Improved Cryptanalysis of SecurID

Scott Contini
Computing Department
Macquarie University
NSW 2109 Australia
scontini@comp.mq.edu.au

Yiqun Lisa Yin
EE Department
Princeton University
Princeton, NJ 08540
yyin@princeton.edu

September 28, 2003

Abstract

SecurID is a widely used hardware token for strengthening authentication in a corporate environment. Recently, Biryukov, Lano, and Preneel presented an attack on the alleged SecurID hash function [1]. They showed that *vanishing differentials* – collisions of the hash function – occur quite frequently, and that such differentials could allow one to recover the secret key in the token much faster than exhaustive search. Based on simulation results, they estimated that given one 2-bit vanishing differential, the time of their attack would be equivalent to about 2^{48} full hash operations.

In this paper, we first give a more detailed analysis of the attack in [1], and then show how to improve it significantly. The bottleneck in their attack is one of the filtering steps. Our modifications to the filtering algorithm speed it up to the point that the time complexity becomes dominated by the number of final candidate keys to be tested, which results in a key recovery attack with time equivalent to about 2^{44} hash operations. No additional enhancement to the filtering steps can further reduce the number of candidates with only a single 2-bit vanishing differential.

We then investigate into the use of 4-bit vanishing differentials and multiple vanishing differentials, both of which do occur in practice. For 4-bit vanishing differentials, the running time of our attack is estimated to be equivalent to about 2^{40} hash operations. Multiple vanishing differentials appear to allow more speedups, but an exact running time is still to be determined.

1 Introduction

The SecurID is a hardware token developed by RSA Security. Its purpose is to strengthen authentication when logging in to remote systems, since passwords by themselves tend to be easily guessable and subject to dictionary attacks.

The SecurID adds an "extra factor" of authentication: one must not only prove themselves by getting their password correct, but also by demonstrating that they have the SecurID token assigned to them. The latter is done by entering the 6- or 8-digit code that is being displayed on the token at the time of login.

Each token has within it a 64-bit secret key and an internal clock. Every minute, or every half-minute in some tokens, the secret key and the current time are sent through a cryptographic hash function. The output of the hash function determines the next two authenticator codes, which are displayed on the LED screen display. The secret key is also held within the "ACE/server", so that the same authenticator can independently be computed and verified at the remote end.

If ever a user loses their token, they must report it so that the current token can be revoked and replaced with a new one. Thus, the user bears some responsibility in maintaining the security of the system. On the other hand, if the user were to temporarily leave his token in a place where it could be observed by others and then later recover it, then it should not be the case that the security of the device is entirely breached, assuming the device is well-designed.

The scenario just described was considered in a recent publication by Biryukov, Lano, and Preneel [1], where they showed that the hash function that is alleged to be used by SecurID [3] (AHSF) has weak properties that could allow one to find the key much faster than exhaustive search. The attack they describe requires recording all outputs of the SecurID using a PC camera with OCR software, and then later searching the outputs for indication of a *vanishing differential* – two closely related input times that result in the same output hash. If one is discovered, the attacker then has a good chance of finding the internal secret key using a search algorithm that they estimated to be equivalent to 2^{48} hash function operations. On a 2.4 GHz PC, 2^{48} hash operations take about 38 years. It would require about 450 of these PC's to find the key in a month, which is attainable by anyone in a typical medium-sized corporation.

The practicality of this attack depends upon how long the attacker must wait for a vanishing differential to occur - the longer the device is out of a user's control, the more likely that the user will recognise it and have it revoked. Simulations have shown that in any two month period, 10% of the SecurID cards will have a vanishing differential; in any two week period, 2.5% of the tokens will have a vanishing differential; and in any two day period, 0.2% of the tokens will have a vanishing differential. Although the attacker's success is not guaranteed, these probabilities are definitely not negligible. The attack described by Biryukov et al appears to be a real, practical threat to the device.

In this paper, we build upon their results. We first go through a deeper analysis of their algorithm, giving further justification of their conjectured running time of 2^{48}. We then show speedups that can be applied to reduce the time complexity down to about 2^{44} hash operations. With these improvements, the device can be attacked in under a month using only 30 modern PC's, which is attainable even within a small corporation. We also analyse special cases that can result in significant speedups to the attack: 4-bit vanishing differentials and multiple vanishing differentials. Both of these do happen in practice. Our results combined with those of [1] suggest that the SecurID cannot be relied on for strong second-factor authentication, assuming the description of the hash function [3] is correct. Note that RSA Security has begun to upgrade their tokens to use an AES based hash. We recommend that all of the older tokens be replaced with the upgraded AES based tokens.

2 The SecurID hash function

In this section, we provide a high level description of the alleged SecurID hash function. Detailed descriptions can be found in [1, 3]. We will follow the same notation as those in [1] wherever possible.

The function can be modeled as a keyed hash function $y = H(k, t)$, where k is a 64-bit secret key stored on the SecurID token, t is a 32-bit time obtained from the clock every 30 or 60 seconds, and y is two 6- or 8-digit codes. The function consists of the following steps:

- an expansion function that expands t into a 64-bit "plaintext",
- an initial key-dependent permutation,
- four key dependent rounds, each of which has 64 subrounds,
- an exclusive-or of the output of each round onto the key,
- a final key-dependent permutation (same algorithm as the initial one), and
- a key dependent conversion from hexadecimal to decimal.

Throughout the paper, we use the following notation to represent bits, nibbles, and bytes in a word: a 64-bit word b, consisting of bytes $B_0, ..., B_7$, nibbles $\mathtt{B}_0, ..., \mathtt{B}_{15}$, and bits $b_0 b_1 ... b_{63}$. The nibble $\mathtt{B}_0$ corresponds to the most significant nibble of byte 0 and the bit b_0 corresponds to the most significant bit. The other values are as one would expect.

For our analysis, only the key-dependent permutation and the key dependent rounds are of interest. In the next two sections, we will describe them in more detail.

2.1 Key dependent permutation

We give a more insightful description of how the ASHF key dependent permutation really works. The original code, obtained by I.C. Wiener [3] (apparently by reverse engineering the ACE/server code), is quite cryptic. Our description is different, but produces an equivalent output to his code.

The key dependent permutation uses the key nibbles $\mathtt{K}_0 \ldots \mathtt{K}_{15}$ in order to select bits of the data for output into a permuted_data array. The data bits will be taken 4 at a time, copied to the permuted_data array from right to left (i.e. higher indexes are filled in first), and then removed from the original data array. Every time 4 bits are removed from the original data array, the size shrinks by 4. Indexes within that array are always modulo the number of bits remaining.

A pointer m is first initialised to the index $\mathtt{K}_0$. The first 4 bits that are taken are those right before the index of m. For example, if $\mathtt{K}_0$ is 0xa, then bits 6, 7, 8, and 9 are taken. If $\mathtt{K}_0$ is 0x2, then bits 62, 63, 0, and 1 are taken. As these bits are removed from the array, the index m is adjusted accordingly so that it continues to point at the same bit it pointed to before the 4 bits were removed.

The pointer m is then increased by a value of $\mathtt{K}_1$, and the 4 bits prior to this are taken, as before. The process is repeated until all bits have been taken.

Note that once the algorithm gets down to the final 3 or less key and data nibbles, the number of data bits remaining is at most 12 yet the number of choices for each key nibble is 16. Hence, multiple keys will result in the same permutation, which we call "redundancy of the key with respect to the permutation." This was used in the attack [2], and to a lesser extent in [1]. Interestingly, [1] mentions that there are 14-bits of redundancy on average, yet the attacks presented so far have exploited only a few of them.

2.2 Key dependent rounds

Each of the four key dependent rounds takes as inputs a 64-bit key k and a 64-bit value b^0, and outputs a 64-bit value b^{64}. The key k is then exclusive-ored with the output b^{64} to produce the new key to be used in the next round.

One round consists of 64 subrounds. For $i = 1, ..., 64$, subround i transforms b^{i-1} into b^i using a single key bit k_{i-1}. Depending on whether the key bit k_{i-1} is equal to b_0^{i-1}, the value b^{i-1} is transformed according to two different functions, denoted by R and S. This particular property causes the hash function to have many easy-to-find collisions (called vanishing differentials) after a small number of subrounds within the *first* round. At the end of each subround, all the bits are shifted one bit position to the left.

We remark that both the R function and the S function are byte-oriented, that is, they update each of the 8 bytes in b separately. After the update, only *two* of the 8 bytes (B_0 and B_4) are modified, and the rest of the 6 bytes remain the same.

3 The attack of Biryukov, Lano, and Preneel

Biryukov, Lano, and Preneel recently presented a full key recovery attack that uses a single 2-bit vanishing differential. The attacker first guesses the subround N in which the vanishing differential occurs, and for each N a filtering algorithm is used to search the set of candidiate keys that make such a vanishing differential possible. In [1], they only described the attack for $N = 1$ and stated that the algorithm would be similar for other N. According to their simulations, one only needs to do up to $N = 12$ to have a 50% chance of finding the key.

Here we give a high-level description of the filtering algorithm for $N = 1$. At the beginning, a table with entries of the form

$$(k_0, B_0, B_4, B'_0, B'_4)$$

is precomputed. The entries contain all combinations of key bit k_0 and data bytes B_0, B_4, B'_0, and B'_4 going into the first round (i.e. after the initial permutation) that will result in a vanishing differential at the end of the first subround. Note that none of the other data bytes have any involvement in the first subround, so whether a vanishing differential can happen or not for $N = 1$ is completely characterised by this table.

The filtering algorithm proceeds in four steps.

- *First Step.* For each entry in the precomputed table, try all possible values of $k_1, ...k_{27}$. Together with k_0, 28 key bits are set, which determines 28 bits of b^0 from the initial key-dependent permutation. Since these bits overlap with the entries in the table by one nibble $\mathtt{B}_9$, key values that do not produce the correct nibble for both plaintexts in the vanishing differential are filtered out.

- *Second Step.* An entry that passes the first step is taken as input and the key bits $k_{28}, ...k_{31}$ are guessed. Filtering is done based on the overlap in nibble $\mathtt{B}_8$.

- *Third Step.* Key bits $k_{32}, ..., k_{59}$ are guessed. Filtering is done based on the overlap in nibble $\mathtt{B}_1$.

- *Fourth Step.* Key bits $k_{60}, ..., k_{63}$ are guessed. Filtering is done based on the overlap in nibble $\mathtt{B}_0$.

Finally, each candidate key that passes the filtering steps is tested by performing a full hash function to see if it is the correct key.

As we can see, the running time of the above attack depends on the time complexity of each filtering step and the number of candidate keys that pass all four filtering steps. Based on simulation results [1], they estimated that the dominant factor is the third filtering step, which is equivalent to about 2^{48} full hash operations for N up to 12.

4 Improved analysis of the Biryukov, Lano, and Preneel attack

Biryukov, Lano, and Preneel only gave simulated results for $N = 1$. They suggested that

> for higher N, the overlap will be higher (because more bits play a role in the vanishing differential) and thus the filtering will be stronger.

and

> For $N > 1$, we expect the complexity of the attack to be lower due to stronger filtering.

Here we show that the results of their simulations can be justified by mathematical arguments, and that the conjecture of the filtering improving for larger N appears to be correct. We first analyse the case $N = 1$, and then generalise the argument to arbitrary N. Our analysis is an average-case analysis. The actual time complexity will depend upon the particular pair of plaintexts that is used in the attack.

There is one subtlety that the reader should keep in mind in our analysis. During the first two filtering steps, only the values (k_0, B_4, B'_4) of the precomputed table are involved. There may be more than one table entry overlapping in these values. In this case, we assume that the multiple entries are grouped together into a single entry until a later filtering step requires testing for the overlap separately. Since, as we will see, the number of multiple entries is very small, we assume that this does not incur a noticeable speed penalty.

4.1 Analysis of the attack for $N = 1$

Their simulations showed that the first step reduced the number of possibilities to 2^{27}, the second step further reduced the count to to 2^{25}, the third step increased the count to 2^{45}, and the fourth step resulted in 2^{41} true candidates. We analyse the second and fourth steps only: the other two can be analysed similarly.

We note that some properties of the precomputed table are necessary in the analysis. In [1], it is stated that the size of the precomputed table is 30 for $N = 1$, which we agrees with our computation. In Appendix A, we provide an analytical way of constructing the entries.

Analysis of the second step: We start by examing the precomputed table to count number of *unique* entries of the form (k_0, B_4, B'_4). In total, there are only 23, which is broken down into 7 with no difference, 16 with a 1-bit difference, and none with 2-bit differences.

There are a total of 2^{32} possible partial keys (each 32 bits) up to step two. Among them,

- A fraction of $\binom{56}{2}/\binom{64}{2} \approx .76$ will put no difference in the tuple (B_4, B'_4).
- A fraction of $\binom{8}{1} \times \binom{56}{1}/\binom{64}{2} \approx .22$ will put a 1-bit difference in (B_4, B'_4).
- A fraction of only $\binom{8}{2}/\binom{64}{2} \approx .01$ will put 2 difference bits in (B_4, B'_4).

Of the $2^{32} \times 0.76$ keys that result in no difference in (B_4, B'_4), only a fraction of $\frac{7}{256}$ will match one of the 7 unique entries in the table for B_4 (which is the same as B'_4). Of those, only half will have the right key bit corresponding to what is stored for that entry of the table. Thus, the expected number of 32-bit keys resulting in no difference in B_4 that pass the second filtering step is

$$2^{32} \times 0.76 \times \frac{7}{256} \times \frac{1}{2} \approx 2^{25.4}.$$

For 1-bit differences, the calculation is similar, except we have 16 unique table entries, and there are 256×8 possible tuples (B_4, B'_4) with $B_4 \bigoplus B'_4$ differing in 1-bit. The expected number here is

$$2^{32} \times 0.22 \times \frac{16}{256 \times 8} \times \frac{1}{2} \approx 2^{21.8}.$$

For 2-bit differences, there are 0 in the table, so none of those will get through.

Combining these results, the expected number of 32-bit keys that pass through step 2 is

$$2^{25.4} + 2^{21.8} \approx 2^{25.5}$$

which closely agrees with the 2^{25} observed by simulation in [1].

Analysis of the fourth step: Without considering outcomes of previous steps, we can directly analyse the fourth step. This is because anything that matches an entry in the precomputed table will result in a vanishing differential for $N = 1$. In other words, the entries in the table are not only a necessary set of cases for a vanishing differential to occur at $N = 1$, but also sufficient. So, analysing the outcome of the fourth step is equivalent to determining the true number of candidates that need to be tested with the full SecurID hash function.

For each of the 30 table entries, we have:

- Only a portion of about $\frac{1}{2^{16}}$ of the 2^{64} keys will permute the bits of the first plaintext so that the bytes (B_0, B_4) match the table entry.
- Of those keys, only a portion of $1/\binom{64}{2}$ will permute the 2 difference bits in the right locations to match the (B'_0, B'_4) of that table entry.
- Only half of those keys will have the right key bit k_0 corresponding to what is in that entry of the table.

Thus, the expected number of final candidate keys is

$$30 \times 2^{64} \times \frac{1}{2^{16}} \times \frac{1}{\binom{64}{2}} \times \frac{1}{2} \approx 2^{40.9}$$

which is approximately 2^{41} that was observed in [1].

Another way of interpreting this result is that the probability of a randomly chosen 2-bit differential disappearing in subround 1 is $\frac{2^{40.9}}{2^{64}} \approx 2^{-23.1}$. This property will be useful in our later analysis.

4.2 Analysis of the attack for $N > 1$

Here we derive general formulas for the number of candidate keys that will pass the second and fourth steps, respectively, as well as the time complexity for the third step. As we discussed before, these are the dominating factors in estimating the running time of the attack. Similar to the case of $N = 1$, the formulas depend upon properties of the precomputed tables.

In the general case, the precomputed tables consist of the following entries:

- legal values for the key bits in indices $0, \ldots, N-1$,
- legal values for the plaintext pairs after the initial permutation in bit indices $32, 33, \ldots, 38 + N$ which we label as (W_4, W_4'), and
- legal values for the plaintext pairs after the initial permutation in bit indices $0, 1, \ldots, 6 + N$ which we label as (W_0, W_0').

By legal values we mean that the combination of key bits and plaintext bits will cause the difference to vanish in subround N. The words W_0, W_0', W_4, W_4' each consist of $7 + N$ bits and the number of key bits is N. Using this notation, observe that when $N = 1$ we have $(W_4, W_4') = (B_4, B_4')$ and $(W_0, W_0') = (B_0, B_0')$.

Analysis of the second step: Of the of 2^{32} key bits considered up to step 2,

- A fraction of $\binom{57-N}{2}/\binom{64}{2}$ will put no difference in the tuple (W_4, W_4').
- A fraction of $\binom{7+N}{1} \times \binom{57-N}{1}/\binom{64}{2}$ will put a 1-bit difference in (W_4, W_4').
- A fraction of only $\binom{7+N}{2}/\binom{64}{2}$ will put a 2-bit difference in (W_4, W_4').

Define C_0 to be the number of unique table entries of the form $(k_0, \ldots, k_{N-1}, W_4, W_4')$ where $W_4 = W_4'$, C_1 similarly except $W_4 \bigoplus W_4'$ having hamming weight 1, and C_2 similarly except $W_4 \bigoplus W_4'$ having hamming weight 2. The expected number of keys causing no bit difference in (W_4, W_4') that will pass the filter in step two is:

$$2^{32} \times \frac{\binom{57-N}{2}}{\binom{64}{2}} \times \frac{C_0}{2^{7+N}} \times \frac{1}{2^N}$$

$$= 2^{19-2N} \times \frac{3192 - 113N + N^2}{63} \times C_0.$$

For 1-bit differences, the equation is

$$2^{32} \times \frac{\binom{7+N}{1} \times \binom{57-N}{1}}{\binom{64}{2}} \times \frac{C_1}{2^{7+N} \times \binom{7+N}{1}} \times \frac{1}{2^N}$$

$$= 2^{20-2N} \times \frac{57 - N}{63} \times C_1.$$

For 2-bit differences, the equation is

$$2^{32} \times \frac{\binom{7+N}{2}}{\binom{64}{2}} \times \frac{C_2}{2^{7+N} \times \binom{7+N}{2}} \times \frac{1}{2^N}$$

$$= 2^{20-2N} \times \frac{C_2}{63}.$$

Hence, the expected number of candidates to pass the second step is then

$$T = \frac{2^{19-2N}}{63} \times \left[(3192 - 113N + N^2)C_0 + (114 - 2N)C_1 + 2C_2\right]. \qquad (1)$$

In [1], the third step is the most time consuming. For each candidate that passes the second step, they must guess 28 bits of key and then perform a fraction of $\frac{28}{64}$ of the permutation for both plaintexts. Under the assumption that the permutation is 5% of the time required to do the full SecurID hash, the running time is equivalent to

$$T \times 2^{28} \times \frac{28}{64} \times 2 \times 0.05$$

full hash operations.

Note that when deriving the above formula, we assumed that for larger N, exactly four filtering steps (same as what was done when $N = 1$) were used. The filtering algorithm was not completely described for $N > 1$ in [1], but it is likely that they imagined that the number of filtering steps would increase. In particular, one may presume that the third step would involve guessing enough key bits so that the resulting permuted data bits just begin to overlap with W_0 and W_0', and an additional layer of filtering would be added for each key nibble guessed beyond that[1]. This speeds up the third step, which we will assume is still the most time consuming of the remaining filtering steps[2]. We proceed under this assumption.

[1] The same idea should be applied to the first filtering step as well, but for the sake of brevity, we avoid making the description too complex.

[2] Omitting details, the third step will be the most time consuming if the fraction of values that remain is less than $\frac{\lfloor \frac{29-N}{4} \rfloor}{16}$ of the values considered. This is usually the case. In the rare exceptions, the fourth step will be slightly more time consuming. One should allow for a very small error (example: half a bit) in our final run time estimate because of this short-cut in the analysis.

N	table size	C_0	C_1	C_2	T	Time for third step	Time for last step	Total time
1	30	7	16	0	$2^{25.5}$	$2^{47.6}$	$2^{40.9}$	$2^{47.6}$
2	350	24	128	84	$2^{25.4}$	$2^{44.2}$	$2^{41.5}$	$2^{44.4}$
3	2366	171	660	248	$2^{26.1}$	$2^{45.0}$	$2^{41.2}$	$2^{45.1}$
4	16784	1047	3778	1392	$2^{26.7}$	$2^{45.5}$	$2^{41.0}$	$2^{45.6}$
5	116184	6349	22700	8264	$2^{27.2}$	$2^{46.1}$	$2^{40.8}$	$2^{46.1}$
6	729236	37257	125824	42836	$2^{27.7}$	$2^{42.7}$	$2^{40.5}$	$2^{43.0}$

Table 1: Computing the running time estimates of algorithm [1] for $N = 1..6$.

In this way, the exact number of key bits guessed in the third step is $4 \times \lfloor \frac{29-N}{4} \rfloor$, and its running time is

$$T \times 2^{4\times\lfloor \frac{29-N}{4} \rfloor} \times \frac{4 \times \lfloor \frac{29-N}{4} \rfloor}{64} \times 2 \times 0.05 \times s \tag{2}$$

full hash operations, where s is the speedup factor that can be obtained by taking advantage of the redundancy in the key with respect to the permutation. The value of s is $\frac{96}{256}$ for $N = 1$, $\frac{12}{16}$ for $N = 2..5$, and 1 for all other values.

Analysis of the fourth step: Following section 4.1, the general formula for the number of final candidates is:

$$\text{table size} \times 2^{64} \times \frac{1}{2^{2N+14}} \times \frac{1}{\binom{64}{2}} \times \frac{1}{2^N} \approx 2^{39.0-3N} \times \text{table size.} \tag{3}$$

Combined analysis: The running time of algorithm [1] for a particular value of N is expected to be the approximately the sum of equations 2 and 3. For $N = 1..6$, these running times are given in Table 1.

Notice that even though the number of candidates T after the second filter steps are approximately the same as N goes from 1 to 2 and also from 5 to 6, the running times of the third steps drop greatly. This is because one less nibble of the key is being guessed, and an extra filtering step is being added. In general, we see the pattern that larger values of N are contributing less and less to the sum of the running times, which agrees with the conjecture from [1]. The total running time for $N = 1$ to 6 is $2^{48.5}$ and larger values of N would appear to add minimally to this total.

5 Faster filtering

As illustrated in the previous section, the trick to speeding up the key recovery attack in [1] is faster filtering. We have found three ways in which their filtering can be sped up:

1. In the original filter, a separate permutation is computed for each trial key. This is inefficient, since most of the permuted bits from one particular permutation will overlap with those from many other permutations. Thus, we can amortize the cost of the permutation computations.

2. We can detect ahead of time when a large portion of keys will result in "bad" permutations in steps 1 and 3, and the filtering process can skip past chunks of these bad permutations.

3. For $N = 1$, we can further speed up the third step of filtering by using a table-lookup to determine what the legal choices are for $\mathtt{K}_{14}$ (this would apply to other steps as well, but the memory requirements quickly become quite large). Each table lookup replaces trying 8 choices for the nibble $\mathtt{K}_{14}$.

In what follows, we describe each of the above techniques in more detail.

The first technique is aimed at reducing the numerator of the factor $\frac{4\times\lfloor\frac{29-N}{4}\rfloor}{64} = \frac{\lfloor\frac{29-N}{4}\rfloor}{16}$ in equation 2. To do this, we view the key as a 64-bit counter, where k_0 is the most significant bit and k_{63} is the least. In step three of the filter, the bits $k_0, \ldots, k_{31}$ are fixed and so are some of the least significant bits (the exact number depends upon N), so we can exclude these for now. The keys are tried in order via a recursive procedure that handles one key nibble at a time. At the j^{th} recursive branch, each of the possibilities for nibble $\mathtt{K}_{7+j}$ are tried. The part of the permutation for that nibble is computed, and then the $j+1^{\text{st}}$ recursive branch is taken. The level of recursion stops when key nibble $\mathtt{K}_{7+\lfloor\frac{29-N}{4}\rfloor}$ is reached. Thus, the $\lfloor\frac{29-N}{4}\rfloor$ from equation 2 gets replaced with the average cost per permutation trial, which is

$$\sum_{i=0}^{\lfloor\frac{29-N}{4}\rfloor-1} 2^{-4i} \approx 1.07.$$

Observe that when $N = 1$, this results in a factor of $\frac{7}{1.07} \approx 6.5$ speedup. This trick alone knocks more than 2 bits off the running time.

The second speedup is dependent upon the first. It will apply to both the first and third filtering steps. During the process of trying a permutation, there will be large chunks of bad trial keys that can be identified immediately, and skipped. For example, consider $N = 1$ in the first filtering step. Whenever one of the difference bits is put into any of the bit indices 40..63 of the permuted data array, it can be skipped because the difference is not in a legal position. More generally, in the recursive procedure for key trials, we check during each trial key nibble whether it will result in a difference bit being put in an illegal place. If affirmative, then any key having the same most significant bits will also result in misplacing the difference bit, so the recursive branch for that key

N	Time for third step	Time for last step	Total time
1	$2^{37.8}$	$2^{40.9}$	$2^{41.0}$
2	$2^{38.0}$	$2^{41.5}$	$2^{41.6}$
3	$2^{39.0}$	$2^{41.2}$	$2^{41.5}$
4	$2^{39.9}$	$2^{41.0}$	$2^{41.6}$
5	$2^{40.7}$	$2^{40.8}$	$2^{41.8}$
6	$2^{37.8}$	$2^{40.5}$	$2^{40.7}$

Table 2: Running times using our improved filter, for $N = 1..6$.

nibble can be skipped. This substantially reduces the number of trial keys. In the first step, this will skip past all but a fraction $\binom{39+N}{2}/\binom{64}{2}$ of the candidates. More importantly, between the first and third steps the amount of keys looked at in the search becomes a fraction $\binom{14+2N}{2}/\binom{64}{2}$ of the amount for the attack in [1].

These two strategies combined result in the following running time for the third filtering step:

$$T \times \frac{\binom{14+2N}{2}}{\binom{64}{2}} \times 2^{4\times\lfloor\frac{29-N}{4}\rfloor} \times \frac{1.07}{16} \times 2 \times 0.05 \times s \tag{4}$$

where T is still the T from equation 1 (though it no longer represents the number of candidates from step one) and s is $\frac{96}{256}$ for $N = 1$, $\frac{12}{16}$ for $N = 2..5$, and 1 for all other values.

We only apply the third speedup for $N = 1$, due to increasing memory requirements. When we arrive at a leaf to try $\mathtt{K}_{14}$, there are only 8 data bits remaining to choose from. Let x represent the final 8 bits for the first plaintext, and x' for the second. We could precompute the legal choices for $\mathtt{K}_{14}$ for each possible (k_0, B_4, B'_4, x, x') where (k_0, B_4, B'_4) are from the 23 unique choices in the main filtering precomputation table. Thus the legal choices for $\mathtt{K}_{14}$ are obtained from a single table lookup, which replaces trying all possibilities. For $N = 1$, this gives a time of approximately:

$$T \times \frac{\binom{16}{2}}{\binom{64}{2}} \times 2^{24} \times \frac{1.07}{16} \times 2 \times 0.05 \times \frac{12}{16}.$$

The combined speedups give the run times in Table 2. In all cases, the third filtering step has become faster than the time for the last step.

The total time for $N = 1..6$ is $2^{44.0}$, and larger values of N are expected to add minimally to it since all steps are getting faster.

N	table size	run time
1	910	$2^{37.5}$
2	9202	$2^{37.9}$
3	53358	$2^{37.4}$
4	311566	$2^{37.0}$

Table 3: Cost of the final step using a 4-bit differential for $N = 1..4$.

6 Vanishing differentials with four-bit difference

For both the attacks in [1] and our improved attack presented in Section 5, only a *single* vanishing differential with a *two*-bit difference were used. Although such differential provides sufficient information to derive the entire secret key, it also limits us in terms of the efficiency of a search algorithm since the number of final candidates is always on the order of 2^{41} for the N of interest. By allowing other forms of vanishing differentials, we have a chance of further reducing the complexity.

According to our simulations, about 25% of the first collisions (first occurrence of a vanishing differential for a give key) are actually from a 4-bit difference. We would expect that our filtering algorithm performs exceptionally well in this circumstance. For example, when $N = 1$ we expect our second filtering speedup to skip all except a fraction of $\binom{16}{4}/\binom{64}{4} \approx 2^{-8.4}$ of the incorrect keys between filter steps one through three. Without going through the analysis, it seems reasonable to assume that the final testing of candidates is still the bottleneck.

The formula for number of final candidate keys can be derived similar to that of equation 3:

$$\text{table size} \times 2^{64} \times \frac{1}{2^{2N+14}} \times \frac{1}{\binom{64}{4}} \times \frac{1}{2^N}.$$

The formula is the same as that for 2-bit differences, except that term $\binom{64}{2}$ has been replaced by $\binom{64}{4}$, giving a factor of $2^{8.3}$ reduction in the number. Therefore, as long as the table size does not increase significantly, it is conceivable that 4-bit differentials could result in a faster attack than 2-bit differentials. In Table 3 we see that this is indeed the case for $N = 1..4$. The table size for $N = 1$ can also be verified analytically as described in Appendix A. We therefore conjecture that the total run time for an attack using one 4-bit vanishing differentials is equivalent to about 2^{40} hash operations.

Note that for $N = 1$, we have a probability of $\frac{2^{37.5}}{2^{64}} = 2^{-26.5}$ for a 4-bit vanishing differential to occur, and the corresponding probability for a 2-bit vanishing differential is $2^{-23.1}$. It may then seem hard to believe that 25% of the vanishing differentials are 4-bits, as claimed above. However, one should keep in mind that there are more input 4-bit differences because the least significant byte of the time is replicated 4 times in the time expansion function.

7 Multiple vanishing differentials

According to simulation results in [1], a 2-bit vanishing differential occurs with probability 2^{-19}. This translates to the fact that 10% of the tokens will have at least one vanishing differential in a two-month period. If vanishing differentials do occur consistently with the above probability, then we expect that 10% of the tokens will have at least *two* vanishing differentials in a *four*-month period. Experiments are being performed to determine more accurate numbers.

Now we show how to reduce the number of candidate keys significantly given two vanishing differentials by constructing more effective filters in each step. We denote the two pairs of vanishing differentials V_1 and V_2, and their N values N_1 and N_2.

We first make a guess of (N_1, N_2). The number of guesses will be quadratic in the number of subrounds tested up to. The following is a sketch for the new filtering algorithm when $N_1 = N_2 = 1$. Other cases can be handled similarly.

- *First Stage.* Take V_1 and guess the first 32 bits of the key. For each 32-bit key that produces a valid (B_4, B'_4), test it against V_2 to see if it also produces a valid (B_4, B'_4). (This is the first and the second filtering steps in the original attack.)

- *Second Stage.* For 32-bit keys that pass the above stage, do the same thing to guess the second 32 bits of the key. (This is the third and the fourth filtering steps in the original attack.)

The main idea here is to do double filtering within each stage so that the number of candidate keys is further reduced in comparison to when only a single vanishing differential is used.

Based on early analysis (in Section 4), we know that the probability that a 32-bit key passes the first stage is $2^{25.5}/2^{32} = 2^{-6.5}$ (assuming using the original filter of [1] - it is even more reduced using our improved filter), and the probability that a 64-bit key passes both stages is $2^{40.9}/2^{65} = 2^{-23.1}$. If the two vanishing differentials are indeed *independent*, we would expect the number of keys to pass the first filtering to be

$$2^{32} \times 2^{-6.5} \times 2^{-6.5} = 2^{19}, \text{and}$$

and the number of keys to pass both filterings to be

$$2^{64} \times 2^{-23.1} \times 2^{-23.1} = 2^{17.8}.$$

Preliminary experiments show that vanishing differentials often occur, but usually the same difference appears throughout. There are some occasions, however, where different difference pairs happen. We are still in the process of testing whether the figures above are attainable in practice.

We should also mention the caveat that the chances of success using multiple vanishing differentials are lower, since we need *both* difference pairs to disappear within N subrounds. On the other hand, the cost of trying this algorithm for two difference pairs is expected to be substantially cheaper than trying the other algorithm for only one. Therefore, the double filtering should add negligible overhead to the search in the cases that it fails, and would greatly speedup the search when it is successful.

8 Conclusion

The design of the alleged SecurID hash function appears to have several problems. The most serious appears to be collisions that happen far too frequently and very early within the computation. The amount to which such collisions can be taken advantage of is exacerbated by subrounds that only involve a small fraction of data bits. Moreover, the redundancy of the key with respect to the initial permutation adds an extra avenue of attack. Altogether, ASHF is substantially weaker than one would expect from a modern day hash function.

Our research has shown that the key recovery attack in [1] can be sped up by a factor of 16, giving an improved attack with time complexity about 2^{44} hash operations. We have also illustrated special data cases that can be attacked significantly faster, and such cases do happen in practice. For example, 25% of the collisions involve 4-bit vanishing differentials, which seem to allow one to recover the key in 2^{40} hash operations.

The attacks in this paper and in [1] are real. The main obstacle in mounting them is waiting for an internal collision. If the user's token is out of his control for a matter of a few days, then the chances of the collision happening are small, but not negligible. This means that most attackers will not have the opportunity for success, but some will. On the other hand, an AES-based hash function ought to prevent any attacker from having a realistic chance of success. We therefore recommend that all SecurID cards containing the alleged hash function be replaced with RSA Security's newer, AES-based hash.

References

[1] A. Biryukov, J. Lano, and B. Preneel. *Cryptanalysis of the Alleged SecurID Hash Function*, http://eprint.iacr.org/2003/162/, 12 Sep, 2003.

[2] S. Contini, *The Effect of a Single Vanishing Differential in ASHF*, sci.crypt post, 6 Sep, 2003.

[3] I.C. Wiener, *Sample SecurID Token Emulator with Token Secret Import*, post to BugTraq, http://archives.neohapsis.com/archives/bugtraq/2000-12/0428.html , 21 Dec, 2000.

A Analysing precomputed tables

Using computer experiments, we were able to exhaustively search for valid entries in the precomputed table up to $N = 6$ for 2-bit vanishing differentials and up to $N = 4$ for 4-bit differentials at this point. It was predicted in [1] that the size of the table gets larger by a factor of 8 as N grows and it may take up to 2^{44} steps and 500GB memory to precompute the table for $N = 12$.

Here we make an attempt to derive the entries in the table analytically when $N = 1$. If we could extend the method to $N > 1$, we may be able to enumerate the entries analytically without expensive precomputation and storage.

We start with Equation (6) in [1]. Note that we are trying to find constraints for the values in the subround $i-1$. So for simplicity, we will omit the superscript $i-1$ from now on, and Equation (6) becames the following.

$$\begin{aligned} B'_4 &= ((((B_0 >>> 1) - 1) >>> 1) - 1) \oplus B_4, \\ B'_0 &= 100 - B_4. \end{aligned} \tag{5}$$

We first note that B_0 and B'_0 have to be different in the msb. Therefore, there is at least one bit difference in (B_0, B'_0). The other bit difference can be placed either in the remaining 7 bits of (B_0, B'_0) or any of the 8 bits in (B_4, B'_4).

Rewrite Equation 5, we have

$$B_0 = (((B_4 \oplus B'_4) + 1) <<< 1) + 1) <<< 1.$$

Since there are at most one bit difference in (B_4, B'_4), it can only take on 9 possible values: 0 (for no bit difference) or 2^i (for one bit difference in bit i). For each possible value of (B_4, B'_4), we enumerate the possible values of (B_0, B'_0) as follows.

- If $B_4 \oplus B'_4 = 0$, then B_0 =0x06. Since there is no bit difference in (B_4, B'_4), we know that B_0 and B'_0 differ in two bits – one of them must be the msb, and the other can be any of the remaining 7 bits.

$B_4 \oplus B'_4$	B_0	B'_0	k_0
0x00	0x06	0x87, 84, 82, 8e, 96, a6, c6	0

- If $B_4 \oplus B'_4 = 2^i$, then there is only one bit difference in (B_0, B'_0), which is the msb. In this case, there are only one choice for B'_0 for each B_0.

$B_4 \oplus B_4'$	B_0	B_0'	k_0
0x01	0x0a	0x8a	0
0x02	0x0e	0x8e	0
0x04	0x16	0x96	0
0x08	0x26	0xa6	0
0x10	0x46	0xc6	0
0x20	0x86	0x06	1
0x40	0x07	0x87	0
0x80	0x08	0x88	0

Combining the above two cases, we have $8 + 7 = 15$ pairs of (B_0, B_0'), each of which giving a valid tuple $(k_0, B_0, B_4, B_0', B_4')$, where k_0 is the msb of B_0.

Finally, note that if (k_0, a, b, c, d) is a valid tuple, than (k_0, c, d, a, b) is also a valid typle. For example, if $(0, 0x06, 0xdd, 0x87, 0xdd)$ is valid, then $(0, 0x87, 0xdd, 0x06, 0xdd)$ is also valid. Therefore, the table consists of a total of $2 \times 15 = 30$ entries. These entries match the results from our simulation.

Similar analysis also confirms that the size of the table for 4-bit vanishing differentials when $N = 1$ is 910.