phase, recall that \mathcal{A}_{clone}^+ attempts to simulate the oracle $\mathcal{T}_{x^+, y^+, \eta}$. Thus, \mathcal{A}^+ first outputs a blinding factor, which we denote by $\hat{\mathbf{b}}$; then \mathcal{A} provides a challenge value $\hat{\mathbf{a}}$. Finally, \mathcal{A}^+ outputs a response bit \hat{z} . If correct, the value $\hat{z} = (\hat{\mathbf{a}} \cdot \mathbf{x}^+) \oplus (\hat{\mathbf{b}} \cdot \mathbf{y}^+) = (\hat{\mathbf{a}} \cdot \mathbf{s}) \oplus (\hat{\mathbf{b}} \cdot \mathbf{x})$.

\mathcal{A} selects a random pair of challenge values $(\hat{\mathbf{a}}^0, \hat{\mathbf{a}}^1)$. It selects these such that they differ in the j^{th} bit. (Assume w.l.o.g. that $\hat{a}^{(0)}[j] = 0$ and $\hat{a}^{(1)}[j] = 1$.) \mathcal{A} then initiates an interaction with \mathcal{A}^+ . It receives the blinding factor $\hat{\mathbf{b}}$. It then transmits challenge $\hat{\mathbf{a}}^{(0)}$, receiving response bit $\hat{z}^{(0)}$ from \mathcal{A}^+ . It rewinds \mathcal{A}^+ and likewise transmits challenge $\hat{\mathbf{a}}^{(1)}$ to get response bit $\hat{z}^{(1)}$.

Suppose that both $\hat{z}^{(0)}$ and $\hat{z}^{(1)}$ are correct. Then $\hat{z}^{(0)} \oplus \hat{z}^{(1)} = \hat{\mathbf{a}}^{(0)} \cdot \mathbf{s} \oplus \hat{\mathbf{a}}^{(1)} \cdot \mathbf{s} =$

$$\left(\sum_{i \neq j} (\hat{a}^{(0)}[i] \oplus \hat{a}^{(1)}[i]) s[i] \right) \oplus s[j].$$

Since \mathcal{A} knows all bits of s except $s[j]$, it can compute the first term here, and thus the target value $w = s[j]$. If both responses $\hat{z}^{(0)}$ and $\hat{z}^{(1)}$ are *incorrect*, the same computation works: The errors will cancel out.

We must now ask the probability, given that $\hat{\mathbf{a}}^{(0)}$ and $\hat{\mathbf{a}}^{(1)}$ differ in the j^{th} bit, that they yield like responses. In other words, we want to determine the probability that $\hat{z}^{(0)}$ and $\hat{z}^{(1)}$ are simultaneously either correct or incorrect. Let Z_0 and Z_1 be random variables, where for $d \in \{0, 1\}$, we have $Z_d = 1$ if $\hat{z}^{(d)}$ is correct, and vice versa. Thus we want to compute $\Pr[Z_0 = Z_1]$. Since the adversary has no way of knowing j in the course of the simulation, we can suppose in computing this probability that we select j *a posteriori*, i.e., after the cloning phase. It is important to note, however, that Z_0 and Z_1 are not identically distributed. In particular, the responses of the adversary \mathcal{A}^+ are *conditioned on the fact that $\hat{\mathbf{a}}^{(0)}$ and $\hat{\mathbf{a}}^{(1)}$ differ in a single bit*.

For this reason, we must invoke our technical Lemma 2, which bounds the effect of this conditioning. Recall that $\text{Adv}_{\mathcal{A}^+}^{HB\text{-}attack}(k, \eta, q, t_1, t_2) = \zeta$ by assumption; thus, the $p = 1/2 + \zeta$ in our lemma. Hence, for $k \geq 9$, we have that

$$\Pr[Z_0 = Z_1] \geq 1/2 + \frac{\zeta^3}{2} - \frac{\zeta^3 + 1}{k}. \quad (2)$$

We must also compute the probability that our simulation halts. This can happen if rewinding fails, i.e., all of the extra challenge-response pairs in the “pool” are used up. For simplicity, we can bound this above by $q2^{-r}$, namely the probability that any one rewinding results in the discarding of r pairs in the “pool.”, where $r = \log_2 q + 1$. Let us set $r = (\log_2 q + 1)$. It follows that we can bound the halting probability above by $q2^{-r} = q2^{-(\log_2 q + 1)} = 1/2$.

Given this bound and eq. 2, we obtain

$$\text{Adv}_{\mathcal{A}^+}^{HB^+ \text{-}Attack}(k, \eta, q, t'_1, t'_2) \geq 1/2 + \frac{\zeta^3}{4} - \frac{\zeta^3 + 1}{2k},$$

for $t'_1 = kt_1q(2 + \log_2 q)$ and $t'_2 = 2kt_2$. These runtimes are due to the fact that \mathcal{A}_{query}^+ may have to do r rewinds for each of k bits. We'll upper bound the cost of each "rewind" with the cost of a complete invocation of \mathcal{A}_{query} . Similarly, \mathcal{A}_{clone}^+ will run two copies of \mathcal{A}_{clone} for each of k bits. Note that to achieve a positive advantage in the reduction, we need $\zeta^3 > \frac{2}{k-2}$. (When the advantage ζ is small, however, we can boost it using the standard technique of executing the \mathcal{A}^+ multiple times and taking the majority output.) \square

D Lower Bounds on Key Sizes

Similar to security constructions based on factoring or finding discrete logarithms, the length of keys that are secure in practice will depend on the state of the art of algorithms and hardware. As a baseline for comparison, in 1993 DIMACS issued a set of random LPN instances reduced to CNF formulas as a satisfiability problem challenge [20]. These challenge problems used a key length of $k = 32$, $q = 64$ queries, and a noise parameter of $\eta = \frac{1}{8}$.

Solutions were found several years later by specialized exhaustive search algorithms [15, 42]. As a measure of practical hardness, each instance of the DIMACS challenge [20] took approximately 5-10 minutes to solve on a 200 MHz SGI R10k processor using Warners and van Maaren's search algorithm [42]. Although this search algorithm does not necessarily find the same \mathbf{x} used to generate the responses, it is clear that with a key as short as 32 bits, a small set of samples can be trivially broken on a home PC.

We will concretely analyze the runtime of the best known LPN learning algorithm due to Blum, Kalai, and Wasserman (BKW) [4]. As mentioned, this algorithm requires a runtime of $2^{O(\frac{k}{\log \kappa})}$. The BKW algorithm essentially performs Gaussian elimination on a large set of noisy samples, except tries to minimize the number of linear combinations. That minimizes the noise accumulated throughout the algorithm. By repeating randomized trials, the BKW algorithm can produce the secret \mathbf{x} with high probability.

Omitting details of their algorithm, for $\alpha\beta \geq k$, given $q = \alpha^3 m 2^\beta$ queries and running in $t = C\alpha^3 m 2^\beta$ time, where $m = \max\left\{\left(\frac{1}{1-2\eta}\right)^{2\alpha}, \beta\right\}$, the BKW algorithm can correctly extract \mathbf{x} with an error that is negligible in k . In other words:

$$\text{Adv}_{BKW}^{\text{extract}-\mathbf{x}}(k, \eta, q, t) \approx 1 - \text{negl}(k)$$

Suppose $\eta = 1/4$. Then $m = 2^{2\alpha}$. If we let $k = 224$, then the values $\alpha = 4$ and $\beta = 58$ minimize the value of t necessary to completely reveal a 224-bit secret x with high probability. For these values of α and β we have that $C\alpha^3 2^{2\alpha+\beta} \geq 2^{80}$. Thus, a 224-bit LPN secret x with noise parameter $\eta = 1/4$ is secure against an adversary running the improved BKW algorithm that can run for 2^{80} steps.

Just for comparison, if $k = 32$, as in the DIMACS challenge, the values that minimize t are $\alpha = 2$ and $\beta = 16$. This yields a $t \approx 2^{24}$. This could reasonably be solved in 10 minutes on a modern processor, as were the DIMACS challenges

[15, 42]. Some selected HB-protocol key lengths and their estimate amount of computation using the BKW algorithm are given as follows:

Key Length	BKW Runtime
32	2^{24}
64	2^{35}
96	2^{46}
128	2^{56}
160	2^{64}
192	2^{72}
224	2^{80}
256	2^{88}
288	2^{96}

As mentioned, these runtimes are simply a reflection of the cost of the best known algorithm. With performance improvements or a tighter analysis, it is likely that the effective key-length of LPN keys are even shorter. Regardless, these lengths are still a practical range for low-cost devices and can offer adequate security in many settings.