From: zunic@us.ibm.com
X-Lotus-FromDomain: IBMUS
To: aesround2@nist.gov
Date: Mon, 15 May 2000 13:36:24 -0400
Subject: Final Comments

Attached are the final comments for Round 2 from the MARS team.
Nev Zunic

Internet: zunic@us.ibm.com
IBM Crypto Solutions
(914) 435-6949 (T/L 295)

---------------------- Forwarded by Nev Zunic/Poughkeepsie/IBM on
05/15/2000 01:37 PM ---------------------------

Nev Zunic
05/15/2000 01:30 PM

To: jfoti@nist.gov
cc: David Safford/Watson/IBM@IBMUS, shaih@watson.ibm.com@IBMUS
From: Nev Zunic/Poughkeepsie/IBM@ibmus
Subject: Final Comments

Jim,
Attached are our final comments for Round 2. I've also attached two additional documents (one on key agility and the other on linear analysis) which are referenced in the Final Comments. These are complementary documents. I'm attaching three different (doc, pdf, and postscript) filetypes of the Final Comments:

(See attached file: Final Comments.doc)(See attached file: Final Comments.pdf)(See attached file: Final Comments.ps)(See attached file: linear.ps)(See attached file: key-agil.ps)

If you have any questions, please let me know.
Nev

Internet: zunic@us.ibm.com
IBM Crypto Solutions
(914) 435-6949 (T/L 295)

Comments on MARS's linear analysis

The IBM MARS team

May 12, 2000

Abstract

We refine the bounds on the linear analysis of the MARS core, to reflect some comments that were made recently by Robshaw and Yin. They claimed that the original argument could only be used to show a bound of 2^{-49} on the bias, rather than the claimed 2^{-69}. Part of this criticism may be attributed to a poor wording of the argument in our original paper, leading to their misinterpretion of our intent. Other parts represent "slightly more sophisticated" approximations, which were not covered by the original bound. Nonetheless, even with these comments, we can still show a bound of 2^{-61} on the bias of any approximation to the core. Roughly, this means that any such approximation must use all the available plaintext-ciphertext pairs in the codebook.

1 Bounds on the linear approximations of the MARS core

In the MARS submission documents [1], an analysis of the graph structure of the MARS core was presented, which suggested that the core does not have a linear approximation with bias of more than 2^{-69} (under some assumptions on the way this approximation is constructed). Recently, Robshaw and Yin challenged this analysis [3], and claimed that the correct bound should actually be about 2^{-49} rather than 2^{-69}. This discrepancy was claimed to be the result of an erroneous argument in the original graph-based analysis, as well as better analysis of the possible approximations for the various components of the MARS core.

The purpose of this note is to clarify the situation with respect to these bounds. We first note that the claimed "erroneous argument" in the graph-based analysis is merely a poor wording of the argument in the original analysis. In this note we present a more detailed description, which implies that under the assumptions of the original analysis, the graph-based argument indeed yield a bound of at most 2^{-69}. Moreover, even if we modify the analysis to incorporate the better approximation of the components suggested in [3], the graph-based analysis still yields a bound of less than 2^{-61} (which correspond to data complexity of 2^{122}). roughly, this means that any linear approximation of the core that follows the framework of this analysis, must use essentially all the available plaintext-ciphertext pairs in the codebook.

Health warning. Before going any further, we would like to remind the reader what should and shouldn't be read into the results of such analysis. On one hand, obtaining a bound of 2^{-61} *does not mean* that there exists an approximation with this bias. This analysis only yield a bound on the possible bias, not an actual approximation with this bias. On the other hand, it also does not mean that there is no way to devise approximations with better bias, since the analysis relies on some assumptions, which may turn out to be unjustified. The value of this analysis, therefore, is mainly in drawing attention to the limitations of the "obvious approach" for linear analysis of the

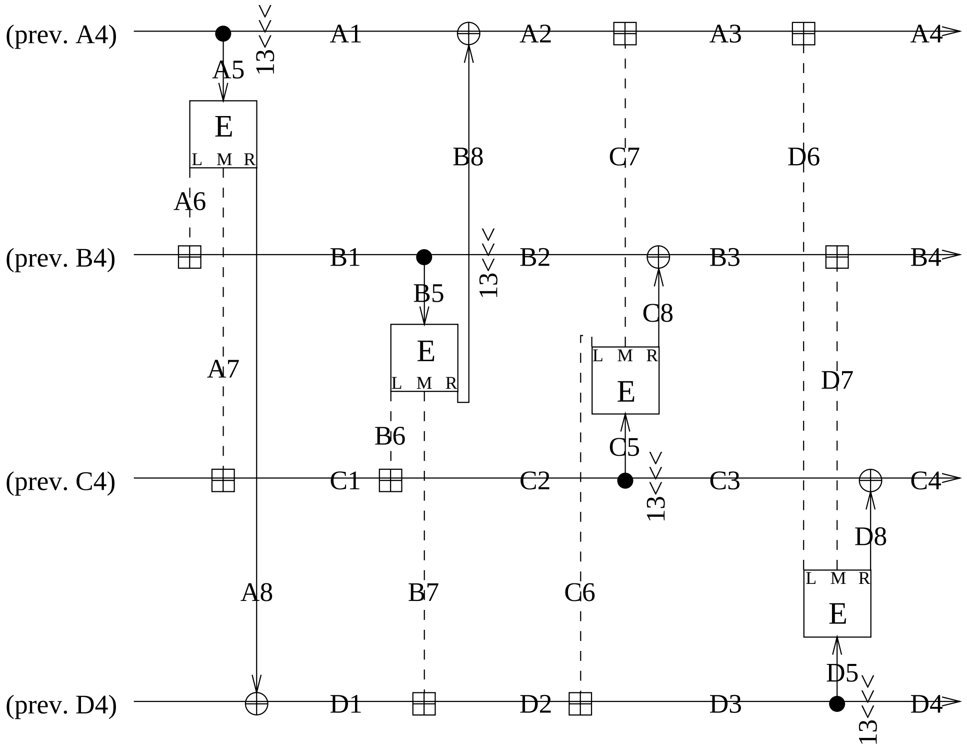

Figure 1: A "super-round" of the MARS core - forward.

MARS core. In essence, these bounds mean that any effort towards linear analysis must include some "non obvious" observations about the MARS core.

For the rest of this note, it is assumed that the reader has read both the linear analysis in the MARS submission document [1, Section 4.1] as well as Robshaw and Yin's comments [3].

1.1 The framework for the graph-based analysis

The graph-based analysis consists of devising estimates for the "best possible approximation" of the various components of MARS, and then analyzing the graph structure of the core to see how these approximation can be combined. The graph structure of the MARS core is depicted in Figures 1 (forward mode) and 6 (backwards mode). In these figures, we use notations similar to those in [1, Figure 7]. For the purpose of analysis, it is convenient to refer to four consecutive E-functions as one "super-round". Using this terminology, Figure 1 describes one super-round in forward mode, and Figure 6 describes one super-round in backwards mode. The MARS core consists of two forward super-rounds, followed by two backwards super-rounds.

One can identify a linear approximation of the core with a subset of the edges in this graph structure (namely, those edges whose values are approximated). An approximation with non-zero bias must correspond to a set of edges S with the following properties (a *valid set* in the language of [1]):

1. S contains at least one input edge and at least one output edge.
2. If S contains an edge which is incident to an xor or an addition operation, then it also contains the other two edges incident to this operation.

3. If S contains an edge which is incident to a copy operation (the •'s in the figures), then it contains at least one of the other two edges incident to this operation.

4. If S contains either input edge I or output edge R of an E-function, then it contains at least one other edge incident to this E-function.

Throughout this note, the edges that correspond to a data-dependent rotation operations are denoted by dashed arrows in all the figures, while all the other edges are denoted by solid arrows.

In [1, Section 4.1.3], it was claimed concerning a valid subset S, that "... in this case S must contain at least three rotation edges in each super-round" (assertion 6). Robshaw and Yin correctly pointed out that there exists a valid set S for a single super-round with only two rotation edges. Indeed, the original write-up was not clear here; we had intended our claim only in the amortized sense: three edges per super-round on average, or twelve edges during the four super-rounds of the complete core. It turns out that this analysis is simple enough to be carried out by hand. In this note we therefore demonstrate that analysis (which is very tedious, but not hard).

Another objections of Robshaw and Yin to the analysis in [1] is that the bounds on the components which were used in that analysis are too low. In particular, for the $\{L, M\}$-approximation of the E-function, they show that a much higher bias is possible than what was assumed in the original analysis. (We note that in some sense, this is outside the model that was considered in the original analysis, since it was explicitly stated there that the bounds apply only to oblivious application of the piling-up lemma. Nonetheless, it is clear that Robshaw and Yin's estimates are "the right ones to use".) In the analysis that we describe in this note, we show that even when using these higher estimates for the bias of the different components, the graph-based argument still yield a bound of less than 2^{-61}.

1.2 Miscellaneous comments

Before presenting the graph-based analysis, let us clarify a few issues regarding the model in which this analysis is applicable.

No "global cancellations". It should be clear from the outset that such graph-based analysis can only cover approximations which are devised by combining many local approximations using the piling-up lemma. Namely, it does not (and cannot) take into account "linear hulls" (or "global cancellations", in the language of [1]).

Indeed, the whole point in presenting lower bounds as in this analysis is to eliminate this simple approach towards analysis of the cipher, and to point out that if an approximation is possible, it would have to use more sophisticated tools. We note, however, that some features of the MARS core makes such high-bias "global cancellations" less likely. (Specifically, the fixed rotation by 13 on the lines, combined with the way different bit positions interact differently with the E-function.)

Treatment of keys. In the original analysis in [1], and also in this note, we use an estimation of the bias "for a random key". This helps in making the analysis more uniform. A different approach was taken by Robshaw and Yin, where they considered "the largest bias that can be obtained with a noticeable fraction of the keys". The latter approach may be useful in identifying "classes of weak keys" with respect to linear analysis, but it does not reflect very well the actual resilience of the core.

We also note that the numerical differences between these approaches are not very large. Taking the approach of Robshaw and Yin would only make a factor of two difference in the

$\{M, R\}$ approximation of the E-function, which would have raised our bound from 2^{-61} to about 2^{-57}.

For the analysis below, we use the following estimates for the bias of the different approximations of the E-function (see Table 1). These estimates represent a modification of Table 7 from [1], to accommodate the results of Robshaw and Yin.

Approximation	Estimated bias
$\{L\}$	2^{-15}
$\{M\}$	2^{-20}
$\{L, M\}$ $\{L, M, R\}$	2^{-12}
$\{L, R\}$ $\{I, L\}$ $\{I, L, R\}$	2^{-8}
$\{M, R\}$	2^{-7}
$\{I, L, M\}$ $\{I, L, M, R\}$	2^{-13}
$\{I, M\}$ $\{I, M, R\}$	2^{-6}
$\{I, R\}$	1/2

Table 1: Bias of approximations for the E-function

1.3 Organization

The analysis below is organized as follows. First we analyze the possible valid sets for a singe forward super-round, and then use it to analyze the possible valid sets for two forward super-rounds together. Similarly, we analyze the possible valid sets for a single backwards super-round, and then use it to analyze the possible valid sets for two backwards super-rounds together. Finally, we combine all these results to analyze the valid sets for the entire core, and then present our conclusions.

2 Valid sets for a single forward super-round

Here we exhibit all the valid sets for a single super-round that only use two rotation edges. To make this exposition tractable, we partition it into several separate cases, depending on which is the "first rotation edge" that is used in this valid set. In this exposition, we use the notations from Figure 1.

First rotation edge is $A6$. We must have also $B1$, and therefore at least one of $B5$ or $B2$ in the valid set. We consider both cases:

- If we have $B2$, then we must include also $C8, B3, D7$ and $B4$. Including $C8$ implies that at least one of $C5, C6$ or $C7$ must also be taken. We cannot use $C6$ or $C7$, since they are rotation edges, which will bring the total number of rotation edges to more than two.

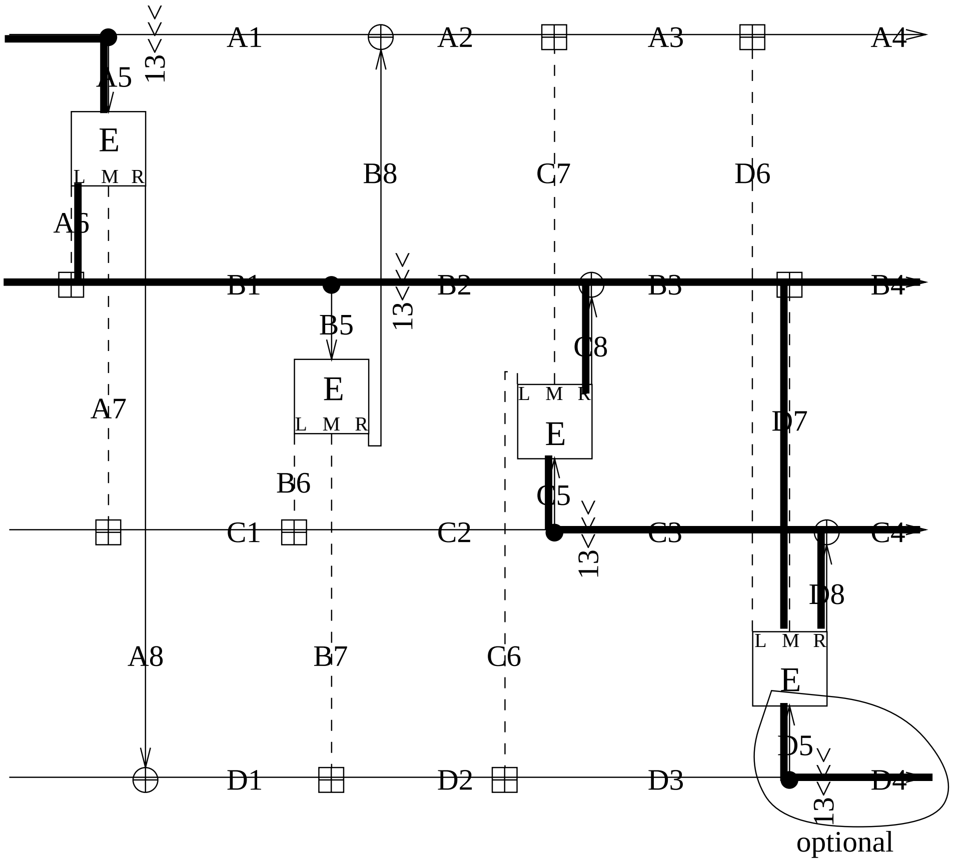

Figure 2: A “high bias” valid set for a forward super-round.

So we must take $C5$, and therefore also either $C2$ or $C3$ (or both). Taking $C2$ implies also $B6$ and $A7$, which brings the number of rotation edges in S above two. Hence, we must use $C3$, which implies also $D8$.

Hence, we have the entire B line ($B0 - B4$), the second half of the C line ($C3 - C4$), rotation edges $A6$ and $D7$, and multiplication edges $C8$ and $D8$. To avoid using an $\{L\}$-approximation of the first E-function, we must take another edge adjacent to that E-function, and the only way to complete the valid set without additional rotation edges is to take $A5$ and $A0$. In addition, we may or may not include edges $D5, D4$. The result is depicted in Figure 2. Notice that this is essentially the example exhibited by Robshaw and Yin, which has bias of at most 2^{-14}.

- If we have $B5$, then we must have there either $B6$, $B7$ or $B8$. Including either $B6$ or $B7$ implies also the inclusion of either $A7$ or $C6$, respectively, which brings the number of rotation edges to more than two. Including $B8$ implies that we must include also $C7$ and $D6$, which again brings the number of rotation edges to more than two. Hence, in this case there is no way to get a valid set with less than three rotation edges.

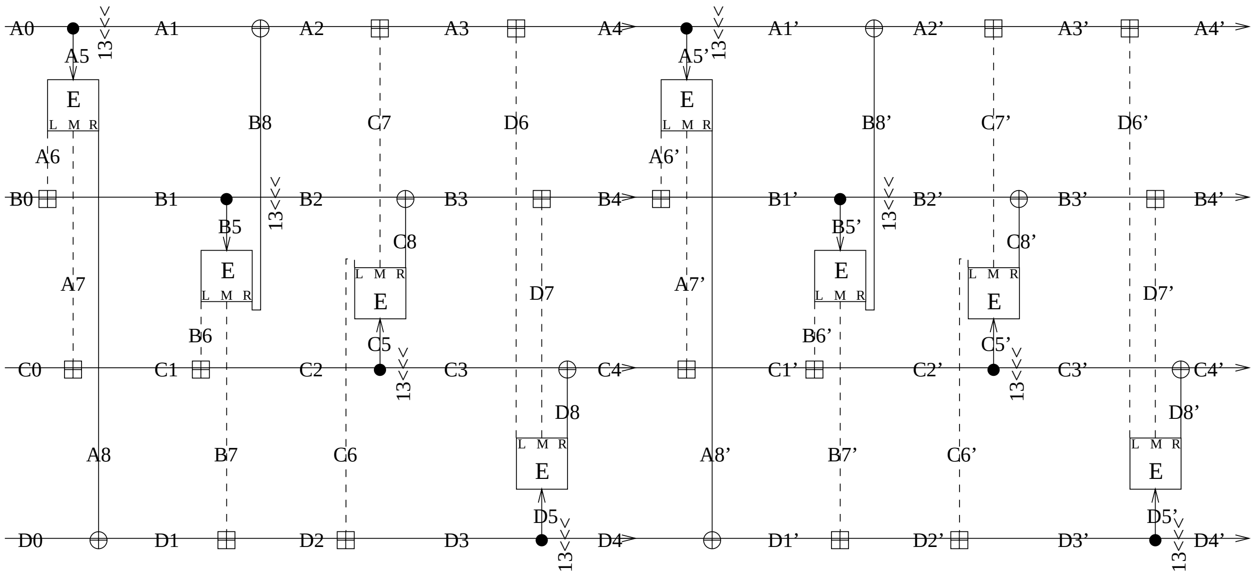

Figure 3: Two successive "super-rounds" of the MARS core.

First rotation edge is $A7$. In this case we must also have $B6$ and at least one of $C3, C5$. If there are no other edges adjacent to the second E-function, then we have an $\{L\}$-approximation of that E-function, which has bias at most than 2^{-15}. So we are forced to include another edge from that E-function. $B7$ is a rotation edge, and inclusion of $B8$ implies also rotation edges $C7$ and $D6$, so these options yield too many rotation edges. Inclusion of $B5$ implies that we must have either $B1$ (and therefore $A6$) or $B2$ (and therefore $C8$ and $D7$). In either case, we end up with at least three rotation edges.

First rotation edge is $B7$. (Notice that the first edge cannot be $B6$, since any valid set that includes $B6$ must also include $A7$.) In this case we must also have rotation edge $C6$. Again to avoid using the $\{M\}$-approximation of the second E-function we must pick at least one more edge adjacent to it, either $B5, B6$ or $B8$. $B6$ is excluded since it is a rotation edge. Using $B8$ implies also edges $C7$ and $D6$, which are rotation edges. Using $B5$ implies that we must have wither $B1$ (and also $A6$) or $B2$ (and also $C8$ and $D7$). In any case, we end up with at least three rotation edges.

First rotation edge is $C7$. (Again, it cannot be $C6$, since any valid set that includes $C6$ must also include $B7$.) In this case we have also $B8$ and $D6$. Having $B8$ means that we must also have either $B5, B6$ or $B7$. Edges $B6$ and $B7$ are excluded because they are rotation edges. If $B5$ is included, then we must take either $B1$ or $B2$. Including $B1$ implies also including $A6$, and including $B2$ implies also including $C8$ and $D7$. In any case, we have more than two rotation edges.

No other options. It is easy to see that we cannot have one of the D edges be the "first rotation edge"(since we must have at least one input edge, $A0, B0, C0$ or $D0$).

Conclusions. We conclude that the only valid subset for a single super round which potentially can have bias of more than 2^{-18} (what you get from three rotation edges) is the one in Figure 2, which is essentially Robshaw and Yin's example.

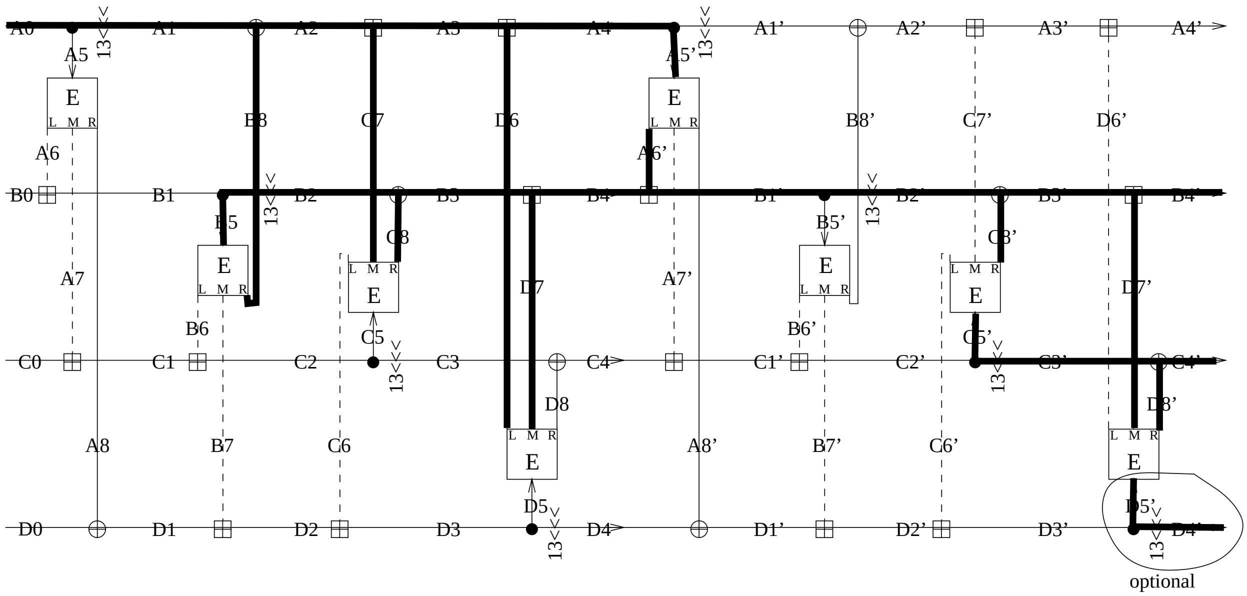

Figure 4: A "high bias" set for two forward super-rounds.

3 Valid sets for two forward super-rounds

We now proceed to analyze two consecutive forward super-rounds. To get bias of more than 2^{-35}, we must use the valid set from Figure 2 in either the first or second super round. For the case analysis below, recall that the valid set from Figure 2 includes the edges $A0, B0$ and $B4, C4$ (and potentially also $D4$).

Using the "high bias case" in the first super round. Since we have $B4 = B0', C4 = C0'$ in the valid set for the first super round, we must have also $A6', A7'$ and $B6'$ in the second super-round (all of which are rotation edges). Adding any edge adjacent to either the first, second or third E-function in the second super-round would imply adding at least one more rotation edge to the valid set, making it at least six rotation edges for the two super-rounds. (We note that with the original estimates from the MARS submission document [1, Table 7], we would have forced to pick another edge adjacent to the first E-function, so as to avoid using the $\{L, M\}$-approximation. Using Robshaw and Yin's estimates, however, we can afford to use that approximation.)

Since $C5'$ is excluded, we must include $C3'$ and $C4'$, and therefore also $D8'$. Now, the only way to complete the valid set without including more rotation edges is to include $D5'$ and $D4'$. The result is shown in Figure 4, and it has bias at most 2^{-32}.

Using the "high bias case" in the second super round. In this case we have $A0' = A4$ and $B0' = B4$ in the valid set for the second super round. Hence, we must have also edges $D7, D6, C8, C7$ and $B8$ from the first super-round. This already gives three rotation edges, and it is easy to see that adding any other edge which is adjacent to either the first, third or fourth E-functions would necessarily add at least one more rotation edge.[1]

Since we have $B2$ in the valid set, we must add either $B1$ or $B5$. Adding $B1$ implies also the rotation edge $A6$, so we add $B5$. The result is shown in Figure 5, It has bias of at most 2^{-31}.

[1] Note again that if we couldn't use the $\{L, M\}$-approximation, we would have forced to add some rotation edges.

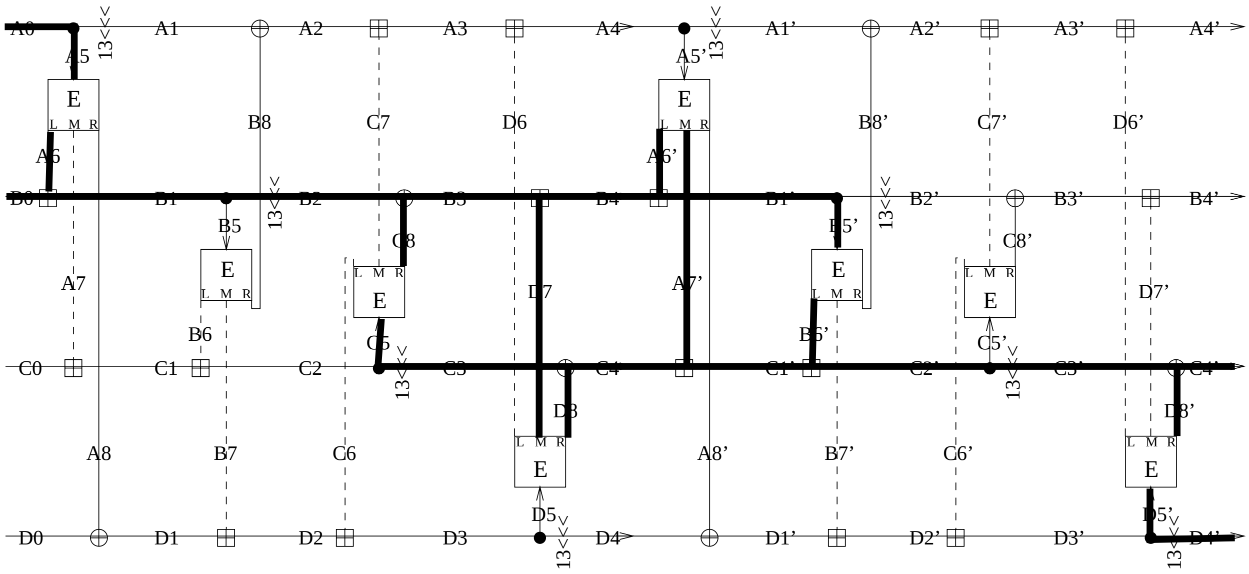

Figure 5: Another "high bias" set for two forward super-rounds.

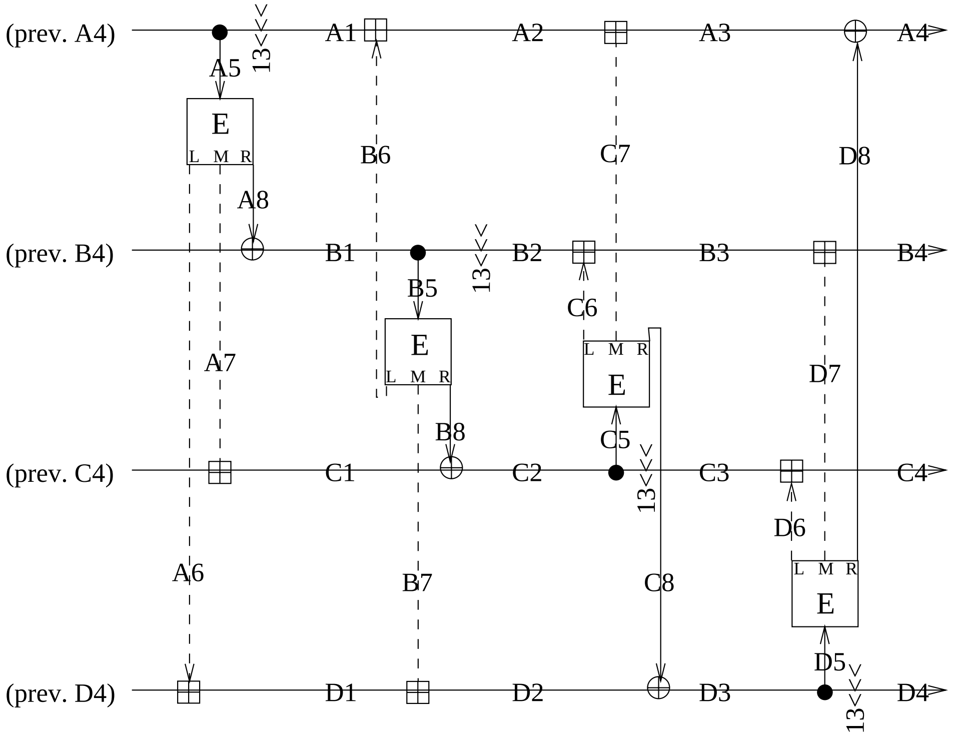

Figure 6: A "super-round" of the MARS core - backwards.

4 Valid sets for a single backwards super-round

We again partition the potential valid sets to several separate cases, depending on which is the "first rotation edge" which is used in this valid set. In this exposition, we use the notations from Figure 6.

First rotation edge is $A6$. We must have also $B7$ and $C8$. Hence we must also have either $C6, C7$ or $C5$. The first two will bring the number of rotation edges above two, while the last implies that we also have $D6$, which again yields more than two rotation edges.

First rotation edge is $A7$. We must have also $C0, C1, B8$, and therefore, also at least one of $B7, B6$ or $B5$. Taking $B6$ implies also $C7$, and taking $B7$ implies also $A6$, both of these have more than two rotation edges. So we must take $B5$. Then we need at least one of $B1, B2$. Taking $B2$ implies $C6$ and $D7$, leading to more than two rotations. Taking $B1$ implies also $A8$, and optionally $A5$ and $A0$.

We also need at least one of $C3, C5$. Taking $C5$ would imply at least one of $C6$ (and therefore also $D7$), $C7$ (and therefore also $B6$) or $C8$ (and therefore also $B7, A6$. Each of these yield more than two rotation edges.

So we must take $C3$, which means that we also have $D6$ and $C4$. To avoid using an $\{L\}$-approximation to the last E-function, we must take at least one more of $D5, D7, D8$, and the only combination that does not imply an additional rotation edge is to take $D5$ and $D4$. The resulting valid set is depicted in Figure 7. Not surprisingly, this is essentially a mirror image of the valid set from Figure 2.

First rotation edge is $B6$. We must have also $C7$ and $D8$. To avoid using an $\{M\}$-approximation of the third E-function, we must include at least one of $C6$ (and therefore also $D7$), $C8$ (and therefore also $A6, B7$), or $C5$ (and therefore at least one of $A7, D6$). In each case, we have more than two rotation edges.

First rotation edge is $C6$. (Notice that the first edge cannot be $B7$, since it implies $A6$.) If we have $C6$, we must have also $D7$. To avoid using an $\{L\}$-approximation of the third E-function, we must include at least one of $C7$ (and therefore also $B6$), $C8$ (and therefore also $A6, B7$), or $C5$ (and therefore at least one of $A7, D6$). In each case, we have more than two rotation edges.

No other options. It is easy to see that we cannot have one of the D edges be the "first rotation edge" (since we must have at least one input edge, $A0, B0, C0$ or $D0$).

5 Valid sets for two backwards super-rounds

We now proceed to analyze two consecutive backwards super-rounds. If we do not use the "high bias case", then we can get bias of at most 2^{-35} for the two rounds. Below we therefore analyze the cases where we use this case in either the first or the second backwards super-round. For the case analysis below, recall that the "high bias case" include the edges $B0, C0$ (and potentially also $A0$) and $C4, D4$.

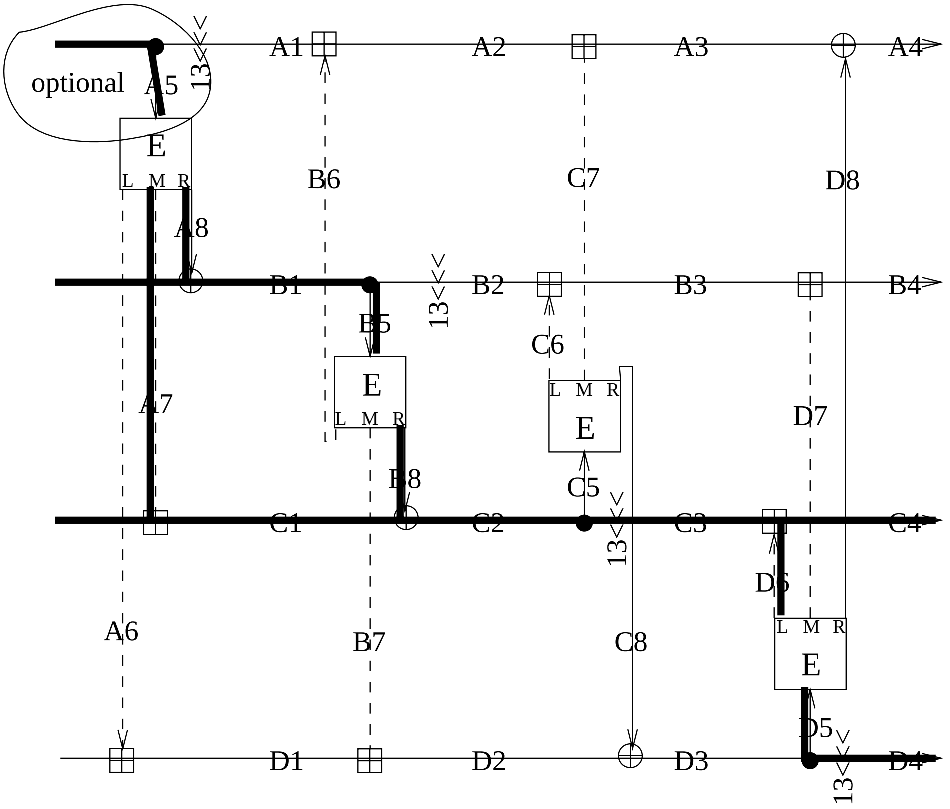

Figure 7: A "high bias" valid set for a backwards super-round.

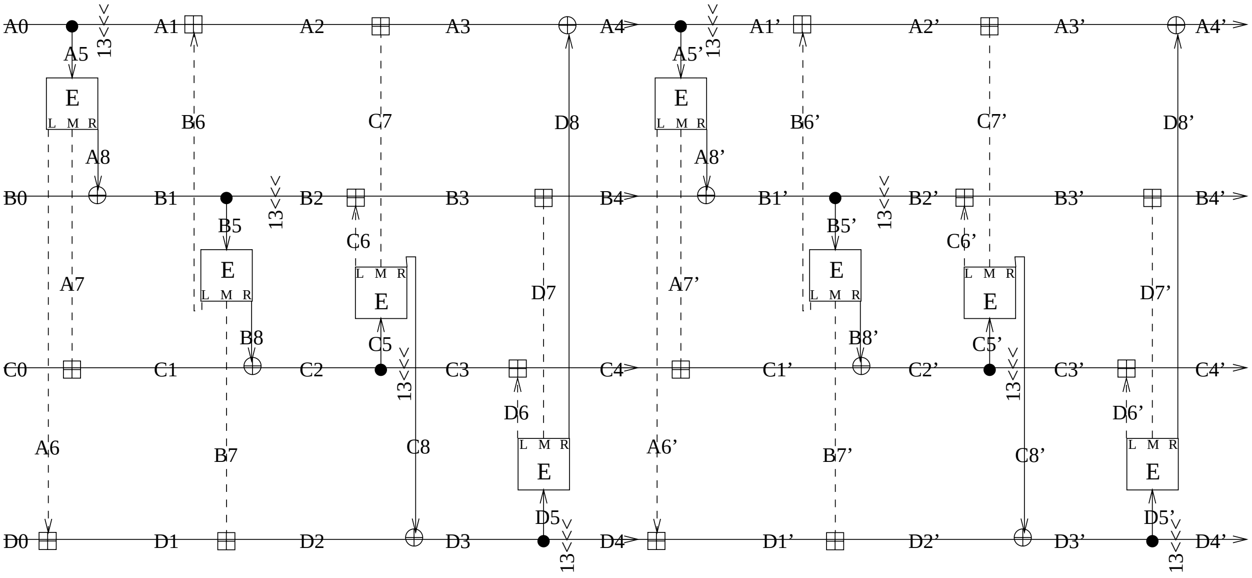

Figure 8: Two successive "super-rounds" of the MARS core - backwards.

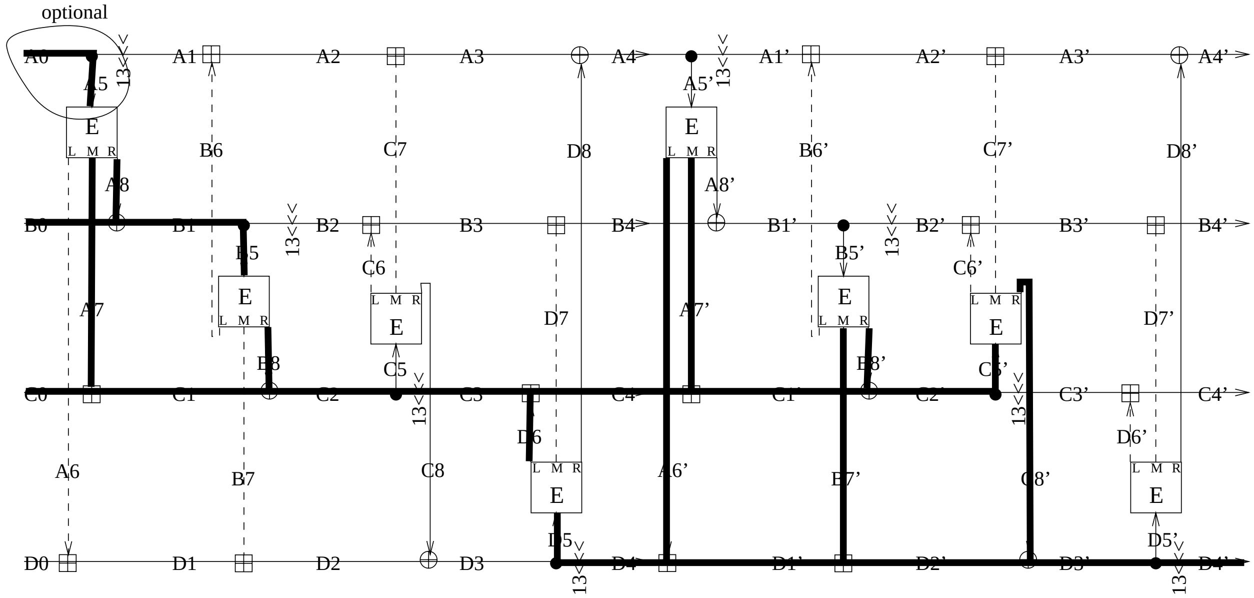

Figure 9: A "high bias" set for two backwards super-rounds.

Using the "high bias case" in the first super round. Since we have $C4 = C0', D4 = D0'$ in the valid set for the first super-round, then we must also have edges $A6'$, $A7'$, $B7'$, $B8'$ and $D8'$. This already includes three rotation edges for the second super-round, and it is easy to see that selecting any other edge adjacent to either the first, second or fourth E-functions in this super-round implies at least one more rotation edge in the set.[2]

To complete the valid set, we must include either $C5'$ or $C3'$, and either $D5'$ or $D4'$. Choosing either $C3'$ or $D5'$ yields additional rotation edges, so the only way to get less than six rotation edges for the two backwards super-rounds is to choose $C5'$ and $D4'$. The result is demonstrated in Figure 9, and it has bias of at most 2^{-31}.

Using the "high bias case" in the second super round. Since we have $B4, C4$ in the valid set for the second super-round, then we must also have edges $D7$, $D6$ and $C6$. This already includes three rotation edges for the first super-round, and it is easy to see that selecting any other edge adjacent to either the fourht, third, or second E-functions in this super-round implies at least one more rotation edge in the set.[3]

Since we have $B2$ in the valid set, we must also include wither $B1$ or $B5$. As we said above, including $B5$ yields too many rotation edges, so we must include $B1$, and therefore also $A8$. Now, the only way to complete the valid set without including additional rotation edges is to include also $A5$ and $A0$. The result is demonstrated in Figure 10, and it has bias of at most 2^{-32}.

6 Valid sets for the entire core

The analysis so far reveals that for each of the two halves of the core, there are only two valid sets which may potentially correspond to approximations with bias of more than 2^{-35}, and even these sets have potential bias of less than 2^{-31}. It can also be seen that the valid sets in Figures 5

[2] Note again that if we couldn't use the $\{L, M\}$-approximation, we would have forced to add some rotation edges.

[3] Note again that if we couldn't use the $\{L, M\}$-approximation, we would have forced to add some rotation edges.

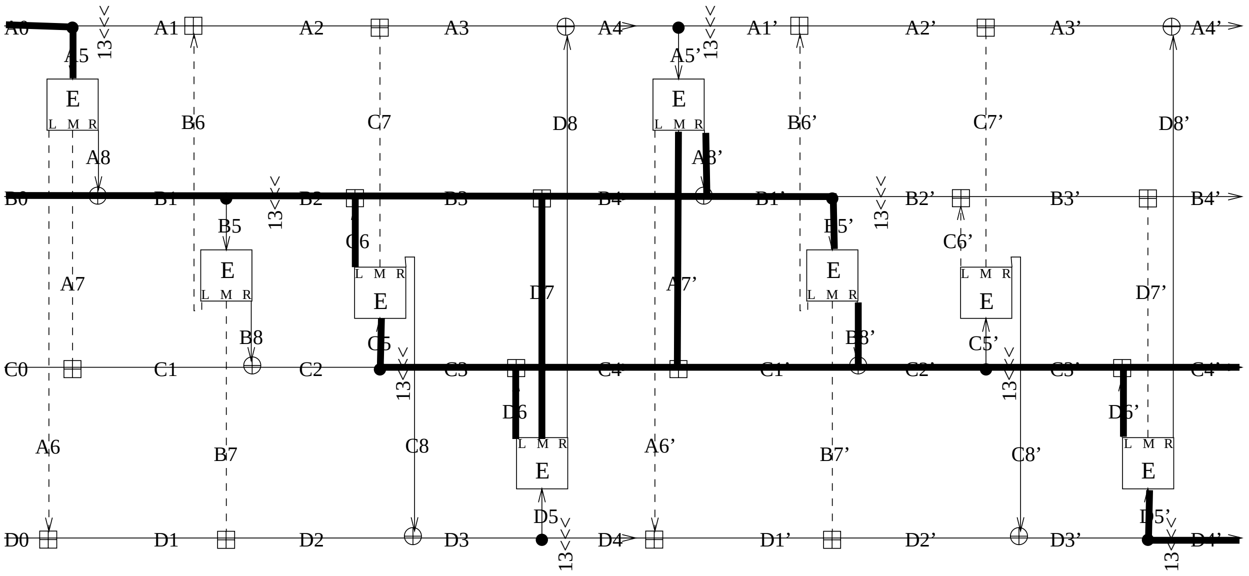

Figure 10: Another "high bias" set for two backwards super-rounds.

and 9 are compatible, so their combination indeed yields a valid set for the core. This valid set has potential bias less than 2^{-61}, and this is the best possible (under the model assumptions of this analysis). If such approximation existed, it would have corresponded to a linear approximation with data complexity of more than 2^{122} (of the core alone, without the mixing). It should be stressed again that this is only an upper bound, and it is unlikely that one can exhibit a linear approximation along these lines with such a "high bias". The value in exhibiting such bound is in demonstrating that any useful linear approximation must use more clever ideas that just the "piling up lemma".

(We remark that all the "high bias" sets contain several $\{L, M\}$ approximations of the E-function. Hence, if we were using the estimates from [1, Table 7] then these sets would correspond to very low bias approximations.)

7 Conclusions

Our conclusions from this analysis is that it validates the original assessment of the resilience of MARS against linear cryptanalysis. Specifically, we have demonstrated that the "obvious ways" of constructing linear approximations are not sufficient against the MARS core. We were able to accomplish this since the MARS core is quite modular, hence, it is possible to devise an analysis "from the bottom up". Clearly, we cannot prove that no good approximations exist, since such a statement is beyond the state of the art. The best proofs that can be offered are ones which cover the obvious linear approximation constructs and demonstrate MARS' resilience to this cryptanalysis. We believe that we have accomplished this goal.

Before concluding, we would like to respond to one more point in the note of Robshaw and Yin. Throughout their note, they repeatedly claim that the complexity of MARS makes it hard to gauge its true strength. We strongly disagree with this assessment. On one hand, the MARS core is not complex. In fact, it was specifically designed in a modular way to facilitate analysis. This is manifested in the bounds which we were able to establish, as well as in the body of work by others.

Even more important, though, we believe that their "thesis", that seemingly simpler ciphers are easier to analyze, is not well supported in reality. To stress this point, we note that even for a seemingly very simple cipher like RC5, a linear analysis was presented (and even published [2]), and three years later was found to be faulted [4]. Such a phenomenon was not encountered for the "more complex" DES. We also note that trying to establish bounds for the seemingly simpler RC6, turns out to be at least as hard (if not harder) than for MARS.

References

[1] C. Burwick, D. Coppersmith, E. D'Avignon, R. Gennaro, S. Halevi, C. Jutla, S.M. Matyas Jr., L. O'Connor, M. Peyravian, D. Safford, and N. Zunic, "MARS - a candidate cipher for AES". In proceedings of the 1st AES conference. NIST, 1998.

[2] B.S. Kalisky and Y.L. Yin. "On differential and linear cryptanalysis of the RC5 encryption algorithm", Crypto'95, LNCS vol. 963, pp. 171-184. Springer-Verlag. 1995.

[3] M. Robshaw and Y.L. Yin. "Potential flaws in the conjectured resistance of MARS to linear cryptanalysis". manuscript. (Presented at the rump session in the 3rd AES conference.)

[4] A.A. Selcuk. "New results in linear cryptanalysis of RC5", FSE'98, LNCS vol. 1372, pp. 1-16. Springer-Verlag. 1998.