

A Technique Security Analysis Using Cryptography of RFID Device

Stephen C. Bono, Matthew Green, Adam Stubblefield
Department of Computer Science, Johns Hopkins University, USA

Abstract

We describe our success in defeating the security of an RFID device known as a Digital Signature Transponder (DST). Manufactured by Texas Instruments, DST (and variant) devices help secure millions of SpeedPass™ payment transponders and automobile ignition keys.

1. Introduction

Radio-Frequency Identification (RFID) is a general term for small, wireless devices that emit unique identifiers upon interrogation by RFID readers. Ambitious deployment plans by Wal-mart and other large organizations over the next couple of years have prompted intense commercial and scientific interest in RFID [23]. The form of RFID device likely to see the broadest use, particularly in commercial supply chains, is known as an EPC (Electronic Product Code) tag. This is the RFID device specified in the Class 1 Generation 2 standard recently ratified by a major industry consortium known as EPCglobal [9, 19]. EPC tags are designed to be very inexpensive – and may soon be available for as little as five cents/unit in large quantities according to some projections [21, 20]. They are sometimes viewed in effect as wireless barcodes: They aim to provide identification, but not digital authentication. Indeed, a basic EPC tag lacks sufficient circuitry to implement even symmetrickey cryptographic primitives [21].

The term RFID, however, denotes not just EPC tags, but a spectrum of wireless devices of varying capabilities. More sophisticated and expensive RFID devices can offer cryptographic functionality and therefore support authentication protocols. One of the most popular of such devices is known as a Digital Signature Transponder (DST). Manufactured by Texas Instruments, DSTs are deployed in several applications that are notable for wide-scale deployment and the high costs (financial and otherwise) of a large-scale security breach. These include:

In this paper, we describe our success in attacking the Texas Instruments DST system. We are able to recover the secret cryptographic key from a target DST device after harvesting just two challenge-response pairs. For arbitrary challenge-response pairs, we are able to recover a key in under an hour using an array of sixteen FPGAs. When the challenge-response pairs derive from pre-determined challenges, i.e., in a chosen-plaintext attack, a time-space trade-off is possible, reducing the cracking time to a matter of minutes. The full details of this chosen-response attack will appear in a future version of this work. Once we have recovered a key, we are able to use an inexpensive, commodity RF device to “clone” the target DST, that is, to simulate its radio output so as to convince a reader.

1.1 Related work

The pre-eminent historical example of black-box reverse-engineering of a cipher was the reconstruction of the Japanese Foreign Office cipher Purple during the Second World War. Under the leadership of William F. Friedman, the United States Signals Intelligence Service performed the feat of duplicating the Purple enciphering machine without ever having physical access to one [13].

There are a number of well known contemporary examples of the reverse-engineering of proprietary cryptographic algorithms. For example, the RC4 cipher, formerly protected as a trade secret by RSA Data Security Inc., was publicly leaked in 1994 as the result of what was believed to be reverse-engineering of software implementations [4]. The A5/1 and A5/2 ciphers, employed for confidentiality in GSM phones, were likewise publicly disclosed as a result of reverse engineering. The exact method of reverse-engineering has not been disclosed, although the source was purportedly “an actual GSM phone” [6].

There are also numerous published fault-induction and side-channel attacks against hardware devices, e.g.,

[5]. These are related to our work, but involve rather different techniques, and generally aim to recover a key, rather than a cipher design.

In fact, we are unaware of any published *black-box* reverse-engineering of a contemporary cipher, whether the original source be a software or a hardware implementation. Having no literature to refer to, we developed techniques for our effort in an ad hoc manner. While our techniques are not easily generalizable, we hope that they will nonetheless aid future researchers at a conceptual level in similar endeavors.

In contrast, the key-recovery techniques we employed are well known. Our parallel FPGA key-recovery system operates on much the same principle, for example, as the well publicized Deep Crack machine that employed hardware-based key-space searching to recover DES keys [10, 17].

As mentioned above, our system for key recovery in software using chosen-challenge pairs exploits the timespace tradeoff developed by Hellman. We also employ the “distinguished point” enhancement of Rivest. We have drawn on previous work by Quisquater *et al.* [18], who created a Hellman-based system for recovering keys from a variant of DES employing 40-bit keys.

We note that we have chosen in this document to reveal information about the DST cipher sufficient to elucidate our reverse-engineering and analysis techniques. We omit details about that would permit direct reconstruction of the cipher. Our goal is to offer our fullest possible contribution to the scientific community, but at the same time to avoid fomenting abuse of our results. Once the industry has had time to take adequate measures, we intend to divulge full cipher details in the interest of stimulating cryptanalytic research by the scientific community.

In general, the problem of achieving authentication in RFID devices – with and without full-blown cryptography – has seen a recent burgeoning in the data-security literature. An up-to-date bibliography may be found at [3].

2 . Practical Significance of Our Results

Our attack on the DST cipher by no means implies wholesale dismantling of the security of the SpeedPass network, nor easy theft of automobiles. The cryptographic challenge-response protocols of DST devices constitute only one of several layers of security in these systems. ExxonMobil informed us that the SpeedPass network has on-line fraud detection mechanisms loosely analogous to those employed for

traditional credit-card transaction processing. Thus an attacker that simulates a target DST cannot do so with complete impunity; suspicious usage patterns may result in flagging and disabling of a SpeedPass device in the network. The most serious system-wide threat lies in the ability of an attacker to target and simulate multiple DSTs, as suggested in our example scenarios below.

In some sense, the threat to automobile immobilizers is more serious, as: (1) An automobile is effectively an off-line security system and (2) A single successful attack on an automobile immobilizer can result in full compromise of the vehicle. While compromise of a DST does not immediately permit theft of an automobile, it renders an automobile with an immobilizer as vulnerable to theft as an automobile without one. Such a rollback in automobile security has serious implications. As noted above, significant declines in automobile theft rates – up to 90% – have been attributed to immobilizers during their initial introduction. Even now, automobile theft is an enormous criminal industry, with 1,260,471 automobile thefts registered by the FBI in 2003 in the United States alone, for a total estimated loss of \$8.6 billion [16].

Extracting the key from a DST device requires the harvesting of two challenge-response pairs. As a result, there are certain physical obstacles to successful attack. Nonetheless, bypassing the cryptographic protections in DST devices results in considerably elevated real-world threats. In this section we elaborate on the context and implications of our work.

2.1 Effective attack range

There are effectively two different methods by which an attacker may harvest signals from a target DST, and two different corresponding physical ranges.

The first mode of attack is *active scanning*: The attacker brings her own reader within scanning range of the target DST. DSTs of the type found in SpeedPass and automobile ignition keys are designed for short range scanning – on the order of a few centimeters. In practice, however, a longer range is achievable. In preliminary experiments, we have achieved an effective range of several inches for a DST on a keyring in the pocket of a simulated victim. A DST may respond to as many as eight queries per second. Thus, it is possible to perform the two scans requisite for our simulation attacks in as little as one-quarter of a second. At the limit of the range achievable by a given antenna, however, scanning becomes somewhat unreliable, and

can require more time.

From the standpoint of an attacker, active scanning has the advantage of permitting a chosen-challenge attack. Hence this type of attack permits the use of precomputed Hellman tables as touched on above. In principle, therefore, it would be possible for an attacker with appropriate engineering expertise to construct a completely self-contained cloning device about the size of an Apple iPod. When passed in close proximity to a target DST, this device would harvest two chosen-challenge transcripts, perform a lookup in an on-board set of precomputed Hellman tables in the course of a minute or so, and then simulate the target DST. We estimate that the cost of constructing such a device would be on the order of several hundred dollars.

The second mode of attack is *passive eavesdropping*. Limitations on the effective range of active scanning stem from the requirement that a reader antenna furnish power to the target DST. An attacker might instead eavesdrop on the communication between a legitimate reader and a target DST during a valid authentication session. In this case, the attacker need not furnish power to the DST. The effective eavesdropping range then depends solely on the ability to intercept the signal emitted by the DST.

We have not performed any experiments to determine the range at which this attack might be mounted. It is worth noting purported U.S. Department of Homeland Security reports, however, of successful eavesdropping of this kind on 13.56 Mhz tags at a distance of some tens of feet [24]. The DST, as we explain below, operates at 134 kHz. Signals at this considerably lower frequency penetrate obstacles more effectively, which may facilitate eavesdropping. On the other hand, larger antennas are required for effective signal interception.

Only careful experimentation will permit accurate assessment of the degree of these two threats. Our cursory experiments, however, suggest that the threats are well within the realm of practical execution.

2.2 Example attack scenarios

For further clarification of the implications of our work, we offer a few examples of possible DST exploits. We let Eve represent the malefactor in these scenarios.

Example 1 (Auto theft via eavesdropping) Eve runs an automobile theft ring. She owns a van with eavesdropping equipment. She parks this near a target automobile so as to eavesdrop on ignition-key-to-reader transmissions. After observing two turns of the ignition key, she is able to extract the cryptographic

key of the DST at her leisure using an FPGA. She returns subsequently to steal the target automobile. To enter the vehicle, she picks or jimmies the door lock. She then hot-wires the ignition and deactivates the immobilizer by simulating the DST of the real key.

Example 2 (Auto theft via active attack) Eve runs an automobile theft ring. She suborns a valet at an expensive restaurant to scan the immobilizer keys of patrons while parking their cars, and to note their registration numbers. Based on these registration numbers, Eve looks up the addresses of her victims (such background checks being widely offered on the Internet). By simulating their DST devices, Eve is able to steal their automobiles from their homes.

Example 3 (SpeedPassTM theft via active attack) Eve carries a reader and short-range antenna with her onto a subway car. (Alternatively, Eve could carry a large “package” with a concealed antenna in some public place.) As she brushes up near other passengers, she harvests chosen challenge-response pairs (along with the serial numbers of target devices) from any SpeedPass tokens they may be carrying. Later, at her leisure, Eve recovers the associated cryptographic keys. She programs the keys into a software radio, which she uses to purchase gasoline. To allay suspicion, she takes care to simulate a compromised SpeedPassTM only once. Additionally, she hides the tag simulator itself under her clothing, interacting with the pump reader via an antenna passing through a sleeve up to an inactive SpeedPassTM casing.

2.3 Fixes

The most straightforward architectural fix to the problems we describe here is simple: The underlying cryptography should be based on a standard, publicly scrutinized algorithm with an adequate key length, e.g., the Advanced Encryption Standard (AES) in its 128-bit form, or more appropriately for this application, HMACSHA1 [15]. From a commercial standpoint, this approach may be problematic in two respects. First, the required circuitry would result in a substantially increased manufacturing cost, and might have other impacts on the overall system architecture due to increased power consumption. Second, there is the problem of backwards compatibility. It would be expensive to replace all existing DST-based immobilizer keys. Indeed, given the long production cycles for automobiles, it might be difficult to introduce a new cipher into the immobilizers of a particular make of vehicle for a matter of years. On the other hand, it may be presumed from the Kaiser presentation [14] that Texas Instruments has plans to

update their cipher designs in the DST. Additionally, TI has indicated to the authors that they have more secure RFID products available at present; in lieu of specifying these products, they have indicated the site www.ti-rfid.com for information.

3. Reverse Engineering

We now present the details of our attack.

The 40-bit cipher underpinning the DST system is proprietary. We refer to this cipher, in accordance with the Texas Instruments documentation, as DST40. The only substantive technical information we were able to locate on DST40 was a rough schematic available in a presentation by Dr. Ulrich Kaiser, which was published on the Internet [14], and in a published conference paper [11] coauthored by Dr. Kaiser with Texas Instruments employees. We show the schematic here in Figure 1. Although the general form and arrangement of functional components of the cipher are clear, critical details about the logic and interconnections of the components are lacking. Moreover, as we shall explain, certain features in the published schematic are inaccurate, in the sense that they do not correspond to the workings of the fielded versions of DST40 we have examined. To distinguish over the DST40 cipher employed in the implementations that we experimented with, we refer to the cipher adumbrated in Figure 1 as the Kaiser cipher.

We considered several initial approaches to reverseengineering DST40. Software packages available from TI may include the DST40 cipher. These packages, however, include license agreements that prohibit reverseengineering, so we did not make use of them. We might alternatively have attempted to probe a hardware implementation of DST-40. Lacking the resources for this approach and aiming, moreover, to explore the minimum resource requirements to attack the DST system successfully, we rejected this option. Instead, we chose to mount an “oracle” or “black-box” attack. This is to say that we chose to uncover the functional details of DST40 simply by examining the logical outputs of a DST device. As the keys of some DST devices are field-programmable, we were able to experiment with a set of chosen keys and inputs for DST40. For our experiments, we purchased a TI Series 2000 - LF RFID Evaluation Kit; this kit contained a reader and antenna, as well as a variety of non-DST RFID devices. We were able to separately purchase small quantities of DST devices in two different form factors.

To explain our effort further, some nomenclature is in order. As the TI schematic of the Kaiser cipher

suggests, DST40 is essentially a feedback shift register. In each round of operation, inputs from the challenge register and key register pass through a collection of logical units. These yield an output value that is fed back into the challenge register. We refer to the full collection of logical units operating in a single round, i.e., operating on a single set of inputs with no feedback, as F . The function F comprises three logical layers.

The first layer, represented in the schematic by boxes $f1...f16$, consists of a collection of sixteen functional units, each of which takes a small number of bits from the key register and a small number of bits from the challenge register and yields a one-bit output. We refer to these functional units as *f-boxes*. Each *f-box* takes either three key bits and two challenge bits or vice versa; two special *f-boxes* take only two bits from each register as input.

The second layer, represented in the schematic by boxes $f17...f20$, consists of four functional units, each of which takes as input the outputs of a set of four *f-boxes*. We refer to each of these second-layer units as a *g-box*. Finally, the third layer consists of a single functional unit, labeled $f21$, into which feed the outputs of the *gboxes*. This last functional unit, which we refer to as the *h-box*, yields the output of the full function F .

There are two main technical lacunae in the TI schematic. Reverse-engineering DST40 required that we focus on these. In particular:

1. The schematic does not describe the logical operations executed by the *f-boxes*, the *g-boxes*, and the *h-box*.
2. The mapping of key and challenge bits to the *fboxes* is governed by a routing array whose organization the TI schematic does not describe. In other words, there is no indication of which bits in the challenge and key registers are input to which *f-boxes*.

In addition, as we shall explain, the TI schematic contains some inaccuracies, in the sense that the Kaiser cipher does not correspond exactly to DST40 as implemented in DST devices. Critically, the function F (and thus the *h-box*) in DST40 actually outputs a *pair* of bits, and the clocking of the cipher is accordingly different. Moreover, two bits in the challenge register are XORed with the output of the *h-box*.

3.1 Obtaining a single-round output

Because we did not know the contents of the *f-boxes*, or other critical details such as the configuration of the routing networks in DST40, we could not directly

verify that production DSTs, such as those that we obtained for our experiments, implemented the Kaiser cipher of Figure 1. Instead we had to test and exploit structural features of the Kaiser cipher, using our evaluation DST as an oracle.

We first noted that the only round dependence of the Kaiser cipher is in the key scheduler. As seen in Figure 1, a $_0$ key, i.e., string of '0' bits, will remain unchanged in the key register throughout the cipher execution. Using this key, it is possible to render each step of the algorithm independent of the round in which it takes place. We used the $_0$ key for the experiments we now describe.

We next observed that each cycle, i.e., each execution of F , results in only a small change to the state of the challenge register: The contents of the register are shifted right by one bit, and the output of the h-box is inserted into the leftmost bit position. Consequently, for any given value submitted to the tag, the challenge register can assume only two possible values after one clock cycle, depending on whether the h-box outputs a '0' or a '1' bit.

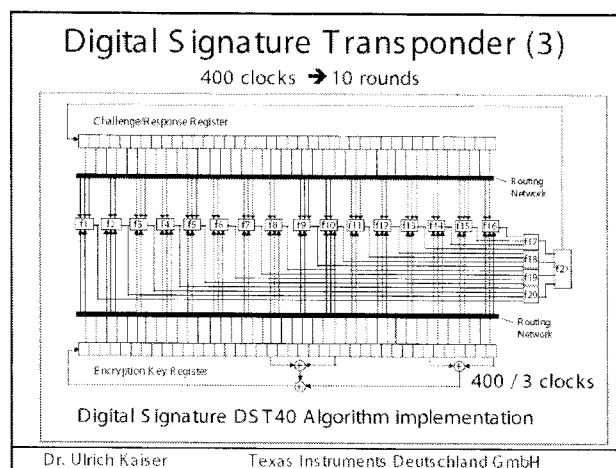


Figure 1. Schematic of Kaiser Cipher.

Using the DST as an oracle, we developed a test to recover the output of the h-box for any value in the challenge/response register (which for brevity we henceforth refer to as the challenge register). Consider a given challenge/response pair $_C,R_$. Upon input of C to the DST device, the challenge register initially contains the bits of C . Let $C_$ denote the sequence of bits in C excluding the last bit, i.e., the first 39 bits in C . Based on our observation above, after a single cycle, the challenge register in the DST contains one of two possible sequences, either $C_0 = 0|C_$ or $C_1 = 1|C_$,

where $|$ denotes concatenation. Therefore, recovering the output of the h-box can be reduced to a determination of whether the challenge register assumes the value C_0 or C_1 after the first cycle.

Fortunately, access to the DST as a challenge/response oracle offers a simple way to make this determination. If C_0 is truly the "next-state" value in the challenge register, then application of the full encryption process, i.e., the full 400 clock cycles on C_0 , will yield a result identical to the encryption of C , but shifted one cycle ahead. The original response R will therefore appear in the challenge register exactly one clock cycle prior to the encryption finishing. Thus, the final result, which we call R_0 , will contain R , but shifted right by one bit. Alternatively, if C_0 is *not* the next-state value of the challenge register, the oracle will likely produce a response unrelated to the original response R , as the single-bit difference will tend to be amplified by the cipher over hundreds of rounds. The analogous observation holds, of course, for C_1 .

Based on the assumption that the encryption circuitry in a production DST tag implements the algorithm of Figure 1, we performed this test using an evaluation DST purchased from Texas Instruments. First, we programmed the DST with the $_0$ key, then submitted a challenge C , along with the two possible "next-state" values C_0 and C_1 .

Unfortunately, this test failed to produce the results we expected, indicating that DST40, the algorithm in the production DST, differs from the Kaiser cipher. After submitting a number of properly-formed challenges to the DST, we saw none of the expected correlations – i.e., neither C_0 nor C_1 produced responses correlated to the original response R .

Through trial and error, we discovered that the method of testing next-state challenge-register values succeeded *when we modeled the output of the h-box as two bits*. For a given challenge C , this required that we instead compute four candidate next-state values, C_{00} , C_{01} , C_{10} , C_{11} , i.e., one for each of the possible output bit-pairs of the h-box. In our experiments, at least one of these four candidates always produced a response corresponding to the initial response R , but shifted right by two bits.² One possible explanation was that the circuit alters its operation every other clock cycle, causing our test to malfunction. It was far more likely, however, that the production cipher DST40 simply produces two bits per cycle. Such a divergence from the diagram of Figure 1 called into question other elements of the diagram, including the number of rounds in the encryption process, and the key update schedule.

With the ability to recover the output of F on a single

cipher iteration, we were able to use our oracle to observe each round of a cipher execution from start to finish, by repeatedly guessing the next state of the challenge register. This approach established that the encryption process took place over 200 cycles, i.e., 200 executions of F , and that the DST draws its response from the rightmost 24 bits of the challenge register on the conclusion of this process.

3.2 Recovering the key schedule

Our experiments, as described above, relied on the assumption that the 0 key remains constant through every cycle of the encryption process. Using only the 0 key, however, would restrict our ability to experiment with the algorithm internals. We required the ability to observe single-round outputs based on different values in the challenge *and* key registers.

Using a non-zero key again made the algorithm rounddependent, with the result that our previous tests would no longer produce useful results. In order for our “nextstate” candidate challenges to be encrypted properly, we needed to provide the oracle with the equivalent nextstate of the key register.

Our next step, therefore, was to reverse-engineer the key schedule. Following the diagram, we assumed that new key bits were computed every few cycles by the exclusive-or of several bits of the key. By querying our oracle, we determined that the key is updated every *three* cycles, beginning with the second cycle – not the first, as suggested by the Kaiser diagram. We also determined that while four bits are indeed exclusiveored together, they are not the bits shown in the diagram. Let k_i denote the i th bit in the key register, beginning with 0. The actual key update is defined by $k_0 = k_{39} \oplus k_{37} \oplus k_{20} \oplus k_{18}$. Note that this design represents an LFSR with a primitive-characteristic polynomial, so all keys other than the 0 key produce maximal length sequences of key register values. Using this model for the key schedule in place of the 0 key, we were able to simulate steps of the algorithm for any key.

With the 0 key, we only had to guess each of the possibilities for the 2-bit output of a single round. To experiment with a non-zero key k , we needed to guess six successive bits (three bit-pairs) of output for the h -box simultaneously, because the key schedule only repeats every three cycles. (6 is the l.c.d. of 2 and 3.) This meant testing 64 possible candidate challenge-register states, $\{Cb_1b_2b_3b_4b_5b_6\}b_1, b_2, \dots, b_6 \in \{0, 1\}$. To test one of these challenge-register states, we programmed into the DST device a key k corresponding to the key-register state after six cipher

cycles applied to k . A DST can process 6-8 challenges per second, so this test requires a minimum of 8 seconds or so. It is thus significantly more time-consuming than previous tests, although it returns the output of three execution cycles, rather than one.

3.3 Uncovering the Feistel structure of DST40

Figure 2 shows the probability that modifying an individual challenge bit results in a change to the output of F . To measure this effect, we generated a random key and challenge, then determined the output of F . Next, for each of the 40 challenge bits, we determined whether the output of F changed upon flipping of the bit. We repeated this test for 150 initial key and challenge settings. (We performed a similar test involving the flipping of key bits, but the results were not significant.)

Let c_i denote the i th bit of the challenge register, starting with 0. The first notable feature of our graph is the effect of bits c_{38} and c_{39} of the challenge register. While the other key and challenge bits have limited influence on the output of a single round, these two bit *always* affect the output of the h -box. Further experimentation revealed that the two bits affect the first and second bit of the two-bit round output respectively. This indicated that the cycle output derived from the exclusive-or of these bits with the output of the F function.

The XOR effect of bits c_{38} and c_{39} shed new light on the algorithm’s design. Not only is the algorithm an invertible permutation, but it is a form of Unbalanced Feistel Network [22].

For the DST, the choice of a Feistel cipher is not a necessary choice, although a useful one. We speculate that the round function was chosen to be a permutation, so that the effect of collisions would not multiply, and the responses would have a uniform distribution.

3.4 Recovering the bit routing networks

After identifying the general structure of the cipher, our next step was to uncover the internal routing network of bits, i.e., which bits act as inputs to each of the f -boxes, as well as the boolean functions computed by each f -box. We made the working assumption (eventually validated) that the h -box (f_{21}) is the only box with a 2-bit output, and the rest each produce a single output bit.

The structure of the Kaiser cipher is such that h receives a single input bit from each of the g -boxes, and produces one or four possible output values. This fact lays the groundwork for identifying which bits of

the challenge and key are routed to each of the g-boxes. It is clear that altering a single input bit of h can at most produce two distinct output values. In consequence, altering the output of only one g-box can never cause h to output more than two distinct values, whereas altering the output of more than one g-box can produce up to four distinct output values.

Using this simple but powerful observation, we devised a test to determine which groups of input bits from the challenge and key are routed into each of the four gboxes. The test involves fixing a set of all but two challenge or key bits, and then iterating through all four combinations of these two bits. If at any time these four bit combinations produce more than two different outputs, then they cannot possibly be routed through the same g-box. It should be noted that this test of g-box membership produces false positives. In particular, it is very possible (and indeed common) that for two test bits that are not routed to the same g-box, and for a given set of fixed bits, different value assignments to the test bits still produce two or fewer distinct outputs from the h-box. Therefore this test requires many repetitions with different sets of fixed bits.

We employed this test first so as to exclude all bits that are *not* in the same g-box as bit 0 of the challenge. After excluding 60 such bits, we discovered all of the bits that *are* routed to g1. We repeated the test for the remaining g-boxes, ignoring bits previously associated with a g box so as to decrease the search space. To our benefit, the routing network of bits that go through each g-box is arranged in a rather regular pattern, and it was not necessary to perform an exhaustive search. After uncovering most of the bits related to g1, we were able to infer and then quickly verify the remainder of the g-boxes.

A slightly more complicated task lay in determining the routing network for DST40, namely which bits of the challenge and key registers serve as inputs to each f-box.




Figure 2. Frequency of change in single round output value on flipping of individual challenge bits.

We already know that altering the output of any given fbox will only affect a single g-box, and therefore

cause the output of h to assume only one of two distinct values. It is helpful, however, to note some other simple facts about the cipher. Let $B = \{b_1 \dots b_5\}$ be a set of challenge and key-register bits. Let B denote all other bits in the challenge and key registers. Our central observation is the output of the cipher will show a special invariant if B is the set of input bits to a single f-box.

A given f-box implements a fixed boolean function z on five bit inputs. (Two of the f-boxes in DST40 have only four inputs, but the principle is the same.) Let us suppose that B is the set of inputs to this f-box. We can then define A_0 to be the set of value assignments to the bits in B such that $z(b_1 \dots b_5) = 0$, and define A_1 analogously for $z(b_1 \dots b_5) = 1$. Observe that for a fixed setting of B , the output of h will be invariant for the setting of B to any value in A_0 . Likewise, for a fixed value assignment to B , the output of h will be invariant for any setting of B to a value in A_1 .

In contrast, suppose that B consists of register bits are are input to two or more f-boxes. In this case, we are unlikely to see the invariant we have just described. For some, and perhaps many settings of B , any given set of value assignments A_0 (or A_1) may induce multiple output values in h .

Using our invariant, we can perform a test to exclude combinations of bits that *cannot* be inputs to the same fbox. We first select a set B of five bits. We fix all their bits B in the challenge and key registers. We iterate over all 32 value assignments to B and record the pattern of outputs from F . It may be the case that only a single output of h results, in which case we repeat the experiment. In the case that there are two distinct outputs from F , we record their correspondence to input values, i.e., we construct a hypothesis for A_0 and A_1 .³ We repeat this experiment over a new setting of B . If we do not see the invariant described above for A_0 and A_1 , then B cannot consist of inputs to a single f-box. We successfully repeated this test until we excluded all possible f-box input combinations except the correct ones.⁴

On first inspection, it would appear as though there is a large number of possible sets of input bits to any given f-box. In fact, though, we can narrow the pool of candidate sets thanks to two observations: (1) The set of inputs to a single f-box must also serve as inputs (at one remove) to the same g-box; and (2) For any f-box, three input bits come from the challenge register and two from the key register (or vice versa). By working with inputs corresponding to a single g-box and by searching in particular for the f-box that includes bit 0 of the challenge register, we started with a search space of size only $(19/2 \times 20/2) + (19/1) \times (20/3) = 54150$. Moreover,

once we identified the inputs to one f-box, each subsequent fbox corresponding to the same g-box had far fewer combinations of input bits to test.

Furthermore, again to our benefit, the f-box inputs in DST40 are ordered in a very regular manner. In particular, given the structure of inputs associated with one gbox, we were readily able to infer those for the remaining g-boxes.

4. Key Cracking

We were able to verify that our reverse engineering of the DST40 algorithm was successful by testing whether the responses computed by a software implementation of our hypothesized algorithm matched those returned by an evaluation DST when given the same challenge and key.

We also wished to test our implementation against actual fielded tags in SpeedPass™ tokens and automobile ignition keys. The cryptographic keys in these devices are immutable once locked at the factory. Without knowing the key on a fielded tag, we had no way to determine whether the algorithm used by such tags was as hypothesized. Therefore, recovering an actual key became necessary.

5. RF Protocol Analysis and Simulation

Low-frequency RFID systems typically make use of powered readers and passive transponders. In the DST system, the reader transmits power to the transponder via a 15-to-50 ms electromagnetic pulse at 134.2 kHz. Once powered, a transponder can receive and respond to commands from the reader, including challenges and read and write operations. The transponder can also execute computations, including as cryptographic operations.

A reader transmits commands as a sequence of amplitude-modulated (AM) bits. Once a power burst has ended, the reader signal will drop significantly in amplitude for some period of time. It is the duration of this “off-time” that communicates a bit value to the transponder. A short off-time duration specifies a ‘0’ bit, while a longer duration specifies a ‘1’ bit. Between each bit transmission, the reader signals returns to full amplitude in order to delimit the off-time intervals and maintain the powered state of the transponder. In some cases, after sending a command to a transponder, a reader will transmit a short, supplementary power burst to energize the tag fully.

Once the transponder has fully received and processed a command, it discharges its stored power, while transmitting its response using frequency

modulated frequency shift keying (FM-FSK). It communicates a bit via 16 RF cycles, specifying a ‘0’ or ‘1’ bit by transmitting at 134.2 kHz or 123.2 kHz respectively.⁵ A preamble of ‘0’ bits followed by a start byte (7E hex) indicates the start of a transmission and allows the reader to synchronize.

6. Conclusion

The weakness we have demonstrated in the TI system is ultimately due to the inadequate key-length of the underlying DST40 cipher. It is quite possible, however, that cryptanalysis will reveal weaknesses in the cipher itself. Indeed, we have preliminary experimental evidence that promises effective cryptanalytic attack. This would improve the efficacy of the attacks we have described. The authors hope that future cryptographic RFID system designers will embrace a critical lesson preached by the scientific community: Cryptographic hardware systems are generally strongest when they employ industry standard cryptographic algorithms with key lengths sufficient to endure over the life of the devices and assets they protect.

7. References

- [1] Automotive immobilizer anti-theft systems experience rapid growth in 1999, 1 June 1999. Texas Instruments Press Release. Available at http://www.ti.com/tiris/docs/news/news_releases/90s/re106-01-99.shtml
- [2] Speedpass™ Press Kit Fact Sheet, June 2004. Referenced at <http://www.exxonmobil.com>
- [3] Security and privacy in rfid systems, 2005. Web-based bibliography. Referenced at <http://lasecwww.epfl.ch/gavoine/rfid/>
- [4] ANONYMOUS USER. RC4?, September 1994. Sci.crypt posting.
- [5] BIHAM, E., AND SHAMIR, A. Differential fault analysis of secret key cryptosystems. In CRYPTO '97 (1997), B. Kaliski, Ed., Springer Verlag, pp. 513-525.
- [6] BIRYUKOV, A., SHAMIR, A., AND WAGNER, D. Real time cryptanalysis of A5/1 on a PC. In Fast Software Encryption (FSE) (2000), pp. 1_18.
- [7] BOURQUE, D. Technology update, chip status and development, October 2004. Slide presentation. Referenced at <http://www.ris.averydennison.com>