Implementable Privacy and Security for Resource-Constrained Devices

Karsten Nohl, Benton Calhoun, David Evans (PI), John Lach, and Abhi Shelat
University of Virginia

Submitted to National Science Foundation CyberTrust Program
24 March 2008

Confidential Proposal — Not for redistribution or citation.

Implementable Privacy and Security for Resource-Constrained Devices
Project Summary

Many emerging microelectronics applications involve tight constraints on cost, size, and power consumption. In particular, passive RFID tags target price points on the order of cents and rely on power received from readers, so must have low power consumption to enable long-range reads. Current design approaches and standard cryptographic primitives fail to satisfy the needs of these systems. Nevertheless, these devices are already widely used for applications with serious security and privacy requirements such as key cards, public transportation tokens, and implantable medical devices. Current implementations resort to "security-by-obscurity" and ad hoc solutions that fail to provide adequate security and are frequently broken in practice.

We propose to develop a comprehensive approach for analyzing, designing, and implementing security and privacy on severely resource-constrained devices. Our research focuses on passive RFID systems because they are the most severely constrained devices that are widely used in security applications. Although the size and cost constraints will gradually diminish with expected improvements in chip fabrication technologies, power consumption is primarily a function of the number of active gates which is not reduced by standard technology advances. Hence, we focus our efforts on developing tools and techniques for designing effective low power circuits. The ultimate goal of our project is to enable designers to create secure, cost-effective, large-scale RFID-deployments by combining primitives and protocols from a library and to implement those designs in a principled and efficient manner.

Intellectual Merit. Providing principled security for resource-constrained devices requires crossing traditional abstraction boundaries and reconsidering privacy metrics, cryptographic primitives, and protocol designs in light of hardware constraints. Our comprehensive approach employs a vertically-integrated design method that enables designers to tradeoff desired properties across traditional abstraction layers. We incorporate innovative privacy metrics based on information leakage distributions and new iterative circuit design techniques for cryptographic primitives. Our project exemplifies the interdisciplinary approach needed to solve the problems inherent in RFID systems, incorporating experts in low-energy integrated circuit design (Calhoun), design automation and embedded systems (Lach), network security (Evans), and cryptography (Shelat).

Broader Impacts. Large-scale RFID systems are being widely deployed without adequate understanding of the privacy risks inherent in current designs or principled design approaches for building secure systems. Current and planned RFID deployments pose serious privacy threats to individuals and security risks to society if these problems are not addressed. The metrics, primitives, and design methods we propse to develop will enable designers of future RFID applications to incorporate better security measures and to better evaluate the tradeoffs in different candidate designs. Because the severe cost and power constraints imposed by real RFID applications are driving our approach, the solutions we develop will be practical and cost-effective. Longer term, the proposed work will lead to deeper understanding of the fundamental limits on security as a function of circuit complexity and new general approaches to vertically-integrated design. Our education plan will develop an RFID security lab course, including materials that enable adoption by other schools. We will also produce and teach outreach courses that use cryptography and RFID systems to excite middle school students about pursuing science and engineering.

Implementable Privacy and Security for Resource-Constrained Devices

Karsten Nohl, Benton Calhoun, David Evans (PI), John Lach, and Abhi Shelat
University of Virginia

Project Description

Decreases in the size and cost of microelectronics over the past several decades have led to tiny devices that can compute and communicate. Devices such as RFID tags are small and cheap enough to be proposed as a replacement for printed bar codes as item identifiers, and have already been widely deployed in security-sensitive applications including public transportation tokens, key cards, and pacemakers. These applications raise important security and privacy concerns, but current design methods provide no principled ways for satisfying these requirements. The power requirements for these devices are too constrained to support implementations of conventional cryptographic building blocks and established designs, so designers resort to ad hoc approaches which have repeatedly been shown to be inadequate (e.g., Bono et al.'s cryptanalysis of Texas Instruments' DST device [8], and our cryptanalysis of the NXP Mifare Classic [74, 20], which is currently the most widely used cryptographic RFID device), or do not even implement any cryptography at all despite being used in applications with critical privacy and security requirements (e.g., implantable medical devices [25] and credit cards [26]).

Current designers suffer from three major deficiencies: (1) the lack of adequate metrics for assessing privacy and security properties of candidate designs, (2) the unavailability of cryptographically strong primitives that can be implemented within the power constraints of RFID tags, and (3) the failure of existing design methods to adequately integrate device and system-wide constraints. We propose to develop a principled and comprehensive design methodology and implementation strategy for large scale deployments of resource-constrained devices. Most of the techniques we develop will apply to a range of resource-constrained devices, but we focus on RFID applications since they are the largest and most important class of security-sensitive applications that involve severely limited devices. We target the two most prevalent security applications of RFID technology: *private identification* (enabling an item to identify itself to a legitimate reader, without leaking too much information to an adversary) and *authorization* (using tags to control access, which requires that they cannot be cloned). These primitives cover the requirements for most current uses of RFID tags including logistics tracking, key cards, and public transportation tokens; most other applications can be built straightforwardly using these primitives.

Roadmap. In the next section, we provide background on RFID systems (background on other topics is included in the relevant sections). The following three sections present our research plan:

- Section 2 describes our ideas for reasoning about privacy properties of large scale deployments and estimating the cryptographic strength of new primitives. Our metrics build on information theory, but incorporate models of sophisticated, but rational, adversary behavior.
- Section 3 explains how we will develop cryptographic building blocks for resource-constrained devices. We address both the logical aspects of designing these primitives, and the low-level hardware implementation issues including the use of novel sub-threshold circuit design techniques.
- Section 4 presents our plans toward a vertically-integrated design methodology for RFID deployments. Our goal is to provide designers with a well understood design space for trading

off costs with privacy and security properties. This requires modeling the interdependencies between the protocol layer and the circuit layer and understanding how a change in one affects both implementation metrics (cost, power, performance) and the security and privacy properties.

In Section 5, we summarize our experimentation plan which includes fabricating and testing chips incorporating our design ideas. Section 6 describes our education and outreach plans, including our plans for an exportable RFID security lab course. Section 7 reports results from prior NSF support. Section 8 summarizes impact and milestones for the proposed work.

Team. Our team includes both computer scientists and electrical and computer engineers, spanning the levels of abstraction needed to design RFID systems. Calhoun specializes in low power integrated circuit design and has relevant experience designing low standby leakage circuits, low power SRAM, and sub-threshold digital ICs. Lach's primary research area is integrated circuit design methodologies, with an emphasis on design automation and embedded systems. Shelat works on fundamental cryptography emphasizing theories that start from practical assumptions. Evans' research is on system security, including work on network security and cryptographic protocols.

1 Background on RFID Systems

RFID systems consist of tags (Section 1.1) that communicate with readers that access and update data from a backend database (Section 1.2). In Section 1.3, we describe relevant previous work on protocols for private identification and authentication.

1.1 RFID Tags

RFID tags are tiny, radio-enabled computers equipped with just enough resources to execute simple protocols. In most applications, including logistics and item tagging, the only function of tags is the ability to transmit a unique identifier (ID) that identifies the item to which the tag is attached. These applications demand per tag prices of at most a few cents. Many of these logistics applications require that the tags can be read from several meters away which places severe power constraints on the tags. The resources available on the tags typically include relatively small finite-state machines, and readable (and sometimes writable) memory on the order of at most a few kilobytes. Writing is a special case of operation that requires much higher power; hence, the writing range is much smaller than the reading range which further constrains the design of protocols.

We focus on *passive tags*, which are the longest-lasting and cheapest tags since they do not include a battery, but must draw all of their power from the reader. Active, near-field tags are used in many other security applications (e.g., access control, passports) but are outside the scope of this research. These near-field tags are much more expensive but can only be read from short distances. Security for near-field cards can be achieved through standard asymmetric cryptography.

Applications that require reading passive tags over a distance impose two severe constraints on tags:

1. *Small area.* Since the silicon chip in the tags is the dominant cost factor, the area of the implementation must be very small. The whole tag (including the radio front-end, memory, control logic, and any custom circuits such as cryptographic functions) needs to fit on at most a square millimeter of silicon, which in current technologies means a few thousand gate equivalents (GEs).
2. *Low power.* Passive tags are powered through the reader and since the strength of the reader field is restricted by regulation, the reading distance mostly depends on the power consumption of the tags.

These constraints are not independent since the power consumption depends on the number of gates and therefore on the area. Recent tag generations have already been manufactured in sizes below 0.3mm^2 that require significant and costly changes to the later stages of the manufacturing process. The area constraint will lose its significance over time because the chip size will reach the minimum size that can be physically manufactured. After a few more generations of chip scaling, the smallest possible chip will already provide plenty of space for more than just the basic tag functionality. Power consumption, on the other hand, is primarily a function of the number of active gates, so will not diminish with expected chip fabrication improvements. Our work, consequently, emphasizes the power constraint rather than the area constraint (whereas most previous research focused exclusively on the area constraint, e.g, [97, 35]).

1.2 RFID readers

Readers in RFID systems supply power to the tags, communicate with tags, and forward information to the backend database for further processing. Readers come in a variety of form factors, ranging from cell phones to large doorways with several antennas. Readers typically communicate with a backend database to access and update information associated with the tag. We assume there exists an authenticated channel over which legitimate readers can access the information in the database. The backend has substantial computational power, but the number of cryptographic operations per read is limited. In Section 4, we consider design techniques for trading off backend load with security properties.

In a secure RFID system, some of the readers belong to the system operator and have access to secret key material. An adversary may also deploy rogue readers, but typically these do not have access to all of the keys in the system. The technical core of the readers, the reader IC, is available for few dollars (e.g., NXP RC632 for $7) and whole readers can be mass-produced for at most several dozens of dollars. These numbers are important for estimating the investment an adversary would need to build an infrastructure of rogue readers. The metrics we propose in Section 2.1 incorporate models of rational adversaries.

1.3 Protocols

Several RFID privacy protocols have been proposed, all of which sacrifice at least one of scalability, availability, or strong privacy. The basic hash protocol [98], in which a tag hashes a random nonce with a secret key, provides strong privacy but the backend workload scales linearly with the number of tags in the system since it needs to try the key for each tag in the system to find the one that matches.

A more scalable private identification protocol, such as the tree protocol introduced by Molnar and Wagner [66], assigns several secrets to each tag and uses information from shared secrets to direct the reader's search. In the tree protocol, the secrets are structured in a tree with the tags as the tree leaves. A tag t_i is assigned the secrets $s_{i,1}, s_{i,2}, \ldots, s_{i,d}$ where d is the depth of the tree (all secrets but the last may be shared with some of the other tags). When queried, the tag responds with a nonce and computed hash for each level in the tree: $H(s_{i,1}, N_1), N_1; H(s_{i,2}, N_2), N_2; \ldots; H(s_{i,d}, N_d), N_d$ where $H(\cdot, \cdot)$ is a strong one-way function and the N_i values are random nonces. The database executes the basic hash protocol for each tree level to find the secret used on each level. Once a leaf is reached, the path from the root to the leaf uniquely identifies the tag. This protocol can be extended to also support authorization by incorporating reader-generated nonces to prevent replay attacks.

This protocol scales well beyond billions of tags since (for a fixed branching factor) the number of cryptographic operations required by the reader scales with the depth of the tree. The drawback of the protocol, however, is that secrets are shared among several tags. An attacker who has physical access to some tags in the system can extract secrets from these tags, and then use those secrets to learn something about where in the tree an observed, but uncompromised, tag is found. A sophisticated attacker may be able to combine the information leaked at the protocol layer with contextual and side-channel information leaks to uniquely identify a tag with high probability [10, 71].

Variations of the tree protocol include the matrix protocol that replicates the same secrets over different tree branches and hence offers less privacy, but more flexible tradeoffs [15]. Other proposed protocols provide strong protection only when rogue reads are rare [95]. If rogue reads occur frequently, these protocols can render the tags dysfunctional and leak some information [37].

2 Research Plan: Metrics for Privacy and Security

Principled design is impossible without tools for measuring properties of candidate designs. Computer system designers have reasonable metrics for measuring properties like performance, cost, and power use, and can use these metrics to evaluate designs and make rational trade-offs between competing desirable properties. When it comes to security and privacy, however, the limited metrics and tools that are available are not well suited to the kinds of large scale deployments and minimal cryptography that can be implemented on RFID tags. The next two sections describe our proposed work toward the goal of giving designers effective ways to measure the privacy properties of candidate protocols and to estimate the security of lightweight cryptographic designs.

2.1 Measuring Privacy

The first papers on RFID privacy focused on the requirement that tags should protect product information from being disclosed [98, 77]. This goal can be achieved by ensuring the tag identities are not tied to the items, but it is a weak notion of privacy since it allows an attacker to trace tags. A stronger property, *unlinkability*, means that an adversary should not be able to differentiate between readings that originated from the same tag and readings that originated from different tags. If such readings can be linked, the tags (and the individuals carrying them) can be traced by an adversary. A system achieves *strong privacy*, as defined by Juels and Weis, when an adversary cannot distinguish between two tags with a probability better than random guessing [37]. This is a useful theoretical notion, but is not achievable in scalable protocols. The only known way to achieve strong privacy for large-scale systems with asymmetric cryptography, but known public key ciphers cannot be implemented on cheap RFIDs [38].

Since practical and scalable protocols must leak some privacy, we need more flexible measures of privacy that distinguish between different amounts of leaked information. Modeling this information leakage is difficult because the number of information sources an attacker might include in the attack is virtually unlimited. Our goal is to develop a privacy metric that can incorporate different sources of information leakage, since a sophisticated adversary will take advantage of all available information sources. We focus on rational attackers who want to collect traces of tags, because these traces have value [4]. For example, knowing about their customer's every move would allow businesses to build detailed profiles which can be used for targeting advertising and price discrimination [76]. RFID traces are similar to web traces in that they describe people's actions. Web traces are already harvested and sold, but unlike the web traces, RFID traces describe movement in the real world, making them potentially even more valuable.

Preliminary Work. We have developed an information theoretic metric for measuring privacy lost at the protocol layer, as well as through other sources [71, 72]. Our notion of privacy is closely related to anonymity, which has been studied in the context of mix-nets which aim to provide anonymous messaging [88, 16]. The anonymity set is defined as the set of all potential senders of a given message. The size of the anonymity set is inversely related to the degree of anonymity. Perfect anonymity is achieved if the set includes all members capable of sending messages in the system. These anonymity metrics are based on Shannon's information theory [89]. They use entropy to describe the number of possible elements in a group. Nohara et al. first used entropy in the analysis of the RFID protocols [69]. They only considered the case of a single compromised tag and concluded that almost no information is leaked if the number of tags in the system is large enough. In most realistic scenarios, however, an attacker can obtain secrets from many tags.

Our metric is similar to Buttyán et al.'s [10], except unlike their metric our privacy measure captures the privacy lost independently of the anonymity set size. We measured information leakage from a privacy protecting RFID protocol in terms of entropy and how this metric translates into a measure for attacker success. This metric allows us to capture privacy lost through both the protocol layer and side channels. We have used it to analyze various scenarios such as when typical individuals in the system carry multiple tags [71] and when an attacker can obtain side channel information based on the antenna characteristics [72].

Privacy Distributions. Our metric has the advantage of being able to capture leakage from multiple source, but like other previous privacy metrics, it suffers from two key limitations: (1) it only provides information on the *average* information leakage, whereas a designer needs to understand how information leakage is distributed across tags in the system; and (2) it does not capture any notion of rational behavior of potential adversaries.

Modeling information disclosure as a probability distribution of leaked information exposes the parts of a system responsible for most of the privacy lost. Incorporating likely threats leads to a more realistic metric that can be used to make pragmatic trade-offs for real systems. We have developed a technique for computing these distributions and used it to analyze secret-trees and side channels. The results led us to identify a small subset of tags as the source of most of the privacy loss and provided new insights into the trade-off between cost and privacy [73]. The privacy distribution of secret trees and many other sources can be approximated by an exponential distribution that depends on the number of broken tags, but not on the system size or the spreading factor. Good designs can hence be found that apply to a wide range of applications. When information from several tags or other sources is combined by an attacker, the overall information leakage can be modeled using a single gamma distribution.

Our approach of expressing all privacy leaks in the form of probability distributions enables designers to identify the weakest link and thereby estimate the privacy of the overall system. When combined with a rational attacker model, identifying the weakest part of the tree protocol enabled us to find new parameters for the tree with much improved privacy.

Our metric incorporates an attacker model takes into account the value of different traces. An attack is considered successful only if the overall value exceeds the attack cost, so rational adversaries will only attempt attackes with positive expected values. Without making restrictive assumptions on the actual incentives of the attacker we can prove an upper bound on the value function that corresponds to the most capable attacker. By varying the size of the groups responsible for most of the information leakage, namely the groups at the bottom level of the tree, we can tradeoff increased computational cost for decreased attack value.

Applying our metric to the tree protocol leads us to identify a simple specialization for the protocol in which the tree is limited to two levels. This change lowers the attack value by up to 80%, incurs no cost on the tag, and only a small overhead in the backend. Even though our parameterization seems obvious in retrospect, previous work proposed binary tree structures which appear to provide the least privacy of all possible setups. This underlines the need for good models for information leakage and understanding of the attacker's incentives when designing privacy protocols.

Our proposed work includes developing the privacy distribution metric to incorporate other types of adversaries and realistic models of contextual and side-channel information sources. We will apply this metric to the design candidates to learn about where privacy is leaked and to produce better designs.

2.2 Estimating Cryptographic Strength

The design methodology we propose (Section 3.2) involves automated tests that determine whether a function is distinguishable from a random function by analyzing its output polynomials. Here, we explain how this test works, why it seems to be a useful measure of cryptographic strength for the kinds of primitives we need for RFID applications, and our research plans for improving techniques for estimating the cryptographic strength of simple primitives.

Preliminaries. A function can be defined by polynomials that express each output bit in terms of the input bits. We consider functions that take two input strings, a secret key k and a data value r, and produce an output string: $y \leftarrow f(k, r)$. Each bit of the output, y_i, can be expressed as a binary function of input bits. Each term, known as a *monomial*, in the algebraic normal form (ANF) of these binary functions is the conjunction of one or more input bits. Each ANF of a function with n input bits (key bits plus data bits) has the general form $y_i = a_{1,1}x_1 + a_{1,2}x_2 + \ldots + a_{1,n}x_n + a_{2,1}x_1x_2 + a_{2,2}x_1x_3 + \ldots + a_{2,(n-1)n}x_{n-1}x_n + \ldots + a_{n,1}x_1x_2 \cdots x_n$ where the j^{th} monomial of degree i has a coefficient $a_{i,j} \in \{0, 1\}$ and the terms correspond to the powerset of the input bits. The string formed by these coefficients, $a_{1,1} \parallel a_{1,2} \parallel \cdots \parallel a_{n,1}$, completely describes the function. We argue that if a random function is indistinguishable from an oracle that is executing the ANF on the input value, then the function itself must be indistinguishable from random as well. This claim is supported by the insight that all known cryptographic attacks (including differential and linear cryptanalysis) are based on structural weaknesses that are reflected in non-randomness of the output monomials.

Our goal is to measure how well a given function produces outputs that are indistinguishable from random outputs. A *pseudorandom function family* (PRF) is a family of function for which the output of a randomly chosen member of the family is indistinguishable from random for any efficient (i.e., polynomial time) distinguisher [61]. A *locally random function family* (LRF) is a function family that is indistinguishable from a PRF up to a certain number of distinct inputs, k [61]. A function with input size n and output size m can at best have a key space of $|Z| \leq \min(2^{mk}, 2^n)$ [65]. Any PRF achieves the upper bound of this inequality. While we would want to use real PRFs, they are too expensive. But since the attackers are usually limited in the number of ciphertexts they can obtain, LRFs with large k suffice for many applications.

Once we identify a function that is indistinguishable from an LRF, we apply the Luby-Rackoff construct to build a more globally pseudorandom permutation [61]. The Luby-Rackoff theorem says that after three iterations of a Feistel network with independent PRFs, the mapping between (random) inputs to the Feistel network and its output cannot be distinguished from a random function by a polynomial time attacker. If the inputs to the Feistel network are not random, a fourth iteration is required to make the function indistinguishable from random by an adaptive polynomial time attacker with access to the decryption process. If LRFs are used instead of PRFs, the resulting cipher is only secure for the number of inputs for which the LRF is secure. Normally, $2n$ rounds resist passive distinguishers from random without known plaintext, $3n$ rounds resist non-adaptive statistical attacks by known or chosen plaintext or ciphertext attacks, and $4n$ rounds resist all adaptive attacks (including square, rectangle, boomerang, etc.). Patarin, Naor, and Reingold provided similar proofs for Benes networks and for unbalanced Feistel networks [67, 82]. The construct we propose in Section 3.2 is an example of a source-heavy unbalanced Feistel network. The technique of analyzing the output distribution of ciphers was used in the design of AES [93] and is also part of our tests.

Developing Security Indicators. To test how well a given design implements a function that is indistinguishable from an LRF we have to compute its complete ANF and test it for randomness. Exhaustively computing and analyzing ANFs of arbitrary functions is intractable, however. Instead, our goal is to find any vulnerability that a polynomial time attacker could potentially exploit. This raises two main research challenges:

Computing the ANF. Finding the complete ANF for any large-input circuit is infeasible since the size of the output monomials are exponential in the input size. We instead compute the monomials for each subset of up to a maximum number of input bits and seek a general principle to determine the number of bits needed for an adequate estimate. In our preliminary experiments, all weaknesses that we were able to detect are found with rather small subsets of inputs, while the largest subsets we tested did not yield any new attacks. This is consistent with results from linear and differential cryptanalysis that always produce attacks involving only small sets of inputs and outputs. Nevertheless, it is conceivable that there are attacks based on large subsets of inputs. Answering this open question is part of our planned research.

We also have to restrict the degree up to which we compute the the monomials in some cases. A subset of n inputs can produce monomials up to degree n, but the higher degree terms (which are conjuncts of many input bits) are almost always zero and hence contribute very little to the randomness of the output. This is consistent with successful algebraic attacks on common ciphers where the higher degree terms are eliminated in order to make the output equations solvable. These attacks rely on the scarcity of the monomials, which is one of the properties that our tests detect. By limiting our analysis to only lower degree terms, our test does not seem to loose efficacy; the same weaknesses are discovered. One open question that our research will address is how analysis time should be divided between analyzing larger subsets of inputs and analyzing higher degree monomials of small subsets.

Randomness Tests. None of the strings we generate are truly random (since we know a way of generating them), but they can still be indistinguishable from random. We instead define a sufficient level of local randomness in a practical way: any level of non-randomness that an attacker could possibly exploit can also be tested for with the right randomness tests. It should, therefore, be possible to prevent all statistical attacks through sufficient randomness testing of the algebraic structure of all the relationships between input and output bits (including key and nonce bites) in all known algebraic forms. The algebraic structure of all the ciphers that have been successfully cryptanalyzed in the past, either failed to become a LRF in less than a quarter of the proposed number of rounds, or used self-similar round functions vulnerable to slide attacks instead of using independent ones (which is also detected by our method). Our preliminary analysis shows that some of the standard randomness tests (i.e., NIST [92] and Diehard [63] test suites) can readily be applied to monomial distributions. We also plan to explore new randomness tests specifically designed to yield faster results for monomial distributions. Our research will identify a set of current and new randomness tests that are appropriate for monomial distributions. While we believe that testing only small subsets of input variables for randomness is sufficient, finding the exact bounds around which a function becomes globally indistinguishable from random is an open research question.

3 Research Plan: Cryptographic Building Blocks

New cryptographic primitives are needed that can be implemented within the size and power constraints of RFID devices. We propose to develop new cryptographic building blocks by employing innovative techniques in two directions: using sub-threshold circuit design to increase the number of logical operations that can be carried out (Section 3.1), and developing a new design method for building cryptographically strong ciphers from simple building blocks (Section 3.2). As chip manufacturing processes improve, the size constraint diminishes since the number of gates available in the smallest manufacturable chip increases, but the power constraint remains. This leads to alternative possible long-term solution in which asymmetric lattice-based cryptography can be used to provide stronger privacy properties (Section 3.3.

3.1 Sub-Threshold Circuit Design

We propose to use sub-threshold digital circuit design techniques to implement security and privacy mechanisms in energy-constrained devices. Sub-threshold circuits use a supply voltage, V_{DD}, that is below the threshold voltage, V_{T}, of the transistors. Although the devices are “off” by traditional definitions, the changing gate voltage still causes enough on/off-current difference to provide digital functionality. We have previously demonstrated the viability of both logic [12] and SRAM memory [11] in functional CMOS chips.

The major advantage of this mode of operation is that it provides a dramatic reduction in the power and energy consumption. Power reduces linearly with the lower V_{DD}, and leakage power (e.g. standby power) decreases super-linearly because of the lower V_{DD} and the exponential decrease in leakage current due to drain induced barrier lowering. In previous work, CoPI Calhoun has shown that the minimum energy per operation (E/op) occurs in the sub-threshold region [12]. Circuits operating at the minimum energy point consume one tenth as much energy per operation compared to the normal operating voltage. This means that we can make dramatically more efficient cryptographic building blocks, which translates directly into longer range between the tag and the reader, shorter interactions between tag and reader, or enhanced security through the use of more complex cryptography.

We propose to develop a set of hardware primitives using sub-threshold operation that can be used to implement cryptographic building blocks including one-way functions, stream ciphers, and random number generators. We will design these primitives in the context of a design flow that allows us to redeploy them rapidly and with low non-recoverable engineering cost in new applications. To build a design flow, sub-threshold circuits must be designed carefully to compensate for lower current and larger sensitivity to variations. One key insight for implementing sub-threshold circuits is that, unlike for strong inversion ($V_{DD}>V_T$), the relative strength of NMOS and PMOS devices can vary widely from one process to another. This is equivalent to the typical NMOS, typical PMOS (TT) corner moving relative to the symmetrical case [9, 87]. This means that the basic functional structures that work well for one process may be very inefficient in a different process. We have defined this phenomenon as process balance and modeled it [87]. Now we propose to use this information to generate different sets of sub-threshold standard cells and memory blocks that are optimized for different process balances. Designs for a new process will use the cells that most closely match the process balance of the current process. This design flow will enable us to build the RFID tags for this project; it also support design porting and custom design in new processes.

Taken together, these proposed advances will allow us to use sub-threshold operation to reduce energy consumption by over an order of magnitude (due to its quadratic dependence on V_{DD}) relative to existing RFID implementations. The primary disadvantage of sub-threshold operation is slower circuit speed, which is not a problem for most RFID applications. For cases where it is, we will use an energy scalable architecture that adjusts the supply voltage to match performance demands.

3.2 Circuit-Level Design of Cryptographic Primitives

We propose to develop a new design methodology for symmetric cryptographic primitives, driven from the start by power and size constraints. Using the metrics described in Section 2.2, we will systematically explore the design space for possible circuits to find building blocks that exhibit a high level of randomness and algebraic complexity. Our approach is based on testing iterative constructs for local randomness as proposed by O'Neil [78]. We test candidate constructs for the randomness of their monomial distributions. Any bias in these distributions could be used by an attacker to distinguish a function from random, from which key recovery attacks can be built. The distribution of monomials in these polynomials has previously been found to correlate directly with the resistance of ciphers against statistical attacks including their resistance to linear and differential cryptanalysis [96]. What is novel in our proposed work is using these techniques directly as part of the design process by incorporating them into an exploration of the design space for possible constructs.

Our goal is to find constructs that can be built mostly from circuitry that already exists. This new approach to finding cryptographic primitives is very different from previous approaches in that we start with a highly constrained design space including only choices that build upon non-cryptographic functions already required in the target applications. Previous cryptographic functions have assumed that their design can be completely separate from other functionality and it was often concluded that cryptography would need to be of a certain size to keep enough state for sufficient security. We challenge this by proposing to instead build cryptographic functions that are horizontally-integrated by designing them along with other functionality, and are vertically-integrated by considering the actual hardware implementation when designing the algorithms. Next, we illustrate our approach with two examples: a cryptographic hash function and a random number generator.

Hash Function. Since the CRC error detection circuit is found on virtually all communicating devices including RFIDs (the EPC Class 1 standard for RFIDs requires a CRC function [18]), we propose to design a hash function that reuses the CRC error detection circuit and some temporary memory. Our automated design approach finds several variations that use a CRC function with slight modifications to build a function that passes available randomness tests (Section 2.2). In particular, we are able to build what appears to be a strong S-box by keeping the lowest $n - 1$ bits of the generator polynomial of an n bit CRC circuit variable. This S-box can readily be used in a number of cipher designs, all of which are variations of the Feistel design that iteratively substitutes state words until sufficient complexity in the mapping between inputs and outputs is achieved. In our hash function that constitutes an unbalanced Feistel network, we load a data word into the CRC state in each round, then load the linear combination of the neighboring state words and a secret key word as the CRC generator as shown in Figure 1. This way of using the CRC function leads to a highly non-linear mapping that introduces a high level of diffusion between inputs and outputs.

In our proposed work, we will more thorougly evaluate the security of the proposed design. Preliminary results suggest that no effective linear and differential cryptanalysis exists against our construct and our randomness tests provide a high level of confidence that the design is also secure against all other known statistical attacks. We will also explore other design opportunities that employ existing circuitry with modifications to produce cryptographic building blocks.

Random Number Generator. Cryptographic protocols for authentication and private identification must include random numbers to distinguish freshly generated responses from replays of old responses. Generating good random numbers has often been found to be hard (e.g., [17]) and seems particularly challenging for constrained devices like RFIDs that can neither store values between queries nor employ expensive hash functions to improve the quality of generated random numbers [64]. Recent breakthroughs in random number generation have shown the potential for generating highly random numbers from simple bistable circuits such as SRAM cells [27] and latches [24]. To produce a new random number, the bistable circuit must be driven into a metastable state and then allowed to resolve. A standard SRAM cell or latch used to produce a random number consists of cross-coupled inverters to store the state, and each inverter contains

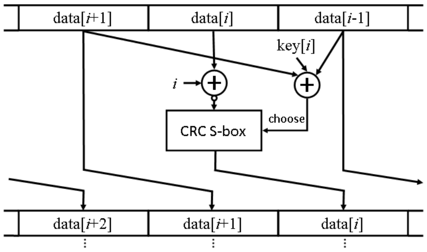

Figure 1: Hash function design incorporating CRC

both PMOS and NMOS devices. The trip point of the cell depends on the relative strengths of the transistors in the cell. Global variations in the devices (e.g., PMOS strength relative to NMOS) and local variations (e.g., one NMOS versus another NMOS) both impact the cell trip point. The extent to which the final value of the bistable circuit is random depends on the relative influence of random processes like thermal noise versus deterministic processes like local device variation and temperature. Only cells that are balanced (e.g., local and global variation) well enough to make thermal noise the primary determinant of the final voltage will produce random outputs.

While previous works have explored the possibility of extracting randomness from existing chips [27, 64], we will design and implement a low power instantiation of RNGs with circuit level improvements. We will improve upon existing RNGs by optimizing them for low power by reducing the supply voltage to sub-threshold levels. For standalone RNG blocks, we will apply leakage reduction approaches such as power gating to reduce the power consumption of the RNG after it has successfully produced the random number. Furthermore, we propose to use storage elements (e.g., registers) that are already integrated in the RFID circuitry (e.g., CRC block) when possible. This dual-use of the registers saves area and power.

We also propose to reduce the vulnerability of the RNG circuits to side channel attacks. Side channel attacks on a RNG might use external factors to attempt to influence the randomness of the bits coming from the RNG. In RFID applications, the attacker has control over external factors such as temperature, clocking, and possibly voltage (depending on how DC regulation is implemented). As noted earlier, cells must be balanced in order to provide random outputs. Changing the temperature of the die will alter the ratio of the PMOS current versus the NMOS current. This is equivalent to changing the global variation, and it can influence bistable circuits to favor a single output value. We will explore three approaches to reducing the influence of external sources on the randomness of on-die RNGs. First, we will investigate new topologies including sense-amplified style latches for the bistable elements. Whereas the stability of standard latches depends on both PMOS and NMOS devices, we have shown that the stability of sense-amplifier style latches depend strongly on only one type of device [86]. Second, we propose to use offset-compensating latches that include special circuits to measure and to offset the impact of imbalancing mechanisms. This approach has been used for many years to allow for more precise sense amplifiers and even for RNG creation [28]. However, no previous work has investigated offset compensation at low supply voltages. Our preliminary results indicate that our offset compensated sub-threshold sense-amplifier latch can reduce the 3-sigma offset from 180mV to 5mV in a 90nm CMOS technology. This will allow us to produce numbers that are more random and to mitigate the influence of any side channel attacks. Finally, we propose to use an on-chip DC-DC converter with a feedback loop to operate the RNG latches at the temperature-independent voltage. Transistor drive strength depends on both carrier mobility and threshold voltage. Mobility increases at lower temperature, while threshold voltage increases. These competing effects lead to each CMOS technology having a specific operating voltage (in the 0.7V to 1.0V range, usually) for which drive current is independent of temperature. We can use a feedback approach to force RNGs only to operate at this voltage.

The three proposed techniques can be used independently or in conjunction. They each improve the randomness and security of RNG circuits, but they each also come at a cost in terms of power and area. We will analyze the overhead of the approaches using simulation and measurements of fabricated ICs (see Section 5) to provide a variety of options to allow a tradeoff between area and power with randomness and security.

3.3 Lattice-Based Cryptography

Traditional public-key encryption schemes rely on algebraic hardness assumptions such as the hardness of factoring, or the hardness of solving Diffie-Hellman-related problems. The algebraic nature of these problems allows for rigorous analysis and security proofs, but results in cryptographic operations that require millions of instructions to implement. For example, a recent analysis of a popular cryptographic protocol reports that a single RSA operation requires approximately 18M instructions [104].

Lattice-based cryptography offers the potential to bring public-key cryptography to smaller devices. Ajtai first proposed using lattice problems to build cryptosystems in 1996 [2]. There are a number of good reasons to consider such schemes, most notably that the operations involved only require modular addition of small numbers. The problem with these schemes, however, is the large size of the public-key and ciphertext. Previous solutions have had blowups on the order of $O(k^4)$ where k is a security parameter. Recent work by Regev [85] and Gentry et al. [21] have made some progress by reducing the size increase to either $O(k^2)$ or even $O(k)$ based on proper setup.

Despite these advances, the large keys needed for lattice-base cryptography are prohibitive for current RFID tag designs. Likely technology advances, however, will enable more room for key storage without supporting more complex function. As Gordon Moore observed, the evolution of microprocessors allows for the same functionality to be implemented in less area with every generation of fabrication technology. Although the implementation footprint shrink by about 50% every 18 months, the power consumption of the circuit stays approximately the same. While this trend leads to increased demand for cooling and power-saving in high-performance applications, it hinders the progress of embedded applications where power is the main constraint. An RFID tag for example, cannot have more functionality, but can only be implemented ever smaller. This trend also has a limit since silicon chips have to be of a certain minimum size to be able to still handle them in the manufacturing process and this lower end has already been reached for the cheapest, long-range RFID tags. In future process generations, the functionality of these tags will stay the same to not exceed the power budget, but its implementation will not fill the available area anymore. This unused space can, however, be filled with additional storage that consumes hardly any power. In particular NAND ROM memory that can only be written to once, consumes very little power and can be switched off when not in use. This storage can then be used to hold large secret keys used for lattice-based ciphers.

There are still several engineering factors which make implementation difficult. In the case of one of the schemes [21], the algorithm for picking a public and private key pair requires a subroutine that is extremely complicated and exists more in the realm of theory than practical code [3]. Thus, a first research question is to determine the practical feasibility of proposed cryptosystems and find a practical instantiation of lattice-based cryptography.

Another interesting research opportunity with such schemes is the ability to approximate some of the modular additions instead of computing them exactly. Indeed, the decryption scheme relies on checking whether a ciphertext is “far” or “close” to a lattice point. Thus, even if some of the operations are slightly incorrect in the lower order bits, the ciphertext can be correctly decrypted. In the case of RFID tags, this ability to be less careful with the modular additions might be exploited to save power and space. Thus, one research question is to characterize the precision necessary in lattice-based cipher computations to assure correct encoding and decoding. The security proofs for these schemes will need to be reconsidered to ensure that approximate addition does not inadvertently leak information.

4 Research Plan: Vertically-Integrated Design

Our goal is to provide a principled design method for developing secure, large-scale, cost-effective RFID systems. We employ a design approach that breaks traditional abstraction barriers to enable better and more efficient designs. An ongoing NSF-funded project led by CoPI Lach [47] is developing this new design methodology as well as a tool for designing dataflow intensive algorithms that enables both algorithm metrics and implementation metrics to be co-optimized during early design exploration. A common data structure, the hierarchical dependency graph (HDG), provides a representation of both the algorithm and the implementation architecture, formally establishing the bidirectional relationship between an algorithm and its implementation.

Using the hierarchical nature of the HDG to manage complexity, algorithm designers and hardware designers can explore the collaborative space we call the *ColSpace* together, trading off various metrics while searching for the best overall design. Our design methodology opens various trade-offs for new security protocols and hence spans a large design space. Often, the same two factors can be traded on different levels. Performance, for instance, can be traded for security on the protocol layer by varying the key length; on the primitive layer by using simpler mathematical operations; and on the implementation layer through serialization and by varying the supply voltage. Hardware designers can identify implementation bottlenecks (i.e. protocol requirements that establish an unacceptable floor for power, area, or latency) using the HDG structures and communicate the desired changes to the protocol designers via HDG augmentations for evaluation of any impacts on security and privacy metrics, including the metrics we propose in Section 2. These metrics can be incorporated into a designer-defined optimization cost function that automatically evaluates solutions in the ColSpace (including solutions that include protocols alterations), searching for the minimum cost solution.

Exploring this design space requires establishing a bidirectional relationship between the protocol layer and the implementation layer, a relationship that is currently seen as unidirectional — alterations to protocols are easily evaluated in terms of impacts on implementation metrics, but the reverse is not true. ColSpace presents all these design tradeoffs in a single design space and enables designers to find the best protocol-primitive-implementation choice across all available degrees of freedom. Next, we illustrate the benefits of ColSpace design with examples of cryptographic primitives and protocols that demonstrate the benefits of crossing traditional abstraction barriers in designing RFID systems.

Cryptographic Primitives. The co-design of cryptographic algorithms and their hardware implementation can sometimes even lead to higher security and smaller implementation at the same time when compared to more traditionally designed cryptographic primitives. One simple example of this opportunity involves the round counters that are a mandatory part of all iterated ciphers such as most block ciphers, stream ciphers, and hash functions. These round counters are used to vary the cryptographic structure of different rounds to prevent slide attacks [5, 6].

Round counters in almost all ciphers are a simple increment that generates the sequence $1, 2, 3, \ldots$, while the cryptographic need for diversity of rounds only requires a sequence of unique values that may as well be unordered. In software, the obvious way to generate a round counter is to just increment an integer. In a hardware implementation, generating an unordered sequence is much cheaper than implementing an increment, and in fact the increment is often implemented by mapping the values from an unordered sequence of unique values. This mapping step does not provide any additional cryptographic strength and can be eliminated from the the implementation leading to an equally secure, but smaller cipher. An optimal sequence of round values has previously been shown to have no linear dependence between values [79], which in the design of these functions has always been assumed to come at an additional cost. In hardware, however, generating an unordered sequence without linear dependencies simply requires moving from an LFSR to a NLFSR which comes at the additional cost of a single NAND gate. The security of any iterated primitive can hence be improved by using an NLFSR-based round counter in hardware more cheaply than adding an increment counter. This fact has not been previously exploited, even though a large number of iterated ciphers have been designed since Feistel first proposed this structure in the in the 1970s, but it was easily found using our integrated design approach.

Another example where the cooperation between hardware designers and cryptographers creates new opportunities are S-boxes. These boxes are used in Feistel-type ciphers to scramble data in every round. Traditionally, ciphers have been proposed with different small S-boxes so that the cipher would still be strong when one of the boxes proves weak. In software, there is no performance penalty for applying different function as compared to applying the same function multiple times. In serialized hardware, however, using several S-boxes makes a significant cost difference. A designer, hence, faces a tradeoff between extra assurance against undetected flaws and implementation cost. In the PRESENT block cipher [7], for example, the DES design is altered to only include one instead of eight S-boxes and the resulting cipher appears to offer a security level similar to that of DES.

Protocols. Extending the vertically-integrated design approach to the protocol layer provides further opportunities. Protocols need to balance the cost and power consumption of the tag with the computational complexity of the reader and back-end system in a way that achieves security and privacy requirements.

One example of a design space that provides strictly better choices by considering parameters on different layers is the improved private identification protocol we proposed in which tag responses incorporate random noise to make it harder for an attacker to distinguish different tags [70]. This increases the computational cost for any legitimate reader as well and the ratio between this cost and the added amount of privacy provides a tradeoff that is vastly superior to previous tradeoffs for the same protocol. The use of noisy responses has previously been proposed as a mean to improve the hardness of certain cheap hash functions [36]. Randomization, hence, provides improvements for both hash functions and protocols in certain applications, but only the synergy of benefiting from noise on both levels at once makes this an attractive design technique for a much larger domain of applications. A design space that does not consider choices on all levels would probably not have found this result.

Another example where a vertically-integrated approach may lead to better designs is the hash functions needed in private identification protocols. Existing protocols are designed assuming a standard cryptographic hash function which is too power-hungry to implement in RFID tags. However, the actual requirements for the hash functions as they are used in the protocol are weaker than those for standard cryptographic hash functions. In particular, they do not require collision resistance, which is typically the hardest property to achieve in hash function design. Hence, it seems likely that less inexpensive building blocks could be designed that serve the needs of private identification protocols. We plan to formally investigate the necessary properties of private hash functions and use the ColSpace approach to design cryptographic building blocks that satisfy those properties.

Our methodology opens corners of the design space that previous less-integrated approaches could not find. Since many of these corners provide parameterizations with superior properties, our approach not only provides a larger number of choices but in many cases better choices.

5 Experimentation Plan

In addition to our theoretical and simulation analyses, we plan to evaluate our work with experiments using fabricated ICs. To verify and quantify the benefits of the proposed techniques, we will fabricate two ICs. The ICs will be designed using industry-standard Cadence design tools. These tools and the necessary computer support for them are available in our labs, and we have previously designed several ICs that were successfully fabricated. We will build test pads into the ICs to facilitate testing the parts using equipment that is already available.

The first IC (year 1) will contain a number of cryptographic primitives including the random number generator and a one-way hash function. This test chip will allow each block to be tested individually. We will measure the power consumption of each block versus supply voltage and experimentally identify the optimum voltage for minimizing energy consumption. Using the metrics developed in this work, we will quantify the security that could be provided by systems incorporating these blocks and examine the tradeoff between power and security. The second IC (year 3) will implement a complete RFID tag including power management, communication circuits, and a full security block that instantiates the findings from the proposed work. The security block will be designed to be flexible in the power/security space so that it can increase security at a cost of power consumption (e.g., requiring shorter range or longer read times) when desired.

6 Education and Outreach Plan

CoPI Lach led an interdepartmental committee to establish the Computer Engineering graduate program which was established in 2002 and is jointly administered by the Computer Science and Electrical and Computer Engineering Departments. The proposed work exemplifies the kinds of problems that require combined expertise in electrical engineering and computer science that the Computer Engineering PhD program seeks to develop in our students. In addition to the graduate student mentoring, we plan to develop and co-teach a new RFID security lab course, involve undergraduates in our research, and conduct outreach activities.

Course Development and Distribution. RFID systems combine circuit design, networking, and cryptography in a way that presents exciting teaching opportunities. We plan to develop a new undergraduate laboratory course in which students will develop and experiment with RFID systems. To enable this, we will develop a combined educational/research laboratory that will include platforms for experimentation, providing hands-on learning experiences for both undergraduate and graduate students. The experimental, hands-on nature of this laboratory is necessary to assess and address the practical limitations of real implementations in a variety of application and deployment scenarios. In the course, we will use an RFID scanner system and RF communication equipment (transmitters, receivers, spectrum analyzers, etc.) to simulate attack and defense scenarios. We will also incorporate rapid prototyping technologies, including FPGA boards for implementing experimental circuit designs and interfaces to RFID platforms for incorporating designs into system experiments. A design tool suite including both commercial tools and custom tools (such as ColSpace), will also be included that allows for designs to be input and augmented at several layers in the abstraction hierarchy, from protocols to chip architectures to low-level circuits, and then rapidly incorporated into prototypes for experimentation.

The course will be developed around a series of lab exercises that will be packaged in a way that supports easy adoption by other universities following the model of the Virginia Internet Teaching Laboratory [51, 52]. Some of these labs will incorporate offensive and defensive measures where groups of students act as red and blue teams in a variety of challenges and contests. The course will be developed and co-taught by Lach and Evans, both winners of the All-University Teaching Award (in 2005 and 2008, respectively).

Undergraduate Research. We have a successful record of actively recruiting and involving undergraduate students in our research groups. Over 50 undergraduates have participated actively in our research groups on NSF-funded projects, including three recent students who have been recognized nationally as CRA outstanding undergraduates: Jon McCune (co-advised by Evans and Lach, CRA outstanding undergraduate honorable mention in 2003, finishing at PhD at CMU), Salvatore Guarnieri (advised by Evans, CRA finalist in 2006, currently a PhD student at U. Washington), and Adrienne Felt (advised by Evans, CRA finalist in 2008, currently on a whirlwind grad school tour). PI Evans is founding director of the BA Computer Science program for College of Arts and Sciences students, which provides students in the College of Arts and Sciences an opportunity to major in Computer Science. Because the BA curriculum is designed to attract students without previous computing experience to computing, includes a gateway course that motivates computing using examples from the arts and sciences with societal impact (such as RFID privacy), and the student body of the College of Arts and Sciences is much more diverse than that of the Engineering school, the BA degree is attracting more underrepresented minorities and women to computing. As advisor to many of these students and director of the Distinguished Majors Program (which includes a substantial research experience), the PI will have excellent opportunities to recruit talented undergraduates from diverse backgrounds to participate in this research project.

K-12 Outreach. RFID systems and cryptography provide wonderful opportunities for getting K-12 students excited about science and engineering. We plan to target our outreach activities primarily toward 8th-grade students and teachers, because studies point to this as the age where students, especially women and underrepresented minorities, most often lose interest in math and science [62]. Over the past 3 years, the PI has developed a two-day course that introduces cryptography to 7th and 8th graders and taught the course to over 200 students at a nearby low-income middle school as well as to summer students in the GEAR-UP program [94]. Typical middle school math classes cover topics with particularly strong connections to cryptography including algebra, probability, combinatorics, and factoring, but often fail to give students much reason to be excited about these topics. Our experience with these courses indicates that cryptography can be used to make mathematics exciting and appealing to a wide range of students. The ubiquity, and physical nature, of RFID tags and readers, presents an exciting opportunity for teaching outreach courses that explore cryptographic protocols, as well as the societal issues associated with RFID systems.

Public Outreach. Many of the privacy and security issues raised by RFID systems involve complex tradeoffs that society must make to balance the interests of freedom, privacy, and commerce. A well-informed public is critical for making political decisions about these tradeoffs, and we believe communicating some of the technical results from the proposed work has the potential to aid public understanding and influence political decision making. As an example, our initial results on the Mifare cryptanalysis were released as the Netherlands were about to deploy a nationwide ticketing system, OV-Chipkaart, based on Mifare technology. Because of concerns about the security and privacy issues raised by our analysis (and later confirmed by other groups), deployment of the system was delayed, there were discussions in the Dutch parliament about the importance of open design, NXP announced a new version of the Mifare tag with improved security [75], and Trans Link Systems (responsible for the OV-Chipkaart deployment) has entered a partnership with us to improve deployment security.

7 Results from Prior NSF Support

David Evans is PI of two completed NSF research grants: *CAREER: Programming the Swarm* ($285K, 3/1/01-2/28/06), 0092945) and *ITR: A Framework for Environment-Aware Massively Distributed Computing* (with Abdelzaher and Brogan, $400K, 9/15/02-8/31/05, 0205327). These projects explored programming and security issues involved in building dependable software systems. Our work to date has produced a biological and environmental programming models [1, 22, 23]; protocols for secure wireless networks [29, 31, 30]; an analysis of security principles for virtual machines [83, 84]; dynamic inference techniques [100, 99, 101] and the Splint tool [50, 19] for detecting security vulnerabilities using lightweight static analysis that is included in most Linux distributions. Our work has been cited over 1000 times in the research literature, and the tools we produced including CellSim, MCL, and Splint are used by dozens of other research groups. Evans is also a CoPI on the ongoing CyberTrust grant, *CT-T: A System Structure for Secretless Security* (with Knight, Davidson, Nguyen-Tuong, and Rowanhill, $1.65M, 10/01/05 - 09/30/08, 0524432). This project developed the N-variant framework for protecting vulnerable services from compromise by executing processes with artificial variation and ensuring they behave consistently and a kernel-level implementation of that framework. The work has been published in USENIX Security [13] and DSN [68], and Evans has presented it in invited seminars in China and at MIT, Harvard, Purdue, UTSA, and SDWest.

John Lach has been the PI on four NSF grants and a co-PI on four more. One relevant grant, *Small-Scale Dynamic Reconfigurability for Large-Scale Benefits* (John Lach (PI), Kevin Skadron, and Mircea Stan, $419,784, 9/1/01-8/31/04, 0105626) focuses on developing the notion of small-scale reconfigurability (SSR) for efficient dynamic adaptation in microprocessors and other integrated circuits. Many applications have characteristics that share broad similarities but differ in their detailed hardware requirements. SSR enables the hardware (e.g., datapath, branch prediction units, caches, and other processor components) to either be configured on a per-application basis, or dynamically adapt based on runtime characteristics to find the best hardware configuration for a specific application. These ideas were applied in GPPs, ASICs, and application specific instruction processors (ASIPs), The project produced ideas and techniques disseminated in tool development and publications [14, 33, 32, 34, 49, 48, 53, 54, 56, 57, 55, 58, 59, 60, 80, 81, 90, 91, 39, 40, 42, 41, 46, 45, 43, 44, 102, 103] and helped support initial development of the HotSpot thermal modeling software (downloaded 1500 times since its release in June 2003 and used in numerous academic publications as well as some industrial settings).

Benton Calhoun joined UVa in 2006, and **Abhi Shelat** joined UVa in 2007. They have no prior NSF support.

8 Impact Summary

The proposed work will lead to improved metrics for evaluating the privacy of RFID deployments and the security of cryptographic primitives, develop a new approach to designing cryptographic building blocks that incorporate sub-threshold circuit design and a new iterative design method, and vertically-integrated design methods for opening new areas of the design space. Although our focus is on RFID systems because of their immediate societal importance, many of the research topics we will investigate have other applications and we hope the metrics, cryptographic structures, and design methods we develop will lead to advances for other applications. The major milestones for the project are summarized below.

Year 1

- Refine privacy distribution metric and apply it to example protocols.
- Develop algorithms for computing approximations of the ANF for candidate circuits and analyze tradeoffs between complexity and efficacy.
- Design low-power circuits for hash function and random number generation.
- Augment ColSpace design tool to incorporate quantitative privacy and security metrics and proposed low-power circuit design techniques.
- Create ColSpace HDG design models for security and privacy circuit building blocks.
- Tape out the first IC containing security and privacy circuit building blocks.
- Set up RFID security lab and offer co-taught RFID security lab course.
- Develop and teach outreach course incorporating RFID privacy and security protocols.

Year 2

- Analyze attack scenarios and incorporate different attack assumptions into privacy metrics.
- Create HDG design models for initial RFID tag design with integrated security.
- Redesign the components from the first IC using the design models.
- Demonstrate in simulation a DC regulator for RFID power management, passive communication system, and a scalable security block.
- Test the IC produced in year 1 for correctness, power consumption, and security properties.
- Complete design of a RFID tag incorporating several different cryptographic structures.
- Offer RFID security lab course incorporating red-team exercises.

Year 3

- Create improved versions of the cryptographic building blocks based on results from experiments and improved design tools.
- Use ColSpace and initial HDG models to explore design options for RFID tag.
- Tape out the second IC containing a complete RFID tag with integrated security.
- Test the second IC for functional correctness, power consumption, and security properties.
- Release exportable lab manual based on materials from RFID security course.

References Cited

[1] Tarek Abdelzaher, B. Blum, Q. Cao, Y. Chen, D. Evans, J. George, S. George, L. Gu, T. He, S. Krishnamurthy, L. Luo, S. Son, J. Stankovic, R. Stoleru, and A. Wood. EnviroTrack: Towards an Environmental Computing Paradigm for Distributed Sensor Networks. In *24th International Conference on Distributed Computing Systems*, Mar 2004.

[2] Miklós Ajtai. Generating Hard Instances of Lattice Problems. In *ACM Symposium on Theory of Computing (STOC)*, 1996.

[3] Miklós Ajtai. Generating Hard Instances of the Short Basis Problem. In *26th International Colloquium on Automata, Languages and Programming*, 1999.

[4] Matthias Bauer, Benjamin Fabian, Matthias Fischmann, and Seda Gürses. Emerging Markets for RFID Traces. *arXiv (http://arxiv.org/abs/cs.CY/0606018)*, August 2006.

[5] Alex Biryukov and David Wagner. Slide Attacks. In *Sixth International Workshop on Fast SoftwareEncryption (FSE)*, Mar 1999.

[6] Alex Biryukov and David Wagner. Advanced Slide Attacks. In *Advances in Cryptology — Proceedings of EUROCRYPT 2000*, May 2000.

[7] A. Bogdanov, L.R. Knudsen, G. Leander, C. Paar, A. Poschmann, M.J.B. Robshaw, Y. Seurin, and C. Vikkelso. PRESENT: An Ultra-Lightweight Block Cipher. In *Cryptographic Hardware and Embedded Systems (CHES)*, Aug 2007.

[8] Stephen C. Bono, Matthew Green, Adam Stubblefield, Ari Juels, Aviel D. Rubin, and Michael Szydlo. Security Analysis of a Cryptographically-Enabled RFID Device. In *14th USENIX Security Symposium*, Aug 2005.

[9] Andres Bryant, Jeffrey Brown, Peter Cottrell, Mark Ketchen, John Ellis-Monaghan, and E. J. Nowak. Low-Power CMOS at $V_dd = 4kT/q$. In *Device Research Conference*, Jun 2001.

[10] Levente Buttyán, Tamás Holczer, and István Vajda. Optimal Key-Trees for Tree-Based Private Authentication. In *International Workshop on Privacy Enhancing Technologies (PET)*, 2006.

[11] Benton H. Calhoun and Anantha P. Chandrakasan. A 256kb Sub-threshold SRAM in 65nm CMOS. In *IEEE International Solid-State Circuits Conference*, Feb 2006.

[12] Benton H. Calhoun, Alice Wang, and Anantha P. Chandrakasan. Modeling and Sizing for Minimum Energy Operation in Sub-threshold Circuits. *IEEE Journal of Solid-State Circuits*, 40(9):1778–1786, Sep 2005.

[13] Benjamin Cox, David Evans, Adrian Filipi, Jonathan Rowanhill, Wei Hu, Jack Davidson, John Knight, Anh Nguyen-Tuong, and Jason Hiser. N-Variant Systems: A Secretless Framework for Security through Diversity. In *15th USENIX Security Symposium*, Aug 2006.

[14] Puyan Dadvar and Kevin Skadron. Potential Thermal Security Risks. In *Semiconductor Thermal Measurement, Modeling, and Management Symposium*, 2005.

[15] Ivan Damgård and Michael Østergaard Pedersen. RFID Security: Tradeoffs between Security and Efficiency. In *RSA Conference, Cryptographers' Track*, 2008.

[16] Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards Measuring Anonymity. In *Privacy Enhancing Technologies Workshop (PET)*, 2002.

[17] Leo Dorrendorf, Zvi Gutterman, and Benny Pinkas. Cryptanalysis of the Windows Random Number Generator. In *14th ACM Conference on Computer and Communications Security*, 2007.

[18] EPCglobal. EPC Radio-Frequency Identity Protocols: Class-1 Generation-2 UHF RFID. *http://www.epcglobalinc.org/standards/uhfc1g2/*, Dec 2005.

[19] David Evans and David Larochelle. Improving Security Using Extensible Lightweight Static Analysis. *IEEE Software*, Jan 2002.

[20] Sharon Gaudin. RFID hack could crack open 2 billion smart cards: One European government sent armed guards to protect facilities using the card. *ComputerWorld*, Mar 2008.

[21] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for Hard Lattices and New Cryptographic Constructions. In *ACM Symposium on Theory of Computing (STOC)*, May 2008.

[22] Selvin George, David Evans, and Lance Davidson. A Biologically Inspired Programming Model for Self-Healing Systems. In *Workshop on Self-Healing Systems*, Nov 2002.

[23] Selvin George, David Evans, and Steven Marchette. A Biological Programming Model for Self-Healing. In *First ACM Workshop on Survivable and Self-Regenerative Systems*, Oct 2003.

[24] Jorge Guajardo. The Life Cycle of an FPGA Intrinsic PUF: From Start-Up To Butterfly. In *Secure Component and System Identification Workshop (SECSI)*, Mar 2008.

[25] Daniel Halperin, Thomas S. Heydt-Benjamin, Benjamin Ransford, Shane S. Clark, Benessa Defend, Will Morgan, Kevin Fu, Tadayoshi Kohno, and William H. Maisel. Pacemakers and Implantable Cardiac Defibrillators: Software Radio Attacks and Zero-Power Defenses. In *IEEE Symposium on Security and Privacy*, May 2008.

[26] Thomas S. Heydt-Benjamin, Daniel V. Bailey, Kevin Fu, Ari Juels, and Tom O'Hare. Vulnerabilities in First-Generation RFID-enabled Credit Cards. In *Eleventh International Conference on Financial Cryptography and Data Security*, Feb 2007.

[27] Daniel E. Holcomb, Wayne P. Burleson, and Kevin Fu. Initial SRAM State as a Fingerprint and Source of True Random Numbers for RFID tags. In *Conference on RFID Security*, Jul 2007.

[28] J. Holleman, B. Otis, S. Bridges, A. Mitros, and C. Diorio. A 2.92 μW Hardware Random Number Generator. In *32nd European Solid-State Circuits Conference (ESSCIRC)*, 2006.

[29] Lingxuan Hu and David Evans. Secure Aggregation for Wireless Networks. In *Workshop on Security and Assurance in Ad Hoc Networks*, Jan 2003.

[30] Lingxuan Hu and David Evans. Localization for Mobile Sensor Networks. In *10th Annual International Conference on Mobile Computing and Networking (MobiCom)*, Sep 2004.

[31] Lingxuan Hu and David Evans. Using Directional Antennas to Prevent Wormhole Attacks. In *Network and Distributed System Security Symposium (NDSS)*, Feb 2004.

[32] Wei Huang, Mircea Stan, and Kevin Skadron. Physically-Based Compact Thermal Modeling-Achieving Parameterization and Boundary Condition Independence. In *THERMINIC*, 2004.

[33] Wei Huang, Mircea Stan, Kevin Skadron, Karthik Sankaranarayanan, Shougata Ghosh, and Sivakumar Velusamy. Compact Thermal Modeling for Temperature Aware Design. In *Design Automation Conference*, 2004.

[34] Philo Juang, Kevin Skadron, Margaret Martonosi, Zhigang Hu, Douglas Clark, Philip Diodato, and Stefanos Kaxiras. Implementing Branch Predictor Decay Using Quasi-Static Memory Cells. *ACM Transactions on Architecture and Code Optimization*, 1(2):180–219, June 2004.

[35] Ari Juels. RFID Security and Privacy: A Research Survey. *Journal of Selected Areas in Communication*, 24(2):381–395, Feb 2006.

[36] Ari Juels and Stephen Weis. Authenticating Pervasive Devices with Human Protocols. In *25th Annual International Cryptography Conference (CRYPTO)*, 2005.

[37] Ari Juels and Stephen Weis. Defining Strong Privacy for RFID. In *Fifth Annual IEEE International Conference on Pervasive Computing and Communications Workshops*, Mar 2007.

[38] Sandeep Kumar and Christof Paar. Are Standards Compliant Elliptic Curve Cryptosystems Feasible on RFID? In *Workshop on RFID Security*, 2006.

[39] V. Vijay Kumar and John Lach. Designing, Scheduling, and Allocating Flexible Arithmetic Components. In *International Conference on Field Programmable Logic and Applications*, 2003.

[40] V. Vijay Kumar and John Lach. Fine-Grained Self-Healing Hardware for Large-Scale Autonomic Systems. In *International Workshop on Autonomic Computing Systems*, 2003.

[41] V. Vijay Kumar and John Lach. Flexible Arithmetic Components for Area-Efficient Fault Tolerance. In *International Conference on Military and Aerospace Programmable Logic Devices*, 2003.

[42] V. Vijay Kumar and John Lach. Heterogeneous Redundancy for Fault and Defect Tolerance with Complexity Independent Area Overhead. In *International Symposium on Defect and Fault Tolerance in VLSI Systems*, 2003.

[43] V. Vijay Kumar and John Lach. Highly Flexible Multi-Mode Digital Signal Processing Systems Using Adaptable Components and Controllers. *EURASIP Journal on Applied Signal Processing*, 2005.

[44] V. Vijay Kumar and John Lach. Highly Flexible Multi-Mode System Synthesis. In *International Conference on Hardware/Software Co-Design and System Synthesis*, 2005.

[45] V. Vijay Kumar and John Lach. IC Modeling for Yield-Aware Design with Variable Defect Rates. In *Annual Reliability and Maintainability Symposium*, 2005.

[46] V. Vijay Kumar, Rashi Verma, John Lach, and Joanne Dugan. A Markov Reward Model for Reliable Synchronous Dataflow System Design. In *International Conference on Dependable Systems and Networks*, 2004.

[47] John Lach, Scott Acton, and Kevin Skadron. Hierarchical Dependency Graphs for ColSpace Design with Application to Leukocyte Detection and Tracking. NSF grant IIS-0612049, July 2006 – July 2009, *http://www.ece.virginia.edu/ jcl7d/ColSpace/ColSpace.htm*.

[48] John Lach, Jason Brandon, and Kevin Skadron. A General Post-Processing Approach to Leakage Current Reduction in SRAM-based FPGAs. In *International Conference on Computer Design*, 2004.

[49] John Lach, David Evans, Jon McCune, and Jason Brandon. Power-Efficient Adaptable Wireless Sensor Networks. In *International Conference on Military and Aerospace Programmable Logic Devices*, 2003.

[50] David Larochelle and David Evans. Statically Detecting Likely Buffer Overflow Vulnerabilities. In *10th USENIX Security Symposium*, 2001.

[51] Jörg Liebeherr. VINTLab: The Virginia Internet Teaching Laboratory. *http://www.cs.virginia.edu/vintlab/*.

[52] Jörg Liebeherr and Magda El Zarki. *Mastering Networks: An Internet Lab Manual*. Addison-Wesley, 2004.

[53] Zhijian Lu, J. Hein, M. Stan, John Lach, and Kevin Skadron. Control-Theoretic Dynamic Frequency and Voltage Scaling. In *Workshop on Self-Healing, Adaptive and Self-Managed Systems*, 2002.

[54] Zhijian Lu, Jason Hein, Marty Humphrey, Mircea Stan, John Lach, and Kevin Skadron. Control-Theoretic Dynamic Frequency and Voltage Scaling for Multimedia Workloads. In *International Conference on Compilers, Architecture, and Synthesis for Embedded Systems*, 2002.

[55] Zhijian Lu, Wei Huang, John Lach, Mircea Stan, and Kevin Skadron. Interconnect Lifetime Prediction Under Dynamic Stress for Reliability-Aware Design. In *International Conference on Computer Aided Design*, 2004.

[56] Zhijian Lu, John Lach, Mircea Stan, and Kevin Skadron. Alloyed Branch History: Combining Global and Local Branch History for Robust Performance. *International Journal of Parallel Programming*, 31:137–177, 2003.

[57] Zhijian Lu, John Lach, Mircea Stan, and Kevin Skadron. Reducing Multimedia Decode Power using Feedback Control. In *International Conference on Computer Design*, 2003.

[58] Zhijian Lu, John Lach, Mircea Stan, and Kevin Skadron. Banking Chip Lifetime: Opportunities and Implementation. In *Workshop on High Performance Computing Reliability Issues*, 2005.

[59] Zhijian Lu, John Lach, Mircea Stan, and Kevin Skadron. Temperature-Aware Modeling and Banking of IC Lifetime Reliability. *IEEE Micro*, pages 40–49, November/December 2005.

[60] Zhijian Lu, Yan Zhang, Mircea Stan, John Lach, and Kevin Skadron. Procrastinating Voltage Scheduling with Discrete Frequency Sets. In *Design Automation and Test in Europe*, 2006.

[61] Michael Luby and Charles Rackoff. How to Construct Pseudorandom Permutations from Pseudorandom Functions. *SIAM Journal on Computing*, 17(2):373–386, 1988.

[62] Jane Margolis and Allan Fisher. *Unlocking the Clubhouse: Women in Computing*. MIT Press, Apr 2003.

[63] George Marsaglia. The Marsaglia Random Number CDROM including the Diehard Battery of Tests of Randomness. *http://www.stat.fsu.edu/pub/diehard/*, 1995.

[64] M. Matsumoto, S. Yasuda, R. Ohba, K. Ikegami, T. Tanamoto, and S. Fujita. $1200\mu^2$ Physical Random-Number Generators Based on a SiN MOSFET for Secure Smart-Card Applications. In *International Solid-State Circuits Conference*, Feb 2008.

[65] Ueli Maurer. A Simplified and Generalized Treatment of Luby-Rackoff Pseudorandom Permutation Generators. In *Advances in Cryptology, EUROCRYPT*, 1992.

[66] David Molnar and David Wagner. Privacy and Security in Library RFID: Issues, Practices, and Architectures. In *ACM Conference on Computer and Communications Security*, 2004.

[67] Moni Naor and Omer Reingold. On the Construction of Pseudo-Random Permutations: Luby-Rackoff Revisited. *Journal of Cryptology*, 12(1), 1999.

[68] Anh Nguyen-Tuong, David Evans, John C. Knight, Benjamin Cox, and Jack W. Davidson. Security through Redundant Data Diversity. In *38th IEEE/IFPF International Conference on Dependable Systems and Networks*, Jun 2008.

[69] Yasunobu Nohara, Sozo Inoue, Kensuke Baba, and Hiroto Yasuura. Quantitative Evaluation of Unlinkable ID Matching Schemes. In *Workshop on Privacy in the Electronic Society (WPES)*, 2006.

[70] Karsten Nohl. Privacy through Noise: A Design Space for Private Identification. In *Secure Component and System Identification Workshop (SECSI)*, Mar 2008.

[71] Karsten Nohl and David Evans. Quantifying information leakage in tree-based hash protocols. In *Eigth International Conference on Information and Communications Security (ICICS)*, Dec 2006.

[72] Karsten Nohl and David Evans. Quantifying information leakage in tree-based hash protocols. Technical Report UVA-2006-20, University of Virginia, Dec 2006.

[73] Karsten Nohl and David Evans. Hiding in Groups: On the Expressiveness of Privacy Distributions. In *23rd International Information Security Conference (IFIP SEC)*, Sep 2008.

[74] Karsten Nohl, David Evans, Starbug, and Henryk Plötz. Reverse-Engineering a Cryptographic RFID Tag . In Submission, Jan 2008.

[75] Mary Catherine O'Connor. NXP Announces New, More Secure Chip for Transport, Access Cards. *RFID Journal*, Mar 2008.

[76] Andrew Odlyzko. Privacy, Economics, and Price Discrimination on the Internet. In *Fifth International Conference on Electronic Commerce*, 2003.

[77] Miyako Ohkubo, Koutarou Suzuki, and Shingo Kinoshita. Cryptographic Approach to "Privacy-Friendly" Tags. In *RFID Privacy Workshop*, 2003.

[78] Sean O'Neil. Algebraic Structure Defectoscopy. Cryptology ePrint Archive, Sep 2007.

[79] Sean ONeil, Benjamin Gittins, and Howard A. Landman. VEST Ciphers. *http://www.ecrypt.eu.org/stream/p2ciphers/vest/vest_p2.pdf*, Aug 2006.

[80] Dharmesh Parikh, Kevin Skadron, Yan Zhang, Marco Barcella, and Mircea Stan. Power Issues Related to Branch Prediction. In *International Symposium on High-Performance Computer Architecture*, 2002.

[81] Dharmesh Parikh, Kevin Skadron, Yan Zhang, and Mircea Stan. Power-Aware Branch Prediction: Characterization and Design. *IEEE Transactions on Computers*, 53(2):168–186, February 2004.

[82] Jacques Patarin and Audrey Montreuil. Benes and Butterfly Schemes Revisited. In *International Conference on Information Security and Cryptology*, 2005.

[83] Nathanael Paul and David Evans. .NET Security: Lessons Learned and Missed from Java . In *Twentieth Annual Computer Security Applications Conference*, Dec 2004.

[84] Nathanael Paul and David Evans. Comparing Java and .NET security: Lessons Learned and Missed. *Computers & Security*, 25(5), Jul 2006.

[85] Oded Regev. On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In *ACM Symposium on Theory of Computing (STOC)*, 2005.

[86] Joseph Ryan and Benton H Calhoun. Minimizing offset for latching voltage-mode sense amplifiers for sub-threshold operation. In *International Symposium on Quality Electronic Design*, March 2008.

[87] Joseph F. Ryan, Jiajing Wang, and Benton H. Calhoun. Analyzing and Modeling Process Balance for Sub-threshold Circuit Design. In *ACM Great Lakes Symposium on VLSI*, pages 275–280, March 2007.

[88] Andrei Serjantov and George Danezis. Towards an Information Theoretic Metric for Anonymity. In *Privacy Enhancing Technologies Workshop (PET)*, 2002.

[89] Claude E. Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27, 1948.

[90] Kevin Skadron, Tarek Abdelzaher, and Mircea Stan. Control-Theoretic Techniques and Thermal-RC Modeling for Accurate and Localized Dynamic Thermal Management. In *International Symposium on High-Performance Computer Architecture*, 2002.

[91] Kevin Skadron, Mircea Stan, Wei Huang, Karthik Sankaranarayanan, Zhijian Lu, and John Lach. A Computer-Architecture Approach to Thermal Management in Computer Systems: Opportunities and Challenges. In *International Conference on Thermal, Mechanical and Thermo-Mechanical Simulation and Experiments in Micro-electronics and Micro-systems*, 2004.

[92] Juan Soto. Statistical Testing of Random Number Generators. In *22nd National Information Systems Security Conference*, Oct 1999.

[93] Juan Soto and Lawrence Bassham. Randomness Testing of the Advanced Encryption Standard Finalist Candidates. Technical Report NIST IR 6483, National Institute of Standards and Technology, Mar 2000.

[94] UVa Today. Math Academy Prepares Middle-Schoolers for College, Jul 2007.

[95] Gene Tsudik. YA-TRAP: Yet Another Trivial RFID Authentication Protocol. In *International Conference on Pervasive Computing and Communications*, 2006.

[96] Serge Vaudenay. An Experiment on DES Statistical Cryptanalysis. In *ACM Conference on Computer and Communications Security (CCS)*, 1996.

[97] Stephen Weis. Radio-Frequency Identification Security and Privacy. Master's thesis, Massachusetts Institute of Technology, Jun 2003.

[98] Stephen Weis, Sanjay Sarma, Ronald Rivest, and Daniel Engels. Security and Privacy Aspects of Low-Cost Radio Frequency Identification Systems. In *International Conference on Security in Pervasive Computing*, 2003.

[99] Jinlin Yang and David Evans. Automatically Inferring Temporal Properties for Program Evolution. In *15th IEEE International Symposium on Software Reliability Engineering (ISSRE)*, Nov 2004.

[100] Jinlin Yang and David Evans. Dynamically Inferring Temporal Properties. In *ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering (PASTE)*, Jun 2004.

[101] Jinlin Yang, David Evans, Deepali Bhardwaj, Thirumalesh Bhat, and Manuvir Das. Perracotta: Mining Temporal API Rules from Imperfect Traces. In *28th International Conference on Software Engineering (ICSE)*, May 2006.

[102] Yan Zhang, John Lach, Kevin Skadron, and Mircea Stan. Odd/Even Bus Invert with Two-Phase Transfer for Buses with Coupling. In *International Symposium on Low Power Electronics and Design*, 2002.

[103] Yan Zhang, Zhijian Lu, John Lach, Mircea Stan, and Kevin Skadron. Optimal Procrastinating Voltage Scheduling for Hard Real-Time Systems. In *Design Automation Conference*, 2005.

[104] Li Zhao, Ravi Iyer, Srihari Makineni, and Laxmi Bhuyan. Anatomy and Performance of SSL Processing. In *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Mar 2005.