Branching Heuristics in Differential Collision Search with Applications to SHA-512

Maria Eichlseder(✉), Florian Mendel, and Martin Schläffer

IAIK, Graz University of Technology, Graz, Austria
maria.eichlseder@iaik.tugraz.at

Abstract. In this work, we present practical semi-free-start collisions for SHA-512 on up to 38 (out of 80) steps with complexity $2^{40.5}$. The best previously published result was on 24 steps. The attack is based on extending local collisions as proposed by Mendel et al. in their Eurocrypt 2013 attack on SHA-256. However, for SHA-512, the search space is too large for direct application of these techniques. We achieve our result by improving the branching heuristic of the guess-and-determine approach to find differential characteristics and conforming message pairs. Experiments show that for smaller problems like 27 steps of SHA-512, the heuristic can also speed up the collision search by a factor of 2^{20}.

Keywords: Hash functions · Cryptanalysis · SHA-512 · Collision attack · Guess-and-determine attack · Branching heuristic

1 Introduction

Since 2005, many collision attacks have been shown for commonly used and standardized hash functions. In particular, the collision attacks of Wang et al. [41,42] on MD5 and SHA-1 have convinced many cryptographers that these widely deployed hash functions can no longer be considered secure. As a consequence, NIST has proposed the transition from SHA-1 to the SHA-2 family. Many companies and organization follow this advice and have already migrated to SHA-2. Even more might do so, since Keccak [33] has not been standardized as SHA-3 yet and SHA-2 is faster on several platforms. In particular, SHA-512 is much faster than both SHA-256 and Keccak on most 64-bit platforms [2]. For this reason, it has been suggested to use a truncated version of SHA-512 even for 256-bit hash values [38]. NIST also defines this variant, called SHA-512/256, in FIPS 180-4 [32].

Nevertheless, not many cryptanalytic results on SHA-512 have been published in the last few years. The security of SHA-512 against preimage attacks was first studied by Aoki et al. in [1]. They presented a preimage attack on 46 out of 80 steps. This was later extended to 50 steps by Khovratovich et al. in [19]. Recently, Li et al. showed that particular preimage attacks can also be used to construct a free-start collision attack for up to 57 steps of SHA-512 in [24]. However, all attacks are only slightly faster than the respective generic attack complexities.

C. Cid and C. Rechberger (Eds.): FSE 2014, LNCS 8540, pp. 473–488, 2015.
DOI: 10.1007/978-3-662-46706-0_24

The currently best known practical collision attack on both the SHA-512 hash and compression function is for 24 steps. It has been published independently by Indesteege et al. [16] and by Sanadhya and Sarkar [36]. Both attacks are trivial extensions of the attack strategy of Nikolić and Biryukov [34] which applies to both SHA-256 and SHA-512. Recently, Mendel et al. [27,29] demonstrated how to extend these attacks to get collisions for the SHA-256 compression function on up to 38 steps with practical complexity.

The attacks by Mendel et al. use a guess-and-determine based automatic search tool to find differential characteristics and conforming message pairs for reduced SHA-256. Since the first publication by De Cannière and Rechberger on SHA-1 in [7], such tools have been constantly improved [21,22,27,29]. Nevertheless, the increased search space of SHA-512 (due to larger word sizes) prevented successful attacks without the application of new ideas.

To handle the larger search space of SHA-512, we propose a new branching heuristic for the guess-and-determine strategy used in these attacks. Our approach is inspired from related ideas in SAT solvers [15,23]. The heuristic performs a randomized look-ahead selection of candidates which should be guessed first. Using this approach, we can detect contradictions earlier and reduce the search space faster. More specifically, we are able to speed up the search on SHA-512 by a factor of about 2^{20} (for 27 steps), which allows us to construct practical collisions for 38 steps with a complexity of $2^{40.5}$.

The remainder of this paper is structured as follows. We first give a high-level overview of our attack strategy and related work in Sect. 2. In Sect. 3, we discuss branching heuristics used in SAT solvers and propose our new look-ahead branching heuristic for differential cryptanalysis tools. In Sect. 4, we demonstrate the application of the heuristic to SHA-512 and present a practical semi-free-start collision for 38 steps. Finally, we conclude in Sect. 5.

2 Motivation

In this section, we give a brief overview on the differential cryptanalysis of hash functions and how the guess-and-determine approach is used to search for differential characteristics. Furthermore, we provide a high-level view on optimization options to improve this search.

2.1 The Search for Differential Characteristics

A differential attack consists of two main parts: constructing a differential characteristic and finding a confirming message pair. Since the attacks by Wang et al. [40–42], these parts are further divided to improve the overall attack complexity as follows:

- **Find a differential characteristic**
  1. Construct the high-probability part of a characteristic.
  2. Determine the low-probability part of a characteristic.

– **Find a conforming message pair**
  3. Use message modification in low-probability part.
  4. Perform random trials in high-probability part.

We provide significant improvements in finding dense low-probability differential characteristics. To motivate our work, we first provide an overview of previously published methods and show how we improve upon these methods using improvements in the guess-and-determine approach.

Constructing the differential characteristic for the low-probability part is one of the most difficult tasks in a differential attack. The main reason is that such low-probability characteristics are usually very dense and have many (hidden) relations which need to be taken into account. Wang et al. found the dense low-probability characteristics for the attacks on MD4, MD5, RIPEMD, SHA-0 and SHA-1 mostly by hand [40–42]. However, for more complex hash functions, such an approach is infeasible. Therefore, (semi-)automatic approaches have been published soon afterwards [7,37]. These approaches have then been refined in a number of publications. Recently, more sophisticated approaches have been proposed that enable attacks on more complex hash functions such as SHA-256 [27,29] among many others [20,22,26,28]. All these approaches (including the search by hand) follow the guess-and-determine strategy.

2.2 The Guess-and-Determine Approach

The basic idea of the search algorithm is to pick and guess previously unrestricted bits. After each guess, the information gained from these restrictions is propagated to other bits. If an inconsistency occurs, the algorithm backtracks to an earlier state of the search and tries to correct it. Similar to [27], we denote these three parts of the search by decision (guessing), deduction (propagation), and backtracking (correction). Then, the search algorithm proceeds as in Algorithm 1 given below.

Algorithm 1. Guess-and-Determine Search Algorithm

Let U be a set of undetermined bits
while U contains undetermined bits **do**
  Decision (Guessing)
    1. pick an undetermined bit (randomly or heuristically)
    2. impose new constraints on this bit
  Deduction (Propagation)
    3. propagate the new information to other variables and equations
    4. **if** an inconsistency is detected, start backtracking,
       else continue with step 1
  Backtracking (Correction)
    5. try a different choice for the decision bit and continue with step 3.
    6. **if** all choices result in inconsistencies,
         undo guesses until this critical bit can be resolved

This procedure can also be visualized by a search tree, which is traversed by depth first search. The branching strategy decides on which variable to split the tree next and thus defines the tree's shape. Typically, the complete tree is much too large for complete traversal, so it is crucial that more promising branches are visited first. In addition, the backtracking algorithm can skip parts of the tree in favor of exploring more distant parts. This makes the search incomplete, but in practice greatly improves the performance.

The challenge in finding a long differential characteristic lies in the fine-tuning of the search algorithm. There are many possible variations, and details can determine whether the search succeeds or fails.

2.3 Improving the Guess-and-Determine Approach

Basically, a guess-and-determine is just a repetition of two steps: first, guess the value of some unknowns and second, determine the value of as many unknowns as possible. However, in practice more details need to be considered to mount successful guess-and-determine attacks on complex hash functions. The most important points to consider are given as follows:

1. **Problem Description:** The complexity of a guess-and-determine attack can be significantly improved if we first optimize the problem description. For example, first constructing a characteristic and then searching for a message pair is already such an optimization. Additionally, the choice of intermediate variables and a good starting point are crucial for a guess-and-determine attack to succeed.
2. **Guessing Strategy:** Instead of randomly guessing variables, using high-level information can lead to much better guesses. For example, by preferring bits (or even words) with no differences, characteristics tend to get sparser, have a higher probability, and conforming message pairs are more likely to exist.
3. **Branching Rules:** In every iteration, the guess-and-determine algorithm needs to decide which branch of the search tree to follow. Using a good branching heuristic, contradictions can be found faster and the search space can be reduced more quickly.
4. **Propagation:** Every time a variable is guessed, we need to check whether the guess is invalid, or new information on other variables can be determined. There is a trade-off between the effort we spend in this step and simply guessing more bits. Different propagation methods for ARX-based hash functions are covered in detail in [9,21,22,27].
5. **Backtracking:** To recover from bad search spaces which do not contain many solutions anymore, we need to backtrack. Two extreme options are performing a complete restart or examining the complete search space. A successful backtracking strategy for SHA-2 has been published in [27].

The first two points are very specific to a given problem and cannot be solved in general. In our attacks on SHA-512, a good starting point is constructed using improved local collisions, similar as in the attack on SHA-256 in [29]. The last two

points have already been covered in a number of publications. Additional efforts in these points did not improve the guess-and-determine attack on SHA-512. This leaves the branching rules which have not been optimized yet. In the following, we show that a good branching heuristic can significantly improve the efficiency of a guess-and-determine attack.

3 Branching Heuristics

Branching rules are one of the essential ingredients for guess-and-determine attacks. They define how the search algorithm selects the next variable to guess, and which guess values to try first for this variable. The branching rule aims to keep the search runtime as short as possible. Depending on whether the current partial assignment is correct (satisfiable) or contradictory, this means either that a satisfying solution is found as soon as possible, or that the contradiction is detected quickly. In the latter case, this corresponds to identifying a conflicting subset of unassigned variables and branching on these first in order to prune the search tree. The search trees traversed by different branching rules can vary drastically in size, from constant (for unsatisfiable problems) or linear (for satisfiable problems) to exponential in the number of variables [35].

This section first discusses existing branching rules used in general-purpose SAT solvers and for the cryptanalysis of hash functions. Afterwards, we introduce our randomized look-ahead heuristic.

3.1 Branching Heuristics in SAT Solvers

Most general-purpose SAT solvers are based on the Davis-Putnam-Logemann-Loveland (DPLL) algorithm [6], a guess-and-determine approach for satisfiability problems given in conjunctive normal form (CNF). The problem of choosing optimal branching variables and corresponding assignments for DPLL algorithms has been proven to be both NP-hard and coNP-hard [25]. However, there is a variety of commonly implemented branching rules based on different heuristics to evaluate the urgency or relevance of potential branching variables. In addition, meta-rules to select different branching rules depending on the situation and search history have been proposed [13].

Commonly used SAT branching rules can be categorized according to their target heuristic (current properties, look-ahead or history analysis), their output (a single branching variable/literal or a preselection of candidate variables) and their randomness (deterministic or randomized). Popularly used heuristics include the following:

– **Uniformly random**. A random unassigned variable is picked with uniform distribution. This approach is computationally cheapest. Many modern SAT solvers apply this rule with a small probability and otherwise use a more informed choice. In differential cryptanalysis, this is the most typical rule.

- **Small clauses**. The earliest heuristics greedily favor variables that appear in many small clauses. The rationale for this choice is twofold. First, smaller clauses need to be fulfilled "more urgently" since there are fewer options left that avoid contradictions. Second, even if the guessed literal evaluates to false in binary clauses, unit propagation ensues and curtails the search tree.

  Example heuristics of this category include Böhm's rule [3], MOM (maximum occurrences on clauses of minimum size) [10], and the Jeroslaw-Wang rules [17]. The latter, for example, assign weights $w(c)$ to clauses c that decrease exponentially with the clause length $|c|$. Each literal (OS-JW) or variable (TW-JW) scores according to the weight sum of all clauses it appears in, and the best literal or variable is picked for guessing.

  More recently, small clauses have been used as a preselection heuristic for more expensive look-ahead rules. In differential guess-and-determine attacks, two-bit conditions [27] play a related role. This preselection heuristic also favors variables with a higher number of closely coupled undetermined variables.
- **Literal count**. These heuristics ignore the clause size and simply count unresolved clauses linked to a variable. Examples include DLCS and DLIS as introduced by the GRASP solver [39]. It makes sense in CNF problems, where satisfying one literal resolves the complete clause. This does not apply, for example, for the xor-chains typically found in hash functions. Instead, this heuristic would create a large amount of (hidden) dependencies and reduce the remaining freedom without the positive effect of immediate propagation.
- **Conflict driven**. A more popular variation of literal counting is VSIDS, first implemented in Chaff [30] and later included in MiniSAT [8] and others. Here, the initial literal score of each variable decays over time via multiplication with a constant $\delta < 1$. However, scores are refreshed (bumped) with occurrences in newly learned clauses from the CDCL process. Effectively, the score keeps track of recent contradictions involving the variable. Critical variables with many recent contradictions are guessed first.

  The BerkMin solver extends this concept to bump not only variables from learned clauses, but from any clauses involved in the resolution process [12].

  In differential attacks, the backtracking strategy [27] provides a similar behaviour.
- **Look-ahead**. Instead of judging current properties of the formula or the previous search history, look-ahead heuristics analyse the actual effects of branching in a candidate variable [15, 23]. For example, the Satz solver performs Unit Propagation Look-Ahead: both possible assigments for each free variable are tested for consequences of this decision and the caused unit propagations. If one of two assignments causes a contradiction, the other is fixed; if both are contradictory, backtracking is started; and if both seem valid, the variable v is assigned score

$$\mathcal{M}(v) = w(\neg v) \cdot w(v) \cdot 1024 + w(\neg v) + w(v),$$

where $w(\ell)$ is typically the number of new binary clauses caused by the propagation of literal $\ell \in \{v, \neg v\}$.
- **Locality**. To limit the candidates for expensive look-ahead calculations, the candidate variables can be limited to those occurring in recently changed (reduced) clauses, as implemented in the marchdl solver [14].

Not all of these rules are suitable for general Boolean satisfiability problems that are not given in CNF format, as already indicated in the list above. In particular, if the propagation and learning process differs from the standard SAT case, the above rules can be counterproductive. On the positive side, dedicated solvers for specific applications can apply domain-specific knowledge to guide the search process.

3.2 The Look-Ahead Branching Heuristic

The branching strategy is one of the most promising areas for optimization in differential cryptanalysis tools based on tree search. Ideally, the branching strategy quickly navigates towards a valid assignment of variables and avoids subtrees without solutions. For detecting invalid subtrees, the branching strategy relies on the propagation method to detect contradictions as soon as possible. However, the propagation procedure can not only be used to decide whether previous guesses were contradictory. In addition, we also want to apply it to guide the branching strategy. The goal of this interaction is to minimize the size of the search tree in order to find solutions faster.

The basic principles of our implementation of the look-ahead branching heuristic are given by the following two observations:

- **Productive propagation is good.** Guessing a variable where propagation of the value determines (many) other variables can have multiple advantages compared to variables with less propagation. The most immediate effect is that the remaining search space is reduced. If more variables are determined right now, they will not create unnecessary subtrees for guessing later. The overall tree size and thus the complexity of the remaining search is reduced.
- **Contradictions are even better.** Of course, the overall search aims to find non-contradictory assignments. Nevertheless, discovering contradictory value assignments in the current subtree is consistently helpful for the remaining search. If only one of two possible value assignments is contradictory, the variable certainly needs to be fixed to the other value. If both values are contradictory, we must already have made an error with a previous guess and need to backtrack immediately. In both cases, it is clearly better to address the conflicting bit sooner rather than later.

Note that the first criterion is not beyond controversy. In particular, limiting the search space at the same time reduces the remaining degrees of freedom.

If one value assigned to a specific bit propagates better than the second possible value, then, intuitively speaking, the probability for a solution in the remaining search space for the first option is lower than for the second value.

3.3 Implementation of the Look-Ahead Branching Heuristic

In order to implement the criteria above in a practical branching heuristic, we use a look-ahead approach related to the Unit-Propagation Look-Ahead (UPLA) used in some SAT solvers. When the branching rule needs to select the next variable to guess, each candidate is in turn evaluated.

In more detail, for each candidate, a value is tentatively assigned and the propagation method is applied to determine the consequences of this assignment. If a contradiction occurs, this candidate is selected immediately. Otherwise, the number of propagated variables is calculated. If it is better than the previously favorite candidate, this variable becomes the new favorite.

There are two performance-related problems with this basic approach. First, performing the look-ahead propagation for all free variables is very costly. Second, the basic UPLA approach includes no randomization. However, we need randomization since a complete search of the tree is typically computationally infeasible in differential cryptanalysis. Instead, large tree parts are skipped and the search is restarted regularly. To avoid becoming lost in the same search branches over and over again, it is essential that the branching strategy is sufficiently randomized.

We address both problems at once by selecting only a random subset of variables for closer evaluation. Our branching heuristic is summarized in Algorithm 2.

Algorithm 2. Look-ahead branching heuristic for differential cryptanalysis

Let U be a set of undetermined bits and $s_{\max}$ the limit of look-ahead candidates.
repeat
  Guessing
    1. pick a bit $v \in U$ randomly and increment s
    2. impose new constraints on this bit v
  Propagation
    3. propagate the new information to other variables and equations
    4. **if** an inconsistency is detected, **return** v as the decision bit
       else count the number m of additional variables that were assigned due to this guess and save the pair (v, m) in a list L.
  Update
    5. remove all variables that were assigned due to the guess v from the set U
    6. undo all changes to restore the original assignment
until U is empty or $s \geq s_{\max}$
return v^* from L with the highest score m as the decision bit

The size of the randomly selected subset is an essential parameter for the success of the heuristic. To limit the look-ahead costs, we limit the maximum

subset size by a constant number that is chosen in the beginning of the search procedure, depending on the specific problem instance. In order to also provide sufficient randomization, we additionally bound the size relative to the current number of unguessed variables.

Beside the subset size, the decision which individual variables to select for look-ahead plays a role. UPLA-based solvers use a pre-selection of interesting candidates, for example by locality criteria. In our case, the search performance can be greatly improved by only guessing bits of specific hash function words and favoring bits with more two-bit conditions or bits involved in recent conflicts. However, the selection must remain sufficiently randomized.

Additionally, we do not explicitly evaluate variables that were already determined by the propagation procedure of one of the previous candidates. We mark these as evaluated without calculating a separate look-ahead and without considering them as favorite candidates, since their score is at most as good as the bit that triggered their propagation (at least with respect to one of the assignment options).

4 Application to SHA-512

In this section, we discuss the application of our look-ahead branching heuristics to SHA-512. As a result, we are able to construct the first practical collision on the reduced SHA-512 compression function for 38 out of 80 steps. The best previously published result was on 24 steps.

4.1 Brief Description of SHA-512

SHA-512 is an iterated hash function that processes 1024-bit input message blocks and produces a 512-bit hash value. In the following, we briefly describe the hash function. It basically consists of two parts: the message expansion and the state update transformation. A detailed description of the hash function is given in [31].

Message Expansion. The message expansion of SHA-512 splits the 1024-bit message block into 16 64-bit words M_i and expands them into 80 expanded message words W_i as follows:

$$W_i = \begin{cases} M_i & 0 \leq i < 16 \\ \sigma_1(W_{i-2}) + W_{i-7} + \sigma_0(W_{i-15}) + W_{i-16} & 16 \leq i < 80 \end{cases}$$

The functions $\sigma_0(x)$ and $\sigma_1(x)$ are given by

$$\sigma_0(x) = (x \ggg 1) \oplus (x \ggg 8) \oplus (x \gg 7)$$
$$\sigma_1(x) = (x \ggg 19) \oplus (x \ggg 61) \oplus (x \gg 6).$$

State Update Transformation. The state update transformation starts from a (fixed) initial value IV of 8 64-bit words $A_{-4}, \ldots, A_{-1}, E_{-4}, \ldots, E_{-1}$ and updates them in 80 steps. In each step one expanded message word W_i is used to compute the two state variables E_i and A_i as follows:

$$E_i = A_{i-4} + E_{i-4} + \Sigma_1(E_{i-1}) + \mathrm{IF}(E_{i-1}, E_{i-2}, E_{i-3}) + K_i + W_i$$
$$A_i = E_i - A_{i-4} + \Sigma_0(A_{i-1}) + \mathrm{MAJ}(A_{i-1}, A_{i-2}, A_{i-3}).$$

For the definition of the step constants K_i we refer to [31]. The bitwise Boolean functions IF and MAJ used in each step are defined by

$$\mathrm{IF}(x, y, z) = x \wedge y \oplus x \wedge z \oplus z$$
$$\mathrm{MAJ}(x, y, z) = x \wedge y \oplus y \wedge z \oplus x \wedge z,$$

and the linear functions Σ_0 and Σ_1 are defined as follows:

$$\Sigma_0(x) = (x \ggg 28) \oplus (x \ggg 34) \oplus (x \ggg 39)$$
$$\Sigma_1(x) = (x \ggg 14) \oplus (x \ggg 18) \oplus (x \ggg 41).$$

After the last step of the state update transformation, the initial values are added to the output values of the last step (Davies-Meyer construction). The result is the final hash value or the initial value for the next message block.

4.2 Extending the Attacks on SHA-256 to SHA-512

For our collision attacks on SHA-512, we use the same strategy as in the attack on SHA-256 in [29]. Since the message expansion and state update transformation is the same (except for larger word sizes and different rotation values in Σ_i, σ_i), we can use similar local collisions (with differences in the same message words) to construct semi-free-start collisions for the compression function on up to 38 steps.

The starting point for 38 steps uses a local collision which spans 18 steps, with differences in 6 expanded message words $(W_7, W_8, W_{10}, W_{15}, W_{23}, W_{24})$. For more details on how to select the starting point, we refer to [29]. Once the starting point is fixed, the main task is to find a differential characteristic and confirming message pair for this 18-step local collision.

By using the same guessing, backtracking and propagation strategy, we did not find any results for 38 steps of SHA-512. Due to the large word size and thus, larger search space, contradictions are detected much later in SHA-512. We have tried different approaches on every level, but did not succeed in finding any valid differential characteristics. The solution was to optimize the branching strategy to detect on one hand contradictions earlier and on the other side to reduce the search space faster.

4.3 Improving the Search Using Look-Ahead Branching

To improve the search algorithm, we use the look-ahead branching heuristic proposed in Sect. 3.3. As discussed there, the choice of the subset size $s_{\max}$ is critical for the behaviour of the heuristic. We have evaluated different variants of the heuristic and get the best results for a limit of $s_{\max} = 16$. Larger values of $s_{\max}$ further reduce the tree depth, but due to the additional cost for evaluating more candidates, this does not improve the overall runtime.

Additionally, with larger subset sizes, the search tends to visit very similar subtrees again and again after each restart. This is particularly critical if the search space is limited to a few words, as in the focused search strategy described below. For other hash functions with larger states sizes or less focused search strategies, the optimal value for $s_{\max}$ may be very different.

Similar to [29], the guess-and-determine attack is separated into three stages. The rules of the guessing strategy are given in Table 1 and the three stages are summarized as follows:

Stage 1:
We first search for a consistent differential characteristic in the message expansion. Hence, we only add unconstrained bits ('?') and difference bits ('x') of W to the set U.

Stage 2:
We continue with the search for a differential characteristic in the state update. Hence, we add all unconstrained bits ('?') and difference bits ('x') of A and E to the set U. We pick decision bits more often from A, since this results in sparser characteristics for A. Experiments have shown that in this case, confirming message pairs are easier to find in the last stage.

Stage 3:
In the last stage, we search for confirming message pairs by guessing bits without difference ('-'). We only pick decision bits of A, E and W which are constrained by two-bit conditions, similar as in [27]. This serves as a preselection heuristic for the branching look-ahead.

Table 1. Decision rules in different search stages.

Stage	Decision bit	Decision rule		
		Probability	Choice 1	Choice 2
1–2	?	1	-	x
	x	1/2	u	n
		1/2	n	u
3	-	1/2	0	1
		1/2	1	0

4.4 Results

Using the improvements in the branching heuristic proposed in the previous section, we are able to find semi-free-start collisions for SHA-512 on up to 38 steps. Finding a differential characteristic together with a conforming message pair took 5441 s (≈1.5 h) on a cluster with 40 CPUs. This corresponds to a complexity of about $2^{40.5}$ evaluations of the SHA-512 compression function. The colliding message pair is given in Table 2 and the differential characteristic is shown in Table 3.

Table 2. Example of a semi-free-start collision for 38 steps of SHA-512.

h_0	e8626f53a3771964	2ae427b8c5065790	c8fd5a1628fc3337	0f362d297f82f987
	89166a0c022ffc40	c2c49c30e629239f	d1fa8bd692843025	ad4bba64c797e6ec
m	610519a88f0d2809	3addc83f01c8b179	84afa7a2772c6141	ad539854e64c9cce
	85450b73549b2085	7296b5291f31c0d9	fc978d9624e2c2cc	fffffffffffffffe
	92114cb9d2f4cd9b	34a3198b79871212	cca7f43154e38081	ac0598a589168fe1
	f32ae6a0070a8d2e	755aa5cada87e894	4b9bd7df3c94b667	65291f2b80cc8c51
m^*	610519a88f0d2809	3addc83f01c8b179	84afa7a2772c6141	ad539854e64c9cce
	85450b73549b2085	7296b5291f31c0d9	fc978d9624e2c2cc	0000000000000001
	92114cb9d2f4cd9c	34a3198b79871212	cca8143154e38079	ac0598a589168fe1
	f32ae6a0070a8d2e	755aa5cada87e894	4b9bd7df3c94b667	65291f2b80cc8c50
Δm	0000000000000000	0000000000000000	0000000000000000	0000000000000000
	0000000000000000	0000000000000000	0000000000000000	ffffffffffffffff
	0000000000000007	0000000000000000	000fe0000000000f8	0000000000000000
	0000000000000000	0000000000000000	0000000000000000	0000000000000001
h_1	946a28eedc3b2ff6	c4573d0a13ea6268	11f07b04b06900dd	897c606e4053bbe4
	2406aae9d58504b4	89b237932b061ba8	663402cb4bb1972c	d99c062dce945423

To show the benefit of our new look-ahead branching heuristic, we have performed some comparisons. Without look-ahead branching, we were able to find a semi-free-start collision for 27 steps of SHA-512 using 4 days on a cluster with 40 nodes, which corresponds to a complexity of about $2^{46.5}$. Using look-ahead branching with $s_{\max} = 16$ we can find differential characteristics with conforming message pairs within seconds on a standard PC (complexity $2^{26.5}$).

The heuristic can also be used to improve the search complexity for primitives with a smaller state to a certain extent. For example, experiments show a speedup of more than an order of magnitude for attacks on 27 or 38 steps of SHA-256. However, due to the heuristic nature of the improvement and the general sensitivity of the search procedure to different parameters, the effects are hard to quantify.

Unfortunately, we were not able to extend the semi-free-start collisions to collision attacks on the hash function. The main reason is that the resulting differential characteristics are quite dense and we do not have enough freedom to match the IV with practical complexity.

Table 3. Differential characteristic for a semi-free-start-collision of SHA-512 reduced to 38 steps (bits with two-bit conditions highlighted).

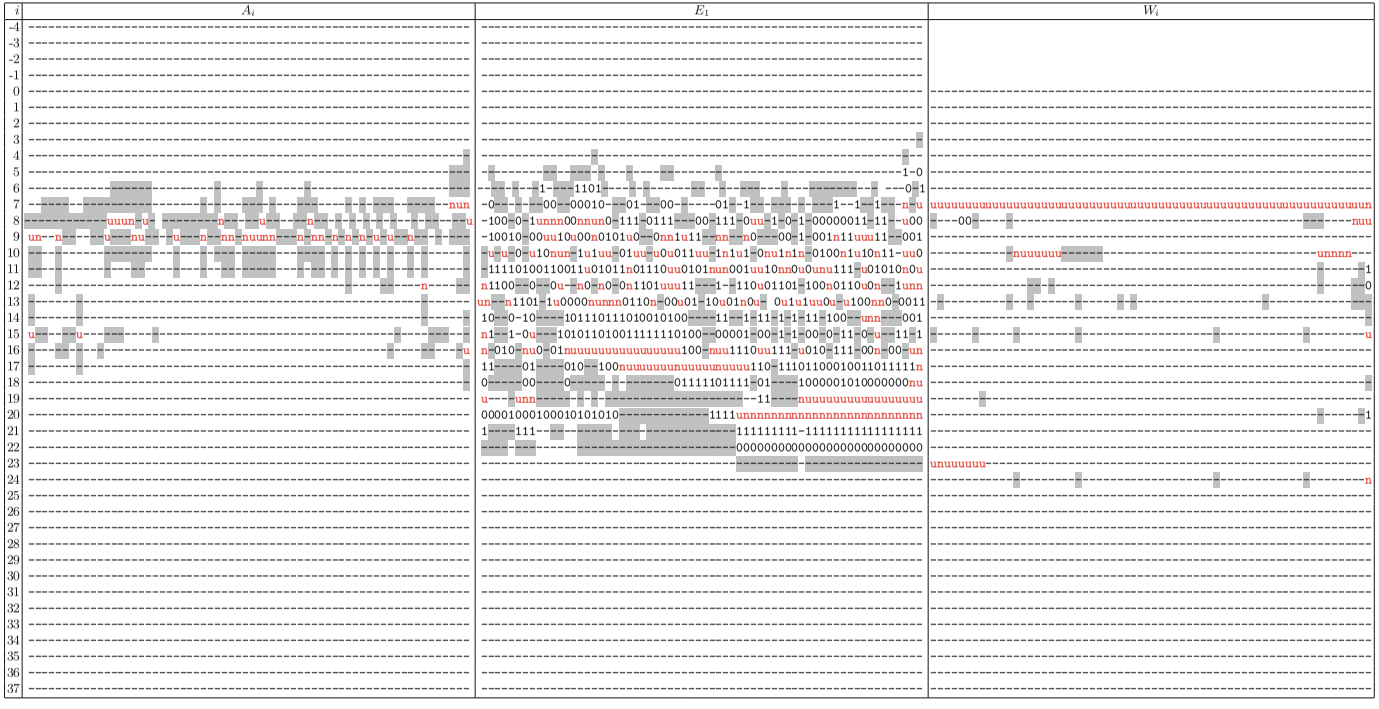

5 Conclusions

In this work, we have improved the best semi-free-start collisions on SHA-512 from 24 to 38 steps. Our attack has a practical complexity of $2^{40.5}$ and we have shown a colliding message pair. We get this result by applying the semi-free-start collision attack on 38 steps of SHA-256 to SHA-512. However, due to the increased word size, and hence increased search space, a straight-forward extension was not possible.

To get these results we have analyzed possible improvements in the guess-and-determine approach to find differential characteristics and conforming message pairs. We got the best results by optimizing the branching heuristic using ideas from SAT solvers. Our heuristic performs a randomized look-ahead selection of candidates which should be guessed first.

Future work includes to apply the look-ahead heuristic to more complex designs. Also, other techniques from SAT solvers may improve guess-and-determine attacks in differential cryptanalysis. However, a direct application of SAT solver techniques without taking high-level information on differential cryptanalysis into account is usually not successful. Finally, an open question is how to use our new results to improve the collision attacks on the SHA-512 hash function.

Acknowledgments. The work has been supported in part by the Secure Information Technology Center-Austria (A-SIT), by the Austrian Government through the research program FIT-IT Trust in IT Systems (Project SePAG, Project Number 835919), and by the European Commission through the FP7 Joint Technology Initiatives (Call ARTEMIS-2012-1, Project Arrowhead, Grant Agreement Number 332987).

References

1. Aoki, K., Guo, J., Matusiewicz, K., Sasaki, Y., Wang, L.: Preimages for step-reduced SHA-2. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 578–597. Springer, Heidelberg (2009)
2. Bernstein, D.J., Lange, T.: eBASH: ECRYPT benchmarking of all submitted hashes, January 2011. http://bench.cr.yp.to/ebash.html
3. Buro, M., Kleine-Büning, H.: Report on a SAT competition. Bull. Eur. Assoc. Theor. Comput. Sci. **49**, 143–151 (1993)
4. Canteaut, A. (ed.): FSE 2012. LNCS, vol. 7549. Springer, Heidelberg (2012)
5. Cramer, R. (ed.): EUROCRYPT 2005. LNCS, vol. 3494. Springer, Heidelberg (2005)
6. Davis, M., Logemann, G., Loveland, D.W.: A machine program for theorem-proving. Commun. ACM **5**(7), 394–397 (1962)
7. De Cannière, C., Rechberger, C.: Finding SHA-1 characteristics: general results and applications. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 1–20. Springer, Heidelberg (2006)
8. Eén, N., Sörensson, N.: An extensible SAT-solver. In: Giunchiglia and Tacchella [11], pp. 502–518
9. Eichlseder, M., Mendel, F., Nad, T., Rijmen, V., Schläffer, M.: Linear propagation in efficient guess-and-determine attacks. In: Budaghyan, L., Helleseth, T., Parker, M. G. (eds.) WCC (2013). http://www.selmer.uib.no/WCC2013/

10. Freeman, J.W.: Improvements to propositional satisfiability search algorithms. Ph.D. thesis, Departement of computer and Information science, University of Pennsylvania, Philadelphia (1995)
11. Giunchiglia, E., Tacchella, A. (eds.): SAT 2003. LNCS, vol. 2919. Springer, Heidelberg (2004)
12. Goldberg, E.I., Novikov, Y.: BerkMin: a fast and robust SAT-solver. In: DATE, pp. 142–149. IEEE Computer Society (2002)
13. Herbstritt, M., Becker, B.: Conflict-based selection of branching rules. In: Giunchiglia and Tacchella [11], pp. 441–451
14. Heule, M., van Maaren, H.: March_dl: adding adaptive heuristics and a new branching strategy. JSAT **2**(1–4), 47–59 (2006)
15. Heule, M., van Maaren, H.: Look-ahead based SAT solvers. In: Biere, A., van Heule, M., Maaren, H., Walsh, T. (eds.) Handbook of Satisfiability. Frontiers in Artificial Intelligence and Applications, vol. 185, pp. 155–184. IOS Press, Amsterdam (2009)
16. Indesteege, S., Mendel, F., Preneel, B., Rechberger, C.: Collisions and other non-random properties for step-reduced SHA-256. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC 2008. LNCS, vol. 5381, pp. 276–293. Springer, Heidelberg (2009)
17. Jeroslow, R.G., Wang, J.: Solving propositional satisfiability problems. Ann. Math. Artif. Intell. **1**, 167–187 (1990)
18. Johansson, T., Nguyen, P.Q. (eds.): EUROCRYPT 2013. LNCS, vol. 7881. Springer, Heidelberg (2013)
19. Khovratovich, D., Rechberger, C., Savelieva, A.: Bicliques for preimages: attacks on Skein-512 and the SHA-2 family. In: Canteaut [4], pp. 244–263
20. Landelle, F., Peyrin, T.: Cryptanalysis of full RIPEMD-128. In: Johansson and Nguyen [18], pp. 228–244
21. Leurent, G.: Analysis of differential attacks in ARX constructions. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 226–243. Springer, Heidelberg (2012)
22. Leurent, G.: Construction of differential characteristics in ARX designs application to skein. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 241–258. Springer, Heidelberg (2013)
23. Li, C.M., Anbulagan: Heuristics based on unit propagation for satisfiability problems. In: IJCAI, vol. 1, pp. 366–371. Morgan Kaufmann, San Francisco (1997)
24. Li, J., Isobe, T., Shibutani, K.: Converting meet-in-the-middle preimage attack into pseudo collision attack: application to SHA-2. In: Canteaut [4], pp. 264–286
25. Liberatore, P.: On the complexity of choosing the branching literal in DPLL. Artif. Intell. **116**(1–2), 315–326 (2000)
26. Mendel, F., Nad, T., Scherz, S., Schläffer, M.: Differential attacks on reduced Ripemd-160. In: Gollmann, D., Freiling, F.C. (eds.) ISC 2012. LNCS, vol. 7483, pp. 23–38. Springer, Heidelberg (2012)
27. Mendel, F., Nad, T., Schläffer, M.: Finding SHA-2 characteristics: searching through a minefield of contradictions. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 288–307. Springer, Heidelberg (2011)
28. Mendel, F., Nad, T., Schläffer, M.: Finding collisions for round-reduced SM3. In: Dawson, E. (ed.) CT-RSA 2013. LNCS, vol. 7779, pp. 174–188. Springer, Heidelberg (2013)
29. Mendel, F., Nad, T., Schläffer, M.: Improving local collisions: new attacks on reduced SHA-256. In: Johansson and Nguyen [18], pp. 262–278
30. Moskewicz, M.W., Madigan, C.F., Zhao, Y., Zhang, L., Malik, S.: Chaff: engineering an efficient SAT solver. In: DAC, pp. 530–535. ACM (2001)

31. National Institute of Standards and Technology. FIPS PUB 180–3: Secure Hash Standard. Federal Information Processing Standards Publication 180–3, U.S. Department of Commerce, October 2008. http://www.itl.nist.gov/fipspubs
32. National Institute of Standards and Technology. FIPS PUB 180–4: Secure Hash Standard. Federal Information Processing Standards Publication 180–4, U.S. Department of Commerce, March 2012. http://www.itl.nist.gov/fipspubs
33. National Institute of Standards and Technology. SHA-3 Selection Announcement, October 2012. http://csrc.nist.gov/groups/ST/hash/sha-3/sha-3_selection_announcement.pdf
34. Nikolić, I., Biryukov, A.: Collisions for step-reduced SHA-256. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 1–15. Springer, Heidelberg (2008)
35. Ouyang, M.: How good are branching rules in DPLL? Discrete Appl. Math. **89**(1–3), 281–286 (1998)
36. Sanadhya, S.K., Sarkar, P.: New collision attacks against up to 24-step SHA-2. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT 2008. LNCS, vol. 5365, pp. 91–103. Springer, Heidelberg (2008)
37. Schläffer, M., Oswald, E.: Searching for differential paths in MD4. In: Robshaw, M. (ed.) FSE 2006. LNCS, vol. 4047, pp. 242–261. Springer, Heidelberg (2006)
38. Shay Gueron, J.W., Johnson, S.: SHA-512/256. Cryptology ePrint Archive, Report 2010/548 (2010). http://eprint.iacr.org/
39. Marques-Silva, J.: The impact of branching heuristics in propositional satisfiability algorithms. In: Barahona, P., Alferes, J.J. (eds.) EPIA 1999. LNCS (LNAI), vol. 1695, pp. 62–74. Springer, Heidelberg (1999)
40. Wang, X., Lai, X., Feng, D., Chen, H., Yu, X.: Cryptanalysis of the hash functions MD4 and RIPEMD. In: Cramer [5], pp. 1–18
41. Wang, X., Yin, Y.L., Yu, H.: Finding collisions in the full SHA-1. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 17–36. Springer, Heidelberg (2005)
42. Wang, X., Yu, H.: How to break MD5 and other hash functions. In: Cramer [5], pp. 19–35