

An Improved Recovery Algorithm for Decayed AES Key Schedule Images

Alex Tsow

The MITRE Corporation*
atsow@mitre.org

Abstract. A practical algorithm that recovers AES key schedules from decayed memory images is presented. Halderman et al. [1] established this recovery capability, dubbed the *cold-boot attack*, as a serious vulnerability for several widespread software-based encryption packages. Our algorithm recovers AES-128 key schedules tens of millions of times faster than the original proof-of-concept release. In practice, it enables reliable recovery of key schedules at 70% decay, well over twice the decay capacity of previous methods. The algorithm is generalized to AES-256 and is empirically shown to recover 256-bit key schedules that have suffered 65% decay. When solutions are unique, the algorithm efficiently validates this property and outputs the solution for memory images decayed up to 60%.

Keywords: anti-tamper, digital forensics, decayed memory, cold-boot attack, AES, key schedule.

1 Introduction

Cold-boot attacks are another troubling example of the increasingly sophisticated threats to security and privacy. In response to these threats we investigate defensive anti-tamper techniques in the hopes of better understanding the potential of specific attack vectors. In this paper we report on our investigation of cold boot attacks and demonstrate that the problem is more serious than previously thought. We present AES key recovery techniques that handle over twice the decay rate of prior methods at comparable computational effort.

The *cold-boot attack* [1] is a serious vulnerability for software-based encryption packages—including BitLocker, FileVault and the open-source project TrueCrypt—where one can recover secret keys from decayed memory images. Decryption with decayed AES keys does not produce original plaintexts. However, the redundancy of key material inherent in the AES key schedule can rectify these faults. When combined with *asymmetric decay*, where bits overwhelmingly decay to their ground state rather than their charged state, this redundancy enables reconstruction of the original key. Heninger and Halderman have developed a recovery algorithm for AES-128 that recovers keys from 30% decayed data in less than 20 minutes about half the time.

* Approved for Public Release; Distribution Unlimited; Tracking Number 09-1872.

Our algorithm recovers keys up to several orders of magnitude faster than Heninger and Halderman’s method. One case that took their algorithm more than 10 days to solve was solved by our improved method in 0.047 seconds. The speed increase enables key recovery from more severely degraded memory images. In an experimental evaluation, our algorithm recovered all keys from a 5,000 case test suite at 70% decay, with 4,927 instances recovered in less than 20 minutes—more than twice decay rate with almost double the success rate in 20 minutes. The speed increase also makes it feasible to enumerate *all* keys from which an image could have decayed, rather than halt on the first key that satisfies the decay and schedule constraints. In particular, the algorithm can determine that a solution is unique. Benchmarks demonstrate feasibility up to 60% decay where there is approximately a $2.5\times$ slowdown compared with the halt-on-first-key search. The algorithm generalizes to 256-bit AES with only a moderate drop in the recovery capability. Empirically, the benchmarks show that AES-256 recovery begins to degrade around 65%; there are no other performance claims in the open literature about AES-256 recovery.

The AES key-schedule is the primary source of key redundancy. For the 128-bit version, the original key is bijectively mapped to 10 additional round-keys [2,3]. The mappings form a system of byte-level equations that constrain the space of likely key candidates.

The asymmetric decay property of DRAM provides a second set of constraints. When the refresh cycle of DRAM is interrupted, the data overwhelmingly decays to 0 (or 1 assuming the complementary encoding for the ground state) because the capacitance is lost over time. Occasionally bits invert to the charged state, although Halderman et al. bound these effects at 0.1%. The asymmetric decay property suggests a compatibility criterion for key candidates: if a candidate schedule subset differs from the decayed memory image only by inversions from the ground state, then it is compatible with the decayed memory image. The performance claims for our algorithm, and indeed those in [1] and [4], are based on the *perfect* asymmetry assumption, where no bit in the ground state ever inverts. The algorithm has also been adapted to accommodate inversions to the charged state by generalizing the compatibility criterion to allow a bounded number of such cases. No other logic changes are necessary.

Key reconstruction is possible because key candidates must satisfy both the asymmetric decay property and the system of equations defined by the AES key schedule. Our algorithm explores a tree of one-byte guesses. At each stage, or tree-depth, the new byte candidate and all bytes implied by the schedule equations are checked against the decayed image. The algorithm guesses bytes in an order such that guess n implies values for n schedule bytes for $n < 11$ in the 128-bit version; guesses 11-14 imply an additional 10 bytes each and the last guess implies the schedule’s remaining 65 bytes.

Byte guessing proceeds in a depth-first manner. Each stage has 256 possibilities, but the schedule and decay-compatibility constraints quickly prune the possibilities, particularly in the later stages where a single byte guess implies several byte values. The selection and order of byte guessing is not unique.

Section 3.4 describes a path selection heuristic for byte guesses which improves recovery times by a factor of several hundred for decay rates over 50%, when compared with the static path implementation.

We make the following contributions: 1) our recovery algorithm is several orders of magnitude faster than the best previously published method, 2) the new method enables key recovery from images with significantly more decay, 3) it enumerates all solutions to a decay image and 4) we generalize the method to AES-256 with little loss of decay performance.

Organization: Section 2 reviews related works. Section 3 describes the algorithm in detail, including a heuristic to optimize the exploration path. Section 4 presents benchmarks for the algorithm with and without the path optimization heuristic. Benchmarks for the unique determination capability and AES-256 key recovery are also presented. Section 5 makes some observations about the benchmark results. Section 6 concludes the paper.

2 Related Work

Halderman et al. [1] established the cold-boot attack as a low-cost way to extract private key information from computers running software encryption. In particular, they extracted private keys for full-disk encryption packages such as BitLocker, FileVault, and the cross-platform open-source project TrueCrypt. Heninger and Halderman released proof-of-concept implementation that recovers 128-bit AES key-schedules.¹ It implements the algorithm from [1] which they have found to recover keys from 30% decayed memory within 20 minutes about half the time. Their archives also contain a recovery algorithm for the RSA cryptosystem.

Heninger and Shacham [4] vastly improved the ability to recover RSA private keys from decayed memory images. They improve recovery from 6% decay (running on the order of minutes) to 46% decay (running on the order of seconds) when p , q , d , d_p , and d_q are in the image. The paper further casts their recovery algorithm in terms of *known bits*, so that the bits may be randomly selected rather than simply the result of an asymmetric memory decay. We note that a perfect memory image maps to 50% known bits under the asymmetric decay assumption, since valid ground-state values theoretically could have decayed.

Nearly a decade before the cold boot attack was demonstrated, Handschuh, Paillier, and Stern modeled *probing attacks* [5] on the square-and-multiply algorithm for modular exponentiation, DES, and RC5. They reconstruct cryptographic secrets by tracing a few critical bits over the target operation's execution. Since cold-boot attacks capture a snapshot of the execution state, these techniques only apply if a trace has been preserved in memory.

Akavia, Goldwasser, and Vaikuntanathan [6] present a model of cold-boot memory recovery in terms of experiments with probabilistic polynomial time players. The recovery player chooses a sequence of probing functions that map private

¹ <http://citp.princeton.edu/memory/code/>

keys to bit vectors; this models key material leakage. They define *adaptive* and *non-adaptive* variants which may or may not alter the choice of probing function in response to the results of previous probes. They further show that the Regev public key cryptosystem [7] is secure under both definitions, but with different leak parameters.

Naor and Segev [8] revisit the above formalism for memory attacks and develop a schema for constructing public key cryptosystems that are resilient against key leakage. The schema relies on the assumptions of a *universal hash proof system* [9] enabling decisional Diffie-Hellman, quadratic residuosity, and Paillier's composite residuosity problem to instantiate the cryptosystem. Alwen, Dodis, and Wichs also examine leakage resilient public key cryptosystems, including identification schemes and authenticated key agreement protocols [10]. They extend their results to the bounded retrieval model, where they consider extremely large keys and an adversary can not learn more than a predetermined bound over a lifetime. Katz further constructs a leakage resistant signature scheme in the standard model [11].

Chari et al. propose the first theoretical model for power analysis [12]. Coron, Naccache, and Kocher develop a similar formalism for characterizing leakage immunity, and present several leakage detection tests [13]. Micali and Reyzin propose a general framework for security against side-channel analysis [14]. These models do not account for memory remanence or cold-boot attacks.

Countermeasures to cold-boot attacks remain scarce within the current technology paradigms. Migrating to hardware embedded encryption, such as that proposed by the Trusted Computing Group's Opal platform [15], will mitigate cold-boot attacks on full-disk encryption. Enck et al. [16] propose an encrypting memory controller that writes only encrypted data to main memory, but decrypts it on reads into the processor or cache. There have also been attempts to manipulate the Intel x86 cache-coherence model to ensure that keys and key-derived state (such as key schedules) remain in L2 caches, but not in main memory.² The feasibility of this approach has yet to be demonstrated with the current architectures, however vendor modification of the instruction set may indeed make this approach a reality.

Intel has developed specialized instructions for executing AES operations [17]. There are six kinds instructions (encrypt round, encrypt last round, decrypt round, decrypt last round, inverse mix columns, and key schedule assist) which use the 128-bit XMM registers to hold round keys and block data. In addition to improving execution speed, these instructions have been designed to eliminate vulnerabilities from *cache attacks* [18], an interprocess side-channel that exploits timing differences for operations dependent upon cached and uncached data. Intel does not claim that these instructions mitigate cold-boot attacks, however we speculate that the schedule derivation assistance may improve the performance of just-in-time round-key derivation, thereby reducing the number of round keys stored in memory.

² Jürgen Pabel. <http://frozencache.blogspot.com/>

3 Algorithmic Description

This primary exposition details the recovery algorithm for AES-128, although the concepts generalize to the 192-bit and 256-bit cases (Section 3.5). Their differing block and key sizes create more potential for confusion when referencing schedule elements. We have implemented the 128-bit and 256-bit cases; Section 4 presents performance results for both cases.

3.1 Preliminaries

A 128-bit AES key schedule expands a four by 32-bit-word key into a 44 word sequence. Schedule components are addressed with the following notation: Let sans-serif variables, S , refer to entire key schedule. Subscripting expresses a hierarchical view of schedule components. Let S_r refer to the four words of round r . Let $S_{r,w}$ refer to word w of round r . Let $S_{r,w,b}$ refer to byte b of word w of round r . It is also convenient to index the schedule in a flat manner. Let S^w_i refer to word i of the schedule. Let S^b_i refer to byte i of the schedule. The notation follows the least-significant-byte-first convention. Some additional function notation is necessary to express the key schedule. Let $\text{sbox}(S^b_i)$ apply the AES substitution box to the byte S^b_i . For convenience, let $\text{sbox}(S^w_i)$ apply the substitution box to each constituent byte when S^w_i is a word. Let $\text{rot}(S^w_i)$ rotate the word S^w_i by eight bit positions of increasing significance; e.g. $\text{rot}(S_{r,w,0}, S_{r,w,1}, S_{r,w,2}, S_{r,w,3}) = (S_{r,w,3}, S_{r,w,0}, S_{r,w,1}, S_{r,w,2})$, in the least-significant-byte-first representation. The round constants, denoted by $\text{rcon}[i]$, are the $(i - 1)^{\text{th}}$ exponent of 2 in the field $GF(2^8)$ for the least significant byte and 0 for the other bytes.

For the 128-bit schedule, the first four words are the key itself. The subsequent words are prescribed by two equation schema.

$$\begin{aligned} S^w_i &= S^w_{i-1} \oplus \text{sbox}(\text{rot}(S^w_{i-3})) \oplus \text{rcon}[i/4], & \text{when } i \bmod 4 = 0 \\ S^w_i &= S^w_{i-1} \oplus S^w_{i-3}, & \text{when } i \bmod 4 \neq 0 \end{aligned} \tag{1}$$

The table below illustrates the indexing schema for the 128-bit key schedule. The column headings indicate the byte and word indices. Row labels indicate the round. This particular table shows how the flat byte-level references relate to the hierarchical tags. For instance, $S_{8,2,3}$ refers to the same byte as S^b_{139} .

$r \setminus b$	0				1				2				3			
	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3
0	S^b_0	S^b_1	S^b_2	S^b_3	S^b_4	S^b_5	S^b_6	S^b_7	S^b_8	S^b_9	S^b_{10}	S^b_{11}	S^b_{12}	S^b_{13}	S^b_{14}	S^b_{15}
1	S^b_{16}	S^b_{17}	S^b_{18}	S^b_{19}	S^b_{20}	S^b_{21}	S^b_{22}	S^b_{23}	S^b_{24}	S^b_{25}	S^b_{26}	S^b_{27}	S^b_{28}	S^b_{29}	S^b_{30}	S^b_{31}
2	S^b_{32}	S^b_{33}	S^b_{34}	S^b_{35}	S^b_{36}	S^b_{37}	S^b_{38}	S^b_{39}	S^b_{40}	S^b_{41}	S^b_{42}	S^b_{43}	S^b_{44}	S^b_{45}	S^b_{46}	S^b_{47}
3	S^b_{48}	S^b_{49}	S^b_{50}	S^b_{51}	S^b_{52}	S^b_{53}	S^b_{54}	S^b_{55}	S^b_{56}	S^b_{57}	S^b_{58}	S^b_{59}	S^b_{60}	S^b_{61}	S^b_{62}	S^b_{63}
4	S^b_{64}	S^b_{65}	S^b_{66}	S^b_{67}	S^b_{68}	S^b_{69}	S^b_{70}	S^b_{71}	S^b_{72}	S^b_{73}	S^b_{74}	S^b_{75}	S^b_{76}	S^b_{77}	S^b_{78}	S^b_{79}
5	S^b_{80}	S^b_{81}	S^b_{82}	S^b_{83}	S^b_{84}	S^b_{85}	S^b_{86}	S^b_{87}	S^b_{88}	S^b_{89}	S^b_{90}	S^b_{91}	S^b_{92}	S^b_{93}	S^b_{94}	S^b_{95}
6	S^b_{96}	S^b_{97}	S^b_{98}	S^b_{99}	S^b_{100}	S^b_{101}	S^b_{102}	S^b_{103}	S^b_{104}	S^b_{105}	S^b_{106}	S^b_{107}	S^b_{108}	S^b_{109}	S^b_{110}	S^b_{111}
7	S^b_{112}	S^b_{113}	S^b_{114}	S^b_{115}	S^b_{116}	S^b_{117}	S^b_{118}	S^b_{119}	S^b_{120}	S^b_{121}	S^b_{122}	S^b_{123}	S^b_{124}	S^b_{125}	S^b_{126}	S^b_{127}
8	S^b_{128}	S^b_{129}	S^b_{130}	S^b_{131}	S^b_{132}	S^b_{133}	S^b_{134}	S^b_{135}	S^b_{136}	S^b_{137}	S^b_{138}	S^b_{139}	S^b_{140}	S^b_{141}	S^b_{142}	S^b_{143}
9	S^b_{144}	S^b_{145}	S^b_{146}	S^b_{147}	S^b_{148}	S^b_{149}	S^b_{150}	S^b_{151}	S^b_{152}	S^b_{153}	S^b_{154}	S^b_{155}	S^b_{156}	S^b_{157}	S^b_{158}	S^b_{159}
10	S^b_{160}	S^b_{161}	S^b_{162}	S^b_{163}	S^b_{164}	S^b_{165}	S^b_{166}	S^b_{167}	S^b_{168}	S^b_{169}	S^b_{170}	S^b_{171}	S^b_{172}	S^b_{173}	S^b_{174}	S^b_{175}

A candidate key schedule byte, $C^b_{i_1}$, with ground states specified by $M^b_{i_1}$ is *compatible with the decayed byte*, $D^b_{i_1}$, when $D^b_{i_1}$ preserves all ground-state bits in $C^b_{i_1}$, or expressed equationally, when $(C^b_{i_1} \oplus D^b_{i_1}) \wedge (C^b_{i_1} \oplus M^b_{i_1}) = 0$.

3.2 Maximizing the Implied Schedule Bytes

For the first guess, the candidate byte, $C^b_{i_0}$, is only constrained by the known bits in corresponding decayed byte, $D^b_{i_0}$. Yet in the second stage, for a properly chosen i_1 , $C^b_{i_1}$ is constrained by $D^b_{i_1}$ and a second byte $D^b_{j_1}$. A properly selected i_1 will instantiate a byte slice of one of the two schedule generating equations, (1). For example $C^b_4 \oplus C^b_{16} = C^b_{20}$ is the first byte-slice of the generating equation for $C^w_5 = (C^b_{20}, C^b_{21}, C^b_{22}, C^b_{23})$. If $i_0 = 4$ and $i_1 = 20$, then the implied byte is at index $j = 16$ and equals $C^b_{20} \oplus C^b_4$. Thus D^b_{16} constrains the implied value of $C^b_{20} \oplus C^b_4$.

This algorithm makes use of the following observations: 1) Each schedule byte $S^b_i, 16 \leq i < 160$ is involved in three equations. 2) There are 256 solutions to each equation when any one variable is fixed. There is a unique solution to each equation when any two variables are fixed. 3) Guessing a single byte at stage n implies up to n other byte values for properly structured guessing orders.

Item 1 follows from the fact that every word is generated by its preceding word and the one 4 words ago; simply limit the scope to the byte-slice of the concerned byte. For example, consider $S_{1,0,0} = S^b_{16}$, the first byte of S^w_4 .

$$\begin{aligned}
 S_{0,0,0} \oplus \text{sbox}(S_{0,3,1}) \oplus 0x01 &= S_{1,0,0} \\
 S_{0,1,0} \oplus S_{1,0,0} &= S_{1,1,0} \\
 S_{1,0,0} \oplus \text{sbox}(S_{1,3,1}) \oplus 0x02 &= S_{2,0,0}
 \end{aligned}
 \tag{2}$$

Item 2 is implied by the fact that \oplus is the field addition operation for $GF(2^8)$ and that $\text{sbox}()$ is a bijection in $GF(2^8)$.

To see item 3, consider the following candidate exploration order, C , for a decayed schedule D . First choose a candidate for $C_{0,0,0}$; it is only constrained by the known bits at that position. The next guess, $C_{1,0,0}$, is constrained by the known bits from $D_{1,0,0}$. Additionally, one can solve the equation $C_{0,0,0} \oplus$

Table 1. An order for byte-guesses and consequent values in AES-128; the end of Section 3.2 describes the scripting notation

w	0				1				2				3			
	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3
0	0 ₀	14 ₁₀				13 ₁₀				12 ₁₀				1	14 ₉	
1	1 ₀	13 ₉				12 ₉				2 ₂	14 ₈			2 ₁	13 ₈	
2	2 ₀	12 ₈				3 ₃	14 ₇			3 ₂	13 ₇			3 ₁	12 ₇	
3	3 ₀	4 ₄	14 ₆			4 ₃	13 ₆			4 ₂	12 ₆			4 ₁	5 ₅	14 ₅
4	4 ₀	5 ₄	13 ₅			5 ₃	12 ₅			5 ₂	6 ₆	14 ₄		5 ₁	6 ₅	13 ₄
5	5 ₀	6 ₄	12 ₄			6 ₃	7 ₇	14 ₃		6 ₂	7 ₆	13 ₃		6 ₁	7 ₅	12 ₃
6	6 ₀	7 ₄	8 ₈	14 ₂		7 ₃	8 ₇	13 ₂		7 ₂	8 ₆	12 ₂	14 ₁	7 ₁	8 ₅	9 ₉
7	7 ₀	8 ₄	9 ₈	13 ₁		8 ₃	9 ₇	12 ₁	14 ₀	8 ₂	9 ₆	10 ₁₀	13 ₀	8 ₁	9 ₅	10 ₉
8	8 ₀	9 ₄	10 ₈	12 ₀	15 ₁₀	9 ₃	10 ₇	11 ₁₀	15 ₉	9 ₂	10 ₆	11 ₉	15 ₈	9 ₁	10 ₅	11 ₈
9	9 ₀	10 ₄	11 ₇	15 ₇		10 ₃	11 ₆	15 ₆		10 ₂	11 ₅	15 ₅		10 ₁	11 ₄	15 ₄
100	10 ₀	11 ₃	15 ₃			11 ₂	15 ₂			11 ₁	15 ₁			11 ₀	15 ₀	

$\text{sbox}(C_{0,3,1}) = C_{1,0,0}$ for $C_{0,3,1}$, since $C_{0,0,0}$ and $C_{1,0,0}$ have candidate values. Thus the second guess is constrained by the known values at $D_{0,3,1}$ as well. By the same logic, $D_{2,0,0}$, $D_{1,3,1}$, and $D_{1,2,1}$ constrain the compatible guesses for $C_{2,0,0}$. Because $C_{r-1,3}$ rotates when computing $C_{r,0}$, continuing to propose bytes in the column $C_{r,0,0}$ causes the implied byte indices increment modulo 4 when the word index wraps around modulo 4.

Table 1 illustrates this behavior by enumerating the order of byte guesses and their consequent bytes. In the table entries, the full-sized number indicates the guessing stage and the subscript indicates the sequence of implied bytes. In particular, i_0 indicates a guess for stage i and i_1 is the first implied byte from this candidate. For example, $C_{5,0,0}$ is guess number 5. This value combined with the value for $C_{4,0,0}$ (chosen in step 4), allows one to solve for $C_{4,3,1}$. The following equations make the order of solution explicit.

$$\begin{aligned}
 C_{4,3,1} &= \text{sbox}^{-1}(C_{5,0,0} \oplus C_{4,0,0} \oplus 01) \\
 C_{4,2,1} &= C_{4,3,1} \oplus C_{3,3,1} \\
 C_{4,1,1} &= C_{4,2,1} \oplus C_{3,2,1} \\
 C_{4,0,1} &= C_{4,1,1} \oplus C_{3,1,1} \\
 C_{3,3,2} &= \text{sbox}^{-1}(C_{4,0,1} \oplus C_{3,0,1} \oplus 01)
 \end{aligned} \tag{3}$$

After selecting candidates for bytes 0-10, there are a number of ways to guess bytes 11-15. Table 1 illustrates a choice for these positions that implies values for an additional 10 bytes for each guess. The last guess implies 65 bytes because it causes round 8 to be fully specified; the entire schedule may be derived from any complete round.

3.3 The Recovery Algorithm

The algorithm, `recoverKeyRec`, explores the candidate space one byte at a time. It exploits the constraints on guesses and their consequent bytes to prune its exploration tree. For each guess, `recoverKeyRec` considers all 256 possible values. If the candidate satisfies all constraints imposed by the decayed image, D , then it guesses a value for the next step. Exploration proceeds in a depth-first manner, so that guess i is incremented to the next compatible candidate when all of descendant candidates have been ruled out.

A breadth-first search is also possible, however this strategy greatly increases the memory utilization. The advantage of the breadth-first method is that one could track the distance from decayed data for each candidate and explore the closest options first. The first implementations of this algorithm maintained the candidates on a binary heap indexed by their cost. In practice, there were many cases in the 60-70% decay range, where the process exceeded its 31-bit address space and halted before recovering the key. On the other hand, recovery speeds of depth-first search up to the 70% rate have proven fast enough in most cases and solvable in all attempted cases. The need to solve more cases with less memory dictated a transfer to depth-first search which consumes a fixed amount of memory, about 4.2 MB in the experimental implementation. Because depth-first search has a small memory footprint, it is also inexpensive to halt tree

```

recoverKeyRec(CandidateMatrix c, DecaySchedule d):
  if (c.length()==16):
    return c.key()
  for i=0 to 255:
    if(d.isCompatible(c.guess(i))):
      key = recoverKeyRec(c.guess(i),d)
      if (key != NULL)
        return key
  return NULL

```

Fig. 1. Recursive expression of key recovery; Appendix A details the core methods

exploration and resume later. The only data necessary to save is the candidate being examined at the halting time. Breadth-first search would require one to save the binary heap of candidates.

The following describes the semantics for tokens in `recoverKeyRec` (Figure 1). For simplicity of expression assume that operations do not mutate objects, but return newly constructed objects. Italicized tokens refer to classes and teletype tokens refer to fields and methods.

Let *c* be a *CandidateMatrix* that contains candidates for schedule bytes in the order indicated by Table 1. *CandidateMatrix* maintains a `count` of how many guesses have been made (0-16) and a flat schedule representation of 176 bytes to store byte candidates and their consequent bytes. The method `guess(Byte b)`, returns a new *CandidateMatrix* whose array has been updated to contain *b* at position `count0` and its consequent bytes at positions as specified by the path matrix (e.g., Table 1). In particular, guessing the 16th byte completes the entire schedule. The `key()` method returns the key once all 16 bytes have been guessed and validated. The new `CandidateMatrix()` constructor simply creates an empty array with `count` set to 0.

Let *DecaySchedule* contain the decayed key schedule and a predicate, `isCompatible(CandidateMatrix)`, that indicates whether or not a guess and its consequent bytes are compatible the decayed schedule. Compatibility is determined by checking that the *CandidateMatrix* contains all the known bits from the corresponding bytes of the decayed schedule.

The function `recoverKeyRec(CandidateMatrix,DecaySchedule)` returns a key whose schedule is decay-compatible and is the result of extending the incoming *CandidateMatrix*. It returns NULL when there is no compatible key schedule extension to the specified candidate prefix. Proper usage asserts that the starting *CandidateMatrix* and *DecaySchedule* are compatible. The initial call to `recoverKeyRec` begins with an empty *CandidateMatrix*. The recursive expression in Figure 1 makes the control logic explicit. Figure 1 halts on the first compatible key schedule, however one could modify it to halt after a full search of the key space; simply replace the `return` with a `print` on the third line. Section 4 benchmarks both variants.

3.4 Path Prioritization

The exploration path illustrated by Table 1 maximizes the number of implied schedule bytes with the goal of minimizing the number of compatible candidates at each stage. There are many ways to grow the selection path, and Table 1 illustrates just one of them. For instance, there are 3 symmetric alternatives obtained by rotating the bytes of each word. Different paths will encounter different constraints and therefore will result in varying recovery times. Within the stages 0-10, the selection order of $S_{x,0,0}$ may be altered and still obtain the same consequent bytes after 10 candidate stages. Growing the path with guesses adjacent to the body of previous guesses will preserve the set of inferred bytes; this claim has been verified experimentally. The following matrix shows that guesses may be grown from the middle of the schedule:

The exploration path in Table 2 starts at $S_{5,0,0}$. To maintain adjacency, the next choice may be $S_{4,0,0}$ or $S_{6,0,0}$. The path grows by extending either the top or bottom of previous choices in $S_{x,0,0}$, where $0 \leq x < 11$. This allows the number of inferred bytes to grow by one at each stage.

A first heuristic for choosing the best path might be to count up the known bits in the exploration path and its consequent bytes. However within a set of guesses and implied bytes, the selection order can also make a difference. Past a certain threshold, adding more constraints does not prune the exploration anymore because the byte is already uniquely determined (or its parent has been ruled out). Consider the case when the decayed data in the candidate position is $0xFF$. If the last guess corresponds to this position, then all of the consequent bytes are wasted constraints because the byte is uniquely determined. On the other hand, that byte position would make an excellent initial position for stage 0 because the first guess never produces any consequent bytes. The final algorithm's path comparison heuristic estimates the number of branches at each stage. At a stage, the estimate simply counts the number of known bits in the initial and consequent positions. The intuition is that each known bit will reduce the number of valid branches by a factor of 2. Thus, the branch estimate for a stage i is $2^{\min(8, k_i)}$ where k_i is the number of known bits in the i^{th} stage's initial and consequent byte positions. This is clearly only an estimate, as some

Table 2. An exploration path starting in round 5; script notation parallels the example at the end of Section 3.2

w	0				1				2				3			
$x \setminus b$	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3
0	10 ₀															
1	7 ₀								10 ₂				7 ₁			
2	6 ₀				10 ₃				7 ₂				6 ₁			
3	3 ₀	10 ₄			7 ₃				6 ₂				3 ₁	10 ₅		
4	2 ₀	7 ₄			6 ₃				3 ₂	10 ₆			2 ₁	7 ₅		
5	0 ₀	6 ₄			3 ₃	10 ₇			2 ₂	7 ₆			1 ₁	6 ₅		
6	1 ₀	4 ₄	10 ₈		4 ₃	7 ₇			4 ₂	6 ₆			4 ₁	5 ₅	10 ₉	
7	4 ₀	5 ₄	8 ₈		5 ₃	8 ₇			5 ₂	8 ₆	10 ₁₀		5 ₁	8 ₅	9 ₉	
8	5 ₀	8 ₄	9 ₈		8 ₃	9 ₇			8 ₂	9 ₆			8 ₁	9 ₅		
9	8 ₀	9 ₄			9 ₃				9 ₂				9 ₁			
10	9 ₀															

Table 3. AES-256 exploration path template; script notation parallels the example at the end of Section 3.2

0				1				2				3				0				1				2				3									
0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3		
0 ₆	15 ₇	23 ₇	31 ₇	14 ₇	22 ₇	30 ₇	13 ₇	21 ₇	29 ₇	12 ₇	20 ₇	28 ₇	10 ₆	18 ₆	26 ₆	9 ₆	17 ₆	25 ₆	8 ₆	16 ₆	24 ₆	25 ₆	8 ₆	16 ₆	24 ₆	31 ₆	8 ₆	16 ₆	24 ₆	31 ₆	11 ₁₅	15 ₆	23 ₆	23 ₆			
1 ₆	14 ₆	22 ₆	30 ₆	13 ₆	21 ₆	29 ₆	12 ₆	20 ₆	28 ₆	10 ₅	18 ₅	26 ₅	9 ₅	17 ₅	25 ₅	8 ₅	16 ₅	24 ₅	31 ₅	8 ₅	16 ₅	24 ₅	31 ₅	3 ₅	15 ₄	23 ₄	30 ₄	3 ₅	15 ₄	23 ₄	30 ₄	2 ₁₄	14 ₅	22 ₅	22 ₅		
2 ₆	13 ₆	21 ₆	29 ₆	12 ₅	20 ₅	28 ₅	10 ₄	18 ₄	26 ₄	9 ₄	17 ₄	25 ₄	8 ₄	16 ₄	24 ₄	31 ₄	3 ₄	15 ₃	23 ₃	30 ₃	3 ₄	15 ₃	23 ₃	30 ₃	4 ₃	14 ₃	22 ₃	29 ₃	4 ₃	14 ₃	22 ₃	29 ₃	1 ₁₃	13 ₄	21 ₄	21 ₄	
3 ₆	12 ₆	20 ₆	28 ₆	10 ₃	18 ₃	26 ₃	9 ₃	17 ₃	25 ₃	8 ₃	16 ₃	24 ₃	31 ₃	4 ₄	15 ₃	23 ₃	30 ₃	4 ₃	14 ₃	22 ₃	29 ₃	4 ₃	14 ₃	22 ₃	29 ₃	5 ₂	13 ₂	21 ₂	28 ₂	5 ₂	13 ₂	21 ₂	28 ₂	4 ₁	12 ₃	20 ₃	20 ₃
4 ₆	10 ₆	18 ₆	26 ₆	9 ₂	17 ₂	25 ₂	8 ₂	16 ₂	24 ₂	31 ₂	5 ₅	15 ₂	23 ₂	30 ₂	5 ₄	14 ₂	22 ₂	29 ₂	5 ₃	13 ₂	21 ₂	28 ₂	5 ₃	13 ₂	21 ₂	28 ₂	5 ₂	12 ₂	20 ₂	26 ₂	5 ₁	10 ₂	18 ₂	18 ₂			
5 ₆	9 ₆	17 ₆	25 ₆	8 ₁	16 ₁	24 ₁	31 ₁	6 ₆	15 ₁	23 ₁	30 ₁	6 ₅	14 ₁	22 ₁	29 ₁	6 ₄	13 ₁	21 ₁	28 ₁	6 ₃	12 ₁	20 ₁	26 ₁	6 ₂	10 ₁	18 ₁	26 ₁	6 ₁	10 ₁	18 ₁	26 ₁	6 ₀	8 ₁	16 ₁	16 ₁		
6 ₆	8 ₆	16 ₆	24 ₆	31 ₀	7 ₇	15 ₀	23 ₀	30 ₀	7 ₆	14 ₀	22 ₀	29 ₀	7 ₅	13 ₀	21 ₀	28 ₀	7 ₄	12 ₀	20 ₀	27 ₀	7 ₃	11 ₀	19 ₀	27 ₀	7 ₂	11 ₀	19 ₀	27 ₀	7 ₁	11 ₀	19 ₀	27 ₀	7 ₀	11 ₀	19 ₀	27 ₀	
7 ₆	11 ₃	19 ₃	27 ₃	11 ₂	19 ₂	27 ₂	11 ₁	19 ₁	27 ₁	11 ₀	19 ₀	27 ₀																									

bits may constrain portions of the guess that have already been determined. The total branching estimate is the product of the stage estimates. Since the stage estimates are all powers of two, it suffices to sum their exponents. Thus, the heuristic ranks paths by the scalar value $\sum_{i=0}^{15} \min(8, k_i)$.

The algorithm considers variants of Table 1 along two axes. One axis is byte slice selection; Table 1 initiates on slice 0, although three other may be obtained by rotating all schedule words by the same amount. The other axis is byte guessing order within stages 0-10 as described above; choose a initial round and then extend the guesses by adding to the adjacent round on the top or the bottom. There are other paths not reached by these variables; e.g., growth may be rooted in words 1-3 of the round key rather than 0. Performance may well improve by selecting these paths from a larger space, however no additional path analysis is investigated in this work.

3.5 Generalizing to Other Instances of AES

Generalizing the algorithm to operate on 192-bit and 256-bit variants of AES is straightforward. One needs to construct a different path template and update the `isCompatible()` method to incorporate the schedule generating equations of the larger keys. Table 3 illustrates the path template used by the 256-bit implementation. The heuristic considers paths on the same axes as the 128-bit version. Since the key size is twice the length of the block size, the matrix is only 8 deep and therefore the heuristic considers fewer paths.

4 Benchmarks

Test cases are generated with OpenSSL’s `RAND_bytes()` function. For a given decay rate d , the test generator derives a key schedule from randomly selected key bytes and then randomly zeroes $d\%$ of the bits. Tests assume a ground-state encoding of 0. Performance was evaluated on Dell Precision Workstation 7400 running a 3.4 Ghz quad-core Xeon processor with 4GB of RAM. The C99 reference implementation of the algorithm is compiled with the MinGW version of `gcc-3.4.5` at the highest level of optimization, `-O4`. All computations cases were run to completion in a serial manner. Time was measured by entry and exit calls to `clock()` from the `time.h` library; clock resolution is 64 Hz. The original

Table 4. Run-time results for four versions of the algorithm; each decay rate test suite contains 10,000 cases

Case	Key size	Path selection	Halting condition
PathOpt-128	128 bits	Heuristically chosen (Section 3.4)	First match
PathOpt-256	256 bits	Heuristically chosen (Section 3.4)	First match
Basic-128	128 bits	Fixed to Table 1	First match
Exhaust-128	128 bits	Heuristically chosen (Section 3.4)	End of key space

PathOpt-128 — Run-time (seconds)

Decay	30%	40%	50%	60%
Total	90.120	93.559	142.322	1,736.321
Avg.	0.009	0.009	0.014	0.174
Med.	0.015	0.015	0.015	0.031
Max	0.015	0.015	0.078	2.094
Min	0.000	0.000	0.000	0.000
St.Dev.	0.007	0.008	0.015	0.772

PathOpt-256 — Run-time (seconds)

Decay	30%	40%	50%	60%
Total	17.046	26.185	123.250	6,954.231
Avg.	0.002	0.003	0.012	0.695
Med.	0.000	0.000	0.000	0.062
Max	0.016	0.062	2.125	352.015
Min	0.000	0.000	0.000	0.000
St.Dev.	0.005	0.006	0.044	5.920

Basic-128 — Run-time (seconds)

Decay	30%	40%	50%	60%
Total	219.204	1,526.308	32,551.469	1,638,788.166
Avg.	0.022	0.153	3.255	163.879
Med.	0.015	0.015	0.078	1.968
Max	9.562	266.390	3,354.890	343,656.375
Min	0.000	0.000	0.000	0.000
St.Dev.	0.140	2.994	55.563	3,753.608

Exhaust-128 — Run-time (seconds)

Decay	30%	40%	50%	60%
Total	96.403	112.350	258.568	4,497.599
Avg.	0.010	0.011	0.026	0.450
Med.	0.015	0.015	0.015	0.110
Max	0.031	0.468	0.875	75.203
Min	0.000	0.000	0.000	0.000
St.Dev.	0.007	0.009	0.036	1.921

keys were found for all test cases, using the heuristically chosen path and halting on the first match.

Table 4 summarizes the benchmark results for four variants of the algorithm: PathOpt-128, PathOpt-256, Basic-128, and Exhaust-128. There are 10,000 cases for each of the four decay rates, 30%, 40%, 50%, and 60%. The 128-bit variants have been run on the same test cases, so their results are directly comparable. All times are measured in seconds. A time of 0.000 means that the computation finished in less than 1/64 second, or about 53 million processor cycles.

Additional testing (Table 5) was performed to estimate the maximum recoverable decay rates for PathOpt-128 and PathOpt-256. Only 5,000 cases were examined due to extended recovery times.

5 Analysis

PathOpt-128 solves all cases at 50% decay and less in under half a second. At 60% decay, PathOpt-128 recovered the worst case in 35.500 seconds while solving the average case in 0.174 seconds. At the extended decay rate of 70%, recovery time averages grew to just over 6 minutes with the median time at just under five seconds. Nearly half of the 17.4 day run was consumed by solving the worst case of the test suite; the second slowest case was over six times faster. 4927 cases were recovered in less than 20 minutes.

PathOpt-128 runs faster than Basic-128 across the board and the speedup quickly grows as the decay rate increases. The speedups for 30%, 40%, 50%, and

Table 5. Extended decay rate runs; each decay rate test suite has 5,000 cases

Case	Total	Avg.	Med.	Max	Min	St.Dev.
PathOpt-128 @ 70% decay	1,504,487.119 s	300.897 s	4.938 s	737,266.687 s	0.000 s	10,677.913 s
PathOpt-256 @ 65% decay	446,879.849 s	89.376 s	0.875 s	194,410.875 s	0.000 s	2,843.061 s

60% are $2.43\times$, $16.3\times$, $228\times$, and $943\times$, respectively. At 70%, only 10 cases had completed after a week when the experiment was terminated. The path selection heuristic makes 70% decay a feasibly solvable problem in the test environment. Even for the low decay rates, Basic-128 has a much higher standard deviation; their worst cases with Basic-128 are several orders of magnitude worse than their worst cases with PathOpt-128.

The profound impact of heuristic path selection at high decay rates suggests that a more thorough search for the best path could further extend the maximum feasible decay capacity. Only a small subset of possible paths are considered. The current analysis takes less than 1/64 second, as evidenced by the 0.000 timing results, so there is ample room for more startup analysis.

Full search of the key space appears to be a small factor slower than stopping at the first compatible key. It widens as the decay increases, but by 60% the Exhaust-128 only takes 2.590 times longer. We note that all 90,000 test cases had precisely one solution, so the exhaustive search seems unnecessary at the tested decay rates.

PathOpt-256 performs well up to 60% decay rates, solving cases in an average of 0.695 seconds and in no more than 352.015 seconds. At 70%, no cases were solved in the test suite during a 1 week trial. At 65% the results are promising: the 5,000 cases have been solved in an average of 89.676 seconds. The longest case took 2.25 days to recover, while 99.4% of the cases have been recovered in less than 20 minutes.

Interestingly, the PathOpt-256 is slightly faster than PathOpt-128 on decay rates at 50% and below. The average solution times for these cases is within two units of the 64 Hz clock resolution. We conjecture that the heuristic path analysis takes less time with the 256-bit version since there are fewer paths to consider, due to the flatness of the path matrix (Section 3.5).

As a point of comparison, the original algorithm [1] was compiled in the same environment and run against the same test suite of 30% and 40% decayed AES-128 schedules. After three weeks of execution, only the first four cases at 30% had been solved and the first case at 40% had not yet finished.

6 Conclusion

We presented a new class of recovery capability that is several orders of magnitude faster than previous methods, particularly for higher decay rates. It more than doubles the decay rate recovery feasibility of prior work. The tree-pruning constraints enable efficient and exhaustive key space searches to determine solution uniqueness. We have generalized the implementation to AES-256 while maintaining excellent performance up to 65% decay.

Acknowledgements

I owe a special debt of gratitude to Dr. Adam L. Young for his insight and wisdom regarding writing, presentation, and the scientific community. I have additionally benefited from the thoughtful review of Kerry A. McKay. I am further indebted to Charles C. Howell, Steven M. Godin, and Eileen M. Boettcher for their strident advocacy—without which this paper would not be possible.

References

1. Halderman, J.A., Schoen, S.D., Heninger, N., Clarkson, W., Paul, W., Calandrino, J.A., Feldman, A.J., Appelbaum, J., Felten, E.W.: Lest we remember: cold boot attacks on encryption keys. In: USENIX Security Symposium, pp. 45–60. USENIX Association, Berkeley (2008)
2. Daemen, J., Rijmen, V.: The block cipher Rijndael. In: Schneier, B., Quisquater, J.-J. (eds.) CARDIS 1998. LNCS, vol. 1820, pp. 277–284. Springer, Heidelberg (2000)
3. Daemen, J., Rijmen, V.: The Design of Rijndael. Springer, Heidelberg (2002)
4. Heninger, N., Shacham, H.: Reconstructing RSA private keys from random key bits. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 1–17. Springer, Heidelberg (2009)
5. Handschuh, H., Paillier, P., Stern, J.: Probing attacks on tamper-resistant devices. In: Koç, Ç.K., Paar, C. (eds.) CHES 1999. LNCS, vol. 1717, pp. 303–315. Springer, Heidelberg (1999)
6. Akavia, A., Goldwasser, S., Vaikuntanathan, V.: Simultaneous hardcore bits and cryptography against memory attacks. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 474–495. Springer, Heidelberg (2009)
7. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC, pp. 84–93. ACM, New York (2005)
8. Naor, M., Segev, G.: Public-key cryptosystems resilient to key leakage. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 18–35. Springer, Heidelberg (2009)
9. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002)
10. Alwen, J., Dodis, Y., Wichs, D.: Leakage-resilient public-key cryptography in the bounded-retrieval model. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 36–54. Springer, Heidelberg (2009)
11. Katz, J.: Signature schemes with bounded leakage resilience. Cryptology ePrint Archive: Report 2009/133 (March 22, 2009)
12. Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards sound approaches to counteract power-analysis attacks. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 398–412. Springer, Heidelberg (1999)
13. Coron, J.S., Naccache, D., Kocher, P.: Statistics and secret leakage. ACM Trans. Embed. Comput. Syst. 3(3), 492–508 (2004)
14. Micali, S., Reyzin, L.: Physically observable cryptography. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 278–296. Springer, Heidelberg (2004)

15. TCG storage security subsystem class: Opal version 1.0,
<http://www.trustedcomputinggroup.org/> (January 27, 2009)
16. Enck, W., Butler, K., Richardson, T., McDaniel, P., Smith, A.: Defending against attacks on main memory persistence. In: ACSAC, Washington, DC, USA, pp. 65–74. IEEE Computer Society, Los Alamitos (2008)
17. Gueron, S.: Advanced encryption standard (AES) instructions set (April 27, 2009), <http://www.intel.com/>
18. Osvik, D.A., Shamir, A., Tromer, E.: Cache attacks and countermeasures: the case of AES. In: Pointcheval, D. (ed.) CT-RSA 2006. LNCS, vol. 3860, pp. 1–20. Springer, Heidelberg (2006)

A Extended Pseudocode

```

path = (( 0),
        ( 16, 13),
        ( 32, 29, 25),
        ( 48, 45, 41, 37),
        ( 64, 61, 57, 53, 49),
        ( 80, 77, 73, 69, 65, 62),
        ( 96, 93, 89, 85, 81, 78, 74),
        (112,109,105,101, 97, 94, 90, 86),
        (128,125,121,117,113,110,106,102, 98),
        (144,141,137,133,129,126,122,118,114,111),
        (160,157,153,149,145,142,138,134,130,127,123),
        (173,169,165,161,158,154,150,146,143,139,135),
        (131,119,107, 95, 82, 70, 58, 46, 33, 21, 9),
        (124,115,103, 91, 79, 66, 54, 42, 30, 17, 5),
        (120,108, 99, 87, 75, 63, 50, 38, 26, 14, 1),
        (174,170,166,162,159,155,151,147,140,136,132))

class CandidateMatrix:
    Int count
    Byte m[176]

    def guess (Byte b):
        c = copy(self)
        c.count = count + 1
        c.m[path[count][0]] = b
        if count == 0:
            return c
        for i = 1 to len(path[count]):
            if defined(c.m[path[count][i]-16]):
                b4 = path[count][i-1]
                b3 = path[count][i]
                b0 = b4-16
                if inFirstWordOfRoundKey(b4):
                    c.m[b3] = unsbox(c.m[b4] XOR c.m[b0]
                                   XOR rcon[getRound(b4)][getBytePos(b4)])
            else:
                c.m[b3] = c.m[b0] XOR c.m[b4]

```

```

    else:
        b0 = path[count][i]
        b3 = path[count][i-1]
        b4 = b0 + 16
        if inFirstWordOfRoundKey(b4):
            c.m[b0] = sbox(c.m[b3]) XOR c.m[b4]
                XOR rcon[getRound(b4)][getBytePos(b4)]
        else:
            c.m[b0] = c.m[b3] XOR c.m[b4]
    if c.count == 16:
        c = deriveFullScheduleFromRound8(c)
    return c

def key ():
    return m[0:16]

class DecaySchedule:
    Byte decaySched[176]
    Byte gndEnc[176]
    def isCompatible (CandidateMatrix candidate):
        for i = 0 to 176:
            if defined(candidate.m[i]):
                if (candidate.m[i] XOR decaySched[i])
                    AND (candidate.m[i] XOR gndEnc[i]):
                    continue
            else:
                return FALSE
        return TRUE

def recoverKeyRec(CandidateMatrix c, DecaySchedule d):
    if (c.length()==16):
        return c.key()
    for i=0 to 255:
        if(d.isCompatible(c.guess(i))):
            key = recoverKeyRec(c.guess(i),d)
            if (key != NULL):
                return key
    return NULL

```

The pseudocode follows a Python-like syntax, but with some additional explicit typing and field declaration. Variables may hold their declared types or undefined values—a property checked by `defined()`.

The `path` variable encodes the guessing and inference order of Table 1 in the flat byte-level schedule view. The recursive exploration function, `recoverKeyRec()`, is the same as in Fig. 1.

The `CandidateMatrix` class is dominated by the `guess` method which guesses a value for the position determined by `path[count][0]` and infers the consequent bytes. The inference logic splits into two cases: when the unknown byte is in the first word, S^w_{i-4} , of the schedule generating equations (see (1)) or

in the middle word, S^w_{i-1} . The variable names, **b0**, **b3**, and **b4** reflect the relative position of their encapsulating words in the schedule; if **b0** comes from an arbitrary word in the first 10 rounds, then **b3** and **b4** come from the words three and four words ahead of **b0**, respectively. The inference sequence in **path** has been chosen to account for the necessary rotations in byte slices when solving the s-box version of the generating equations. Upon completion of the 16 guesses the eighth round becomes fully specified, implying the remainder of the schedule. The **key()** method simply returns the first 16 bytes of completed key schedule.

DecaySchedule holds the observed decayed data and ground state encoding for each byte. Its one method **isCompatible()** checks that each defined bit of the candidate schedule equals the decayed data or the ground state.