From: "mrobshaw" <mrobshaw@supanet.com>
To: <AESround2@nist.gov>
Cc: "Lisa Yin" <yiqun@nttmcl.com>
Subject: comments on mars
Date: Tue, 2 May 2000 08:51:06 +0100
X-MSMail-Priority: Normal
X-Mailer: Microsoft Outlook Express 5.00.2314.1300
X-MimeOLE: Produced By Microsoft MimeOLE V5.00.2314.1300

Dear Sirs,

At the rump session of the third AES conference I presented some joint results due to Lisa Yin and myself on the linear cryptanalysis of MARS.

Please find attached a more complete note and summary of our analysis and findings.

Best wishes,
Matt Robshaw

Potential Flaws in the Conjectured Resistance of MARS to Linear Cryptanalysis

Extended abstract — April 27, 2000

M.J.B. Robshaw[1] and Yiqun Lisa Yin[2]

[1] 88 Hadyn Park Road, London, W12 9AG
mrobshaw@supanet.com
[2] NTT Multimedia Communications Laboratories,
250 Cambridge Ave., Palo Alto, CA 94306
yiqun@nttmcl.com

Abstract. In this note we consider the conjectured resistance of MARS to linear cryptanalysis and discover that some of the existing analysis may well be flawed.

1 Introduction

As the AES process nears its conclusion, five algorithms remain for consideration. In this note we describe the preliminary findings of an investigation of the resistance of MARS to linear cryptanalysis. Our attention focuses on the main source of cryptographic strength in MARS [2], the so-called `E` function and its combination in successive rounds with other instances of the same function.

The designers of MARS justify their claims for the resistance of MARS to differential and linear cryptanalysis by providing [2] "crude (though conservative) bounds on the complexity of such attacks." After some analysis and extensive experimentation, it is our conclusion that these bounds for resistance to linear cryptanalysis could be flawed.

In particular, we have shown that the style of analysis used by the MARS designers seems to provide an upper bound of only 2^{-49} for the bias of a linear approximation to the cryptographic core. (Their analysis showed a bound of 2^{-69}.) Whether or not a linear approximation can be found that meets the bias of 2^{-49} is, of course, unknown. While unlikely, if such an approximation did exist then the data requirements for an attack on the cryptographic core of MARS would be of the order of 2^{98} plaintexts.

2 MARS

MARS is one of the five finalists for the AES and it has many novel and successful design features. It is however, an exceedingly complicated cipher to analyze. Given that the time for public cryptanalysis in the second round is so short, it is questionable how much attention has been made to the detailed design of this and other more complicated AES candidates [9].

Our attention in this note is almost exclusively focused on the role of what is termed the `E` function [2]. It consists of one 32-bit word of input I and produces three 32-bit words of output L, M, and R. The action of this function is illustrated graphically in Figure 1 where left rotations by a fixed amount have the rotation amount indicated and rotation amount for the variable left rotations is indicated by the source of the arrow.

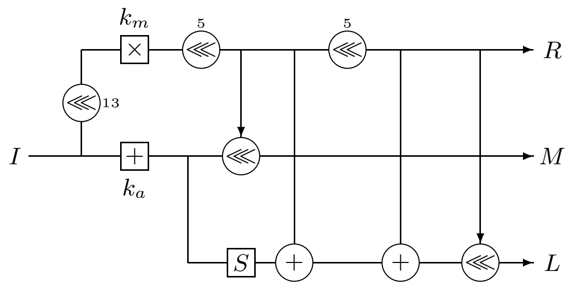

Fig. 1. The `E` function in MARS.

3 Linear cryptanalysis

Linear cryptanalysis [6, 7] is an intriguing style of analysis. While it is very effective against DES [7] it is typically not as successful as differential cryptanalysis [1] in the analysis of other ciphers. It is also a style of analysis for which there remain many complex and unanswered questions [10].

The aim of a linear cryptanalytic attack is to find an effective linear expression connecting some bits of the intermediate text related to the plaintext, some bits of the intermediate text related to the ciphertext, and some bits of the key. When the probability that such an approximation holds is biased, by taking sufficiently many plaintext/ciphertext pairs the correct value of a bit of key information can be identified. The greater the bias, the fewer the number of plaintext/ciphertext pairs needed and the data requirements for a linear cryptanalytic attack are inversely proportional to the square of the bias of the approximation [6].

In this note, we follow the style of analysis of the MARS designers but we demonstrate that some upper bounds on the bias of basic approximations used as building blocks in their assessment are incorrect. We do this using simple and established techniques. We also wonder whether established results on the resistance of MARS to linear cryptanalysis are that conservative when advanced techniques such as *linear hulls* [8], *multiple linear approximations* [3, 4], and the effects of *key dependence* [10], are not considered.

3.1 Notation

The approximations that we consider in this note involve four 32-bit words that are denoted [2] as I, L, M, and R (Figure 1). An approximation involving n bits from these 32-bit words in positions t_i $(1 \leq i \leq n)$ from word I (say) will be denoted as $I[t_1, \ldots, t_n]$. We consider the rightmost bit of a 32-bit word to be the least significant bit and denote its position by 0.

In a slight abuse of notation we will consider a word as a vector in Z_2^{32} and we will use Γ (say) to indicate the bits of the word that are to be used in a linear approximation. This is most conveniently described by means of the *scalar product* of two vectors. Thus the $\{0,1\}$-vector Γ might denote the bits of I to be used in an approximation and $I \cdot \Gamma$ is the value of these bits combined using exclusive-or. An example linear approximation might be written $(I \cdot \Gamma_1) \oplus (M \cdot \Gamma_2) = 0$.

4 The S-box

The S-box $S(\cdot)$ in MARS gives a 32-bit output that is chosen from 512 using a nine-bit input. The construction of $S(\cdot)$ is based on the hash function SHA-1 by choosing an appropriate set of parameters. On page 31 of [2] the designers of MARS conjecture that there are

> " ... no approximations of the S-box with bias of more than 2^{-3}."

Some of the estimated upper bounds for the bias of linear approximations to the E-function are based on this conjecture. In reality, a quick search of a small fraction of the total number of possible linear approximations[3] reveals many linear approximation to the S-box that have a bias greater than 2^{-3}.

Here we describe some experimental results on the bias of selected linear approximations for the S-box. There are a few interesting types of linear approximations to consider which may prove beneficial in a wider linear-cryptanalytic attack. Among them is the interesting case where a linear approximation to the S-box involves no input bits but some of the $2^{32}-1$ possible approximations of the output bits. Let $\Gamma_x \in \{0,1\}^9$ and $\Gamma_y \in \{0,1\}^{32}$. Then an approximation to the S-box is of the form

$$x \cdot \Gamma_x = S(x) \cdot \Gamma_y.$$

Since there are too many possible linear approximations, we can only consider a small subset of them. Below are three types that we considered.

[3] Independently, Knudsen and Raddum [5] have found a linear approximation to the S-box with bias 82/512.

1. Suppose Γ_x has Hamming weight zero or one. For each such Γ_x and all $2^{32} - 1$ non-zero Γ_y, we computed the bias. Table 1 gives the number of approximations with a bias greater than the conjectured 2^{-3} and also the maximum bias found for this type of approximation.
2. There are good analytical reasons to consider the least significant five bits of the output from the S-box in a linear approximation, so we might take $0 \leq \Gamma_x \leq (2^9 - 1)$ and $1 \leq \Gamma_y \leq (2^5 - 1)$. These approximations are further sub-divided into cases where the approximations involve only a single bit for both the input and output, and the case where only output bits are involved. Table 2 gives the distribution of the biases for all such approximations to the S-box.
3. Suppose $\Gamma_x = 0$ and Γ_y consists of a periodic pattern of bits. For example, with $\Gamma_y = 0x15151515$ the bias is $31/512 \approx 2^{-4.0}$, and with $\Gamma_y = 0x88888888$ the bias is $23/512 \approx 2^{-4.5}$. Such a periodic output mask helps increase the bias of the data-dependent rotations and seems to be useful when we consider certain variants of MARS in which the addition operations are changed to exclusive-or. However we will not pursue this here.

Γ_x	`0x000`	`0x001`	`0x002`	`0x004`	`0x008`	`0x010`	`0x020`	`0x040`	`0x080`	`0x100`
#	41	44	50	36	40	37	43	42	45	46
$\text{bias}_{\text{max}} \times 512$	68	73	72	71	70	73	71	72	73	72

Table 1. A partial search of all linear approximations to the S-box reveals numerous approximations with bias greater than the conjectured 2^{-3}. The number of such approximations (after searching over all output masks) for the given input mask is given in the second row while the maximum bias (bias_{max}) for that input mask is given in the third. Overall, the maximum bias we found was 73/512 and one of the masks that gives this bias is $\Gamma_x =$ `0x001` and $\Gamma_y =$ `0xefde00f5`.

5 Linear approximations of `E`

In this section we consider the formation of linear approximations to the `E` function in MARS (Figure 1). The designers of MARS present an analysis of the different ways in which linear approximations can be formed and they present their results in Table 7, on page 36 of [2]. The designers consider each subset of the input I and the three output strands L, M, and R and for each subset they list their

> "... estimate of the highest possible bias which can be obtained with this subset."

bias of approximation	number of all approximations	number of single-bit approximations	number of output-only approximations
$= 0$	522	1	1
$[2^{-9}, 2^{-8})$	1165	5	1
$[2^{-8}, 2^{-7})$	2177	2	4
$[2^{-7}, 2^{-6})$	3955	17	9
$[2^{-6}, 2^{-5})$	5340	10	12
$[2^{-5}, 2^{-4})$	2614	10	4
$[2^{-4}, 2^{-3})$	99	0	0
total	15872	45	31

Table 2. Distribution of the bias for linear approximations of the S-box with $0 \leq \Gamma_x \leq (2^9 - 1)$ and $1 \leq \Gamma_y \leq (2^5 - 1)$. Some of these involve only single bits at the input and output and some approximations have no input bits involved.

In this section we first give immediate theoretical justification for why some of these estimates must be in error. We then provide experimental confirmation of the fact.

5.1 Approximations involving L and M

Here we consider linear approximations involving the outputs L and M from the function E. In [2] it is conjectured that the highest possible bias is 2^{-20}.

Why this bound is immediately suspect. Consider approximations of the E function that involve only L and M. We will use notation from Figure 7 of [2]. Furthermore we will let r_1 and r_2 denote the 32-bit quantities from which the two rotation amounts are induced by the R strand. There is then the following relation between L and M.

$$\begin{aligned} M &= w_1 \lll r_1 \\ w_2 &= S(w_1) \\ w_3 &= w_2 \oplus r_1 \oplus r_2 \\ L &= w_3 \lll r_2. \end{aligned}$$

Consider the least significant five bits of r_1 and r_2 and suppose that $r_1 \bmod 32 = r_2 \bmod 32 = r$. So with probability 2^{-5} we have that

$$(L \ggg r) \equiv S(M \ggg r) \bmod 32.$$

Therefore, we can construct linear approximations for the L and M strands from many linear approximations to $S[\cdot]$ and relate the biases of the two. From Table 2 this suggests that there will be many linear approximations involving L and M with bias less than 2^{-20}. Since the estimates in [2] take account of the integer addition that is required to mix the output from the L and M strands into the data stream, a little more work and experimentation is needed. But already this bound is very suspect.

Linear hulls and key dependency. For some linear approximations the experimental bias can be larger than that predicted by analysis. In addition, there are often some complicated key dependencies when considering the susceptability of a cipher to linear (and for that matter differential) cryptanalysis. As an example, let us consider the relation between $L[0]$ and $M[0]$.

$$\begin{aligned} M[0] &= (w_1[32 - r_1]) \lll r_1 \\ w_2[i] &= S(w_1[j]) \\ w_3[i] &= w_2[i] \oplus r_1[i] \oplus r_2[i] \\ L[0] &= (w_3[32 - r_2]) \lll r_2. \end{aligned}$$

The above set of equations certainly hold for $i = j = r_1 = r_2 = 0$. But they also hold for many other proper choices of i, j, r_1, r_2 and so we see there are many valid approximations that might contribute to the bias demonstrated by the linear hull $L[0] \oplus M[0]$. Furthermore, it is possible to consider the construction of linear approximations that contribute to the linear hull $L[0] \oplus M[0]$ but which also involve local linear approximations across the key-dependent operations. While such approximations might be expected to have very low bias (particularly those involving the integer multiplication) it is sometimes surprising to note the significant variability in the bias observed.

Experimental Results
Choose 300 keys k_a and k_m at random. Check that the multiplicative key word k_m is valid for MARS. Compute the bias of $L[0] \oplus M[0]$ exactly over all 32-bit input words.

bias	*number of keys*
$(2^{-12.4}, 2^{-13.0})$	276
$(2^{-13.0}, 2^{-13.8})$	24

For all the keys tried, the bound for the largest bias to approximations involving L and M in the MARS analysis paper [2] has been contradicted. Sometimes this is by as much as a factor of 2^7. While we only give results on experiments involving $L[0]$ and $M[0]$ similar results hold for approximations involving other bit positions.

5.2 Approximations involving M and R

In Table 7 of [2] it is conjectured that the bias of any approximation for the M, R strands is at most 2^{-7}. Below we will identify two linear approximations of M, R which both have an *average* bias of 2^{-7}. But for half the subkeys the bias is larger than 2^{-7} and for a fraction of 1/8 of the subkey the bias is at least $2^{-6.2}$.

The M and R strands of the `E` function satisfy the following equations.

$$R = ((I \lll 13) \times k_m) \lll 10,$$
$$M = (I + k_a) \lll r_1,$$

where k_a, k_m are the addition and multiplication subkeys, and r_1 is the rotation amount determined by $R \ggg 5$.

There are three perfect linear approximations across the multiplication (since the multiplicative keys in MARS have a special form) and two of them are useful for our purpose. These two approximations correspond to the following perfect approximations for the I and R strands; $I[19] \oplus R[10]$, and $I[20] \oplus R[10, 11]$.

Now let us consider two matching approximations for the I and M strands. Note that we choose $M[0]$ since it will can be used in an approximation to the addition in the update step without losing bias.

$$I[j] \oplus M[0], \; j = 19, 20. \tag{1}$$

In [2] it is claimed that any approximation (except for the least significant bit) for $y = I + k_a$ has a bias of at most 1/4. However, for fixed subkeys, the bias can be larger. In particular, the bias depends on the probability that there is a carry into bit j during addition for a random input I.

Let q represent the least significant j bits of k_a. Then the bias of Approximation 1 across addition is given by $\frac{q-2^{j-1}}{2^j}$. Suppose that $q \leq 2^{j-4}$. Then the bias across addition is at least $\frac{7}{16} = 2^{-1.2}$. Similarly, we have the same result when $k_a(j) \geq 2^j - 2^{j-4}$. Therefore, for a fraction of 1/8 of the addition subkeys[4], the bias across addition is at least $2^{-1.2}$.

Taking into account the bias of the data-dependent rotation operation, we note that the two approximations $M[0] \oplus R[10]$ and $M[0] \oplus R[10, 11]$ have a bias of at least $2^{-6.2}$ for a fraction of 1/8 of the addition subkeys.

Experimental Results

Choose 300 keys k_a and k_m at random. Check that the multiplicative key word k_m is valid for MARS. Compute the bias of $M[0] \oplus R[10]$ exactly over all 32-bit input words.

bias	*number of keys*
$(2^{-6.0}, 2^{-6.2})$	34
$(2^{-6.2}, 2^{-7.0})$	107
$(2^{-7.0}, 0)$	159

[4] Similarly, for a fraction of 1/4 of the addition subkeys, the bias across the integer addition is at least $2^{-1.4}$.

5.3 Approximations involving L, M, and R

The conjectured highest possible bias in [2] for approximations of this type is 2^{-13}. However it is straightforward to observe that any linear approximation for $\mathtt{E}$ of the type $L[\Gamma_L] \oplus M[\Gamma_M]$ can be converted into a linear approximation of the type $L[\Gamma_L] \oplus M[\Gamma_M] \oplus R[\Gamma_R]$ (where $1 \leq \Gamma_R \leq (2^5 - 1)$) with a bias that should be either the same or higher. From the results in Section 5.1 we might immediately suspect the bound given in Table 7 of [2]. Experiments confirm that this bound is regularly contradicted.

Experimental Results
Choose 300 keys k_a and k_m at random. Check that the multiplicative key word k_m is valid for MARS. Compute the bias of $L[0] \oplus M[0] \oplus R[0]$ exactly over all 32-bit input words.

bias	*number of keys*
$(2^{-11.9}, 2^{-12.0})$	11
$(2^{-12.0}, 2^{-12.6})$	289

5.4 Approximations involving L only

Here we consider the bias of the linear hull $L[0]$. The bound for the largest bias [2] is 2^{-15}. While for the majority of cases this seemed to be a reasonable estimate, we found that in 2% of the cases this bound was contradicted.

Experimental Results
Choose 500 keys k_a and k_m at random. Check that the multiplicative key word k_m is valid for MARS. Compute the bias of $L[0]$ exactly over all 32-bit input words.

bias	*number of keys*
$(2^{-13.5}, 2^{-15.0})$	10
$(2^{-15.0}, 2^{-16.0})$	24
$(2^{-16.0}, 0)$	466

5.5 Approximations involving I, L, and M

The conjectured bound for the largest bias [2] is 2^{-13}. Once again, this is not an unreasonable estimate, but we found that in 5% of the cases this bound was contradicted.

Experimental Results
Choose 300 keys k_a and k_m at random. Check that the multiplicative key word k_m is valid for MARS. Compute the bias of $I[0] \oplus L[0] \oplus M[0]$ exactly over all 32-bit input words.

bias	*number of keys*
$(2^{-12.5}, 2^{-13.0})$	15
$(2^{-13.0}, 2^{-15.0})$	253
$(2^{-15.0}, 0)$	32

5.6 Summary of intermediate results

So far we have considered some components of the analysis that led the MARS designers to their conjectured resistance of MARS to linear cryptanalysis. While some of the assumptions and estimates seem to be reasonable, some are clearly in error. Much of this is due to the complexity of MARS and the difficulty in providing accurate estimates as to how some of the internal operations interact.

To be fair, the components that are most in error are not directly used in the estimate to the global resistance of the cryptographic core to linear cryptanalysis. However we will now demonstrate that there are also errors in this later part of the analysis.

6 Global Approximations to MARS

The designers of MARS [2] provide justification for the resistance of MARS to linear cryptanalysis using a graphical approach. The function `E` is used within a network. A very elegant argument is given [2] which provides a lower bound on the different types of approximations that would be present in a linear approximation to the core. Then, since Table 7 in [2] gives estimates for the highest possible bias for these different types of approximations, a conservative bound to the bias of the best linear approximation to the cryptographic core of MARS is derived.

While this approach appears to be reasonable, it appears that its implementation within [2] is incorrect and leads to an estimate for the bias of a linear cryptanalytic attack that cannot be substantiated by the analysis. This graphical approach also ignores the potential effects of linear hulls.

6.1 The super-round

Consider the forward transformation in the cryptographic core. It consists of eight rounds of computation involving the `E` function. This is typically viewed as two super-rounds with each super-round consisting of four of these individual rounds (see Figure 2). The following line of reasoning is given in [2].

1. Analyze the graph-network of a super-round.

2. Count the minimum number of times an approximation needs to involve L and M. (Analysis in [2] reveals 1 and 2 respectively.)
3. From Table 7 in [2] observe the maximum bias for any approximation involving L and M. (From [2] this gives 2^{-8} and 2^{-6} respectively.)
4. Estimate the bias for a super-round to be bounded by $2^{-8} \times 2^{-6} \times 2^{-6} \times 2^{2} = 2^{-18}$, and therefore that the bias for the keyed transformation is at most $(2^{-18})^4 \times 2^3 = 2^{-69}$.

First, we observe that the "conservative bound" of 2^{-6} for the bias of approximations involving L given in Table 7 is derived under the assumptions that the maximum bias of an approximation across a data-dependent rotation is 2^{-6} and that the maximum bias of an approximation across the S-box is 2^{-3}. However, in Section 4 we showed that this latter assumption is incorrect. In fairness, we have to say that the S-box approximations which contradict the conjectured bias for the S-box might not be of immediate use. Nevertheless it provides additional evidence to suggest that this bound for L might not be conservative.

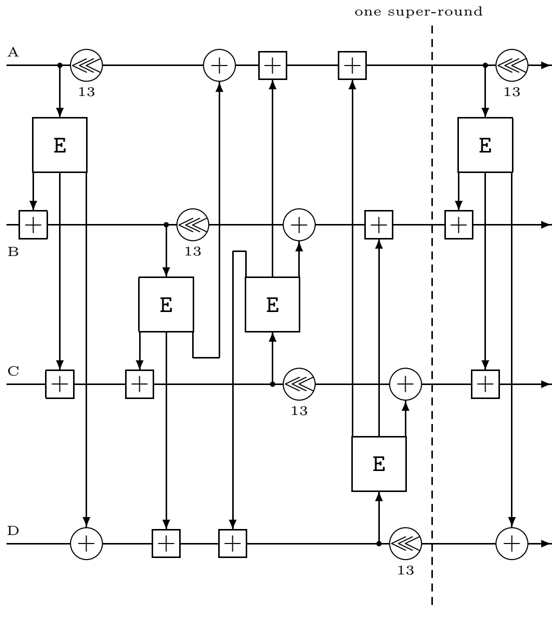

Fig. 2. The forward mode of the keyed transformation in MARS. There are eight forward rounds, consisting of two so-called super-rounds [2]. The details of the `E` function are given in Figure 1. Note that the three outputs from `E` can be read in a natural fashion as L, M, and R.

Second, we observe something far more important. Namely, there exist linear approximations to the super-round that involve only one approximation involving an L strand and one approximation involving an M strand. This contradicts the second point in the reasoning provided and is most easily demonstrated by example. Consider Figure 2. We will denote the input to successive `E` functions by I_j and the outputs by L_j, M_j, and R_j where $0 \leq j \leq 4$. Recall that the structure of MARS is such that I_j is taken from a strand of data that is subsequently combined with the output R_{j+1}. Similar relationships for the other inputs and outputs can be derived. The values of the four data strands at the beginning of the $(j+1)^{\text{st}}$ call to the function `E` will be denoted by A_j, B_j, C_j, and D_j.

The following four approximations to successive `E` functions can be chained together:

$$\begin{gathered} I_0[x] \oplus L_0[y] \\ - \\ R_2[y+13] \oplus I_2[z] \\ M_3[y+13] \oplus R_3[z+13] \end{gathered}$$

Here x, y, and z can be any bit patterns, not necessarily single-bit approximations. When concatenated they lead to the following linear approximation to a super-round

$$A_0[x] \oplus B_0[y] \oplus B_4[y+13] \oplus C_4[z+13].$$

What is important here is that both L and M are each involved only once. Thus, if we were to follow the arguments given in [2] then we would only be able to conclude the following:

1. Every approximation across a super-round must involve at least one approximation involving L and one approximation involving M.
2. All linear approximations across a super-round are likely to have a bias less than $2^{-6} \times 2^{-8} \times 2 = 2^{-13}$. All linear approximations across the cryptographic core are likely to have a bias less than $(2^{-13})^4 \times 2^3 = 2^{-49}$.

With a bias of 2^{-49} the keyed transformation, or cryptographic core, could be compromised with 2^{98} known plaintexts, far less than the notional aim of 2^{128} known plaintexts. Whether or not there is an approximation that achieves this kind of bias is unknown.

6.2 Linear hulls

The graphical approach to considering the bounds of linear approximations to a super-round (or more) neglects to take account of the effect of linear hulls. By considering linear hulls we aim to take account of other internal linear approximations that might otherwise be overlooked. Again, we illustrate our point with an example.

One very clear example of the possibility of a substantial linear hull effect occurs when we consider five successive `E` functions. Once again, see Figure 2.

Using the same notation as we have just introduced the following five linear approximations can be concatenated together.

$$\begin{aligned} &M_0[x] \oplus I_0[y] \\ &R_1[y+13] \oplus L_1[x] \\ &M_2[y+13] \\ &L_3[y+13] \oplus R_3[x+13] \\ &I_4[y+13] \oplus M_4[x+13] \end{aligned}$$

Together these approximations can be combined to give a five-round linear approximation $B_0[x] \oplus B_5[x]$. This is valid for any bit pattern x though of course a single-bit value to x is likely to be the most useful. Note that this linear approximation iterates. The internal bit pattern y can take on any value at all. The bias of the linear hull is likely to be dominated by the action of single-bit values to y but this is not necessarily the case.

Unfortunately it is very hard to get experimental confirmation of the likely extent of this effect. Unlike some other AES finalists, MARS does not lend itself it to the construction of small-scale versions on which to experiment. However it seems to be reasonable to observe that by ignoring the implications of one particular effect, the bounds that are derived for the resistance of MARS to linear cryptanalysis could be over-estimates. The complexity of the cipher hinders us from deciding how important this might be.

7 Conclusions

Our investigation of MARS has been very focused on one particular attribute of the cipher. The very limited time available in the second round of the AES process has meant that it has been very difficult to make any headway in providing an adequately accurate analysis of the cipher. However, our experiences are sufficient to draw the following conclusions.

- The complicated design of MARS seems to force the use of potentially inaccurate models in deriving estimates for the security offered.
- The non-existence of small-versions of the cipher means that analysis and experimentation are severely hampered.
- It is unclear how significant the issues of linear hulls, key-dependency, and multiple linear approximations might be in the linear cryptanalysis of MARS. Both the `E` function and the global design of the network are somewhat novel.

As the AES process draws to a close, we are forced to look at the security of the ciphers with only very partial evidence available. In such circumstances it is important to feel that the evidence at hand is representative of the true behavior of the cipher. However we have shown that the analysis provided in [2] could be flawed in important ways.

After our experience with MARS we are forced to conclude that the complexity of a cryptographic algorithm can be an enormous handicap to accurate analysis.

References

1. E. Biham and A. Shamir. *Differential Cryptanalysis of the Data Encryption Standard*. Springer-Verlag, New York, 1993.
2. C. Burwick, D. Coppersmith, E. D'Avignon, R. Gennaro, S. Halevi, C. Jutla, S. Matyas, L. O'Conner, M. Peyravian, D. Safford, and N. Zunic. MARS - a candidate cipher for AES. June 10, 1998.
3. B.S. Kaliski and M.J.B. Robshaw. Linear cryptanalysis using multiple approximations. In Y.G. Desmedt, editor, *Advances in Cryptology — Crypto '94*, volume 839 of *Lecture Notes in Computer Science*, pages 26–39, New York, 1994. Springer Verlag.
4. B.S. Kaliski and M.J.B. Robshaw. Linear cryptanalysis using multiple approximations and FEAL. In B. Preneel, editor, *Fast Software Encryption*, volume 1008 of *Lecture Notes in Computer Science*, pages 249–264, 1995. Springer Verlag.
5. L.R. Knudsen and H. Raddum. Linear approximations to the MARS S-box. March 15, 2000. Preprint.
6. M. Matsui. Linear cryptanalysis method for DES cipher. In T. Helleseth, editor, *Advances in Cryptology — Eurocrypt '93*, volume 765 of *Lecture Notes in Computer Science*, pages 386–397, 1994. Springer-Verlag.
7. M. Matsui. The first experimental cryptanalysis of the Data Encryption Standard. In Y. G. Desmedt, editor, *Advances in Cryptology — Crypto '94*, volume 839 of *Lecture Notes in Computer Science*, pages 1–11, New York, 1994. Springer-Verlag.
8. K. Nyberg. Linear approximation of block ciphers. In A.D. Santis, editor, *Advances in Cryptology — Eurocrypt '94*, volume 950 of *Lecture Notes in Computer Science*, pages 439–444, 1994. Springer-Verlag.
9. B. Schneier, J. Kelsey, D. Whiting, D. Wagner, C. Hall, and N. Ferguson. Twofish: A 128-bit Block Cipher. 15 June, 1998.
10. A. A. Selcuk. New results in linear cryptanalysis of RC5. In S. Vaudenay, editor, *Fast Software Encryption*, volume 1372 of *Lecture Notes in Computer Science*, pages 1–16, 1998, Springer-Verlag.

This article was processed using the LaTeX macro package with LLNCS style