

Explicit Optimal Binary Pebbling for One-Way Hash Chain Reversal

Berry Schoenmakers

Dept of Mathematics & Computer Science
TU Eindhoven, The Netherlands
berry@win.tue.nl

Abstract. We present explicit optimal binary pebbling algorithms for reversing one-way hash chains. For a hash chain of length 2^k , the number of hashes performed in each output round does not exceed $\lceil k/2 \rceil$, whereas the number of hash values stored throughout is at most k . This is optimal for binary pebbling algorithms characterized by the property that the midpoint of the hash chain is computed just once and stored until it is output, and that this property applies recursively to both halves of the hash chain.

We introduce a framework for rigorous comparison of explicit binary pebbling algorithms, including simple speed-1 binary pebbles, Jakobsson’s binary speed-2 pebbles, and our optimal binary pebbles. Explicit schedules describe for each pebble exactly how many hashes need to be performed in each round. The optimal schedule turns out to be essentially unique and exhibits a nice recursive structure, which allows for fully optimized implementations that can readily be deployed. In particular, we develop the first *in-place* implementations with minimal storage overhead (essentially, storing only hash values), and fast implementations with minimal computational overhead. Moreover, we show that our approach is not limited to hash chains of length $n = 2^k$, but accommodates hash chains of arbitrary length $n \geq 1$, without incurring any overhead. Finally, we show how to run a cascade of pebbling algorithms along with a bootstrapping technique, facilitating sequential reversal of an unlimited number of hash chains growing in length up to a given bound.

1 Introduction

Originally introduced by Lamport to construct an identification scheme resisting eavesdropping attacks [Lam81, Hal94], one-way hash chains have become a truly fundamental primitive in cryptography.¹ The idea of Lamport’s asymmetric identification scheme is to let the prover generate a hash chain as a sequence of successive iterates of a one-way hash function applied to a random seed value, revealing only the last element of the hash chain to the verifier upon registration. Later, during successive runs of the identification protocol, the remaining elements of the hash chain are output by the prover in reverse order, one element on each run.

Due to the one-way property of the hash function, efficient reversal of a hash chain is non-trivial for long chains. In 2002, Jakobsson introduced a simple and efficient pebbling algorithm for reversal of one-way hash chains [Jak02], inspired

¹ Bitcoin’s block chain is probably the best-known example of a hash chain nowadays, each block containing a hash of the previous block in the chain. Unlike in our setting, however, block chains are costly to generate due to the “proof of work” requirement for the hash values.

by the pebbling algorithm of [IR01] for efficient key updates in a forward-secure digital signature scheme. Pebbling algorithms for one-way hash chain reversal strike a balance between storage requirements (measured as the number of hash values stored) and computational requirements (measured as the maximum number of hashes performed in any round). The performance constraint is that each next element of the reversed hash chain should be produced within a limited amount of time after producing the preceding element—without this performance constraint, the problem would indeed be easy, see Appendix A. For a hash chain of length $n = 2^k$, Jakobsson’s algorithm stores $O(\log n)$ hash values only and the number of hashes performed in each round is limited to $O(\log n)$ as well.

The problem of efficient hash chain reversal was extensively studied by Copper-smith and Jakobsson [CJ02]. They proved nearly optimal complexity for a binary pebbling algorithm storing at most $k + \lceil \log_2(k + 1) \rceil$ hash values and performing at most $\lfloor \frac{k}{2} \rfloor$ hashes per round. Later, it was observed by Yum et al. that a greedy implementation of Jakobsson’s original algorithm actually stores no more than k hash values, requiring no more than $\lceil \frac{k}{2} \rceil$ hashes per round [YSEL09].

In this paper we consider the class of *binary* pebbling algorithms, covering the best algorithms of [Jak02,CJ02,YSEL09] among others. A binary pebbling algorithm is characterized by the property that the *midpoint* of the hash chain is computed just once and stored until it is output; moreover, this property applies recursively to both halves of the hash chain. In particular, this means that after producing the last element of the hash chain as the first output, a binary pebbling algorithm stores the k elements at distances $2^i - 1$, for $1 \leq i \leq k$, from the end of the hash chain.

We introduce a simple yet general framework for rigorous analysis of efficient binary pebbling algorithms for hash chain reversal, and we completely resolve the case of binary pebbling by constructing an *explicit* optimal algorithm. The storage required by our optimal algorithm does not exceed the storage of k hash values and the number of hashes performed in any output round does not exceed $\lceil \frac{k}{2} \rceil$. This matches the performance of the greedy algorithm of [YSEL09], which is an optimal binary pebbling algorithm as well. However, we give an exact schedule for all hashes performed by the algorithm (rather than performing these hashes in a greedy fashion). We also believe that our approach is much more accessible than previous ones, leading to high quality algorithms that can readily be deployed.

Our optimal schedule is defined explicitly, both as a recursive definition and as a closed formula, specifying exactly how many hashes should be performed in a given round by each pebble. Furthermore, we will argue that the optimal schedule is essentially unique. Apart from the insightful mathematical structure thus uncovered, the explicit optimal schedule enables the development of fully optimized solutions for one-way hash chain reversal. We construct the first in-place (or, *in situ*) hash chain reversal algorithms which require essentially no storage beyond the hash values stored for the pebbles; at the same time, the computational overhead for each round is limited to a few basic operations only beyond the evaluation of the hash function. Finally, as another extreme type of solution, we

show how to minimize the computational overhead to an almost negligible amount of work, at the expense of increased storage requirements.

Concretely, for hash chains of length 2^{32} using a 128-bit one-way hash, our in-place algorithm only stores 516 bytes (32 hash values and one 32-bit counter) and performs, at most 16 hashes per round. Our results are therefore of particular importance in the context of lightweight cryptography. See, e.g., the references in [PCTS02,YSEL09,MSS13] for a glimpse of the extensive literature on hash chains, covering an extensive range of lightweight devices such as wireless sensors, RFID tags, and smart cards. Moreover, we note that our results are also interesting in the context of post-quantum cryptography, as the security of one-way hash chains is not affected dramatically by the potential of quantum computers.

2 One-Way Hash Chains

Throughout, we will use the following notation for (finite) sequences. We write $A = \{a_i\}_{i=1}^n = \{a_1, \dots, a_n\}$ for a sequence A of length n , $n \geq 0$, with $\{\}$ denoting the empty sequence. We use $|A| = n$ to denote the length of A , and $\#A = \sum_{i=1}^n a_i$ to denote the weight of A . We write $A \parallel B$ for the concatenation of sequences A and B , and $A+B$ for element-wise addition of sequences A and B of equal length, where $+$ takes precedence over \parallel . Constant sequences are denoted by $c = c^{*n} = \{c\}_{i=1}^n$, suppressing the length n when it is understood from context; e.g., $A + c$ denotes the sequence obtained by adding c to all elements of A .

Let f be a cryptographic hash function. The length- 2^k (one-way) hash chain $f_k^*(x)$ for a given seed value x is defined as the following sequence:

$$f_k^*(x) = \{f^i(x)\}_{i=0}^{2^k-1}.$$

For authentication mechanisms based on hash chains, we need an efficient algorithm for producing the sequence $f_k^*(x)$ *in reverse*. The problem arises from the fact that computation of f in the forward direction is easy, while it is *intractable* in the reverse direction. So, given x it is easy to compute $y = f(x)$, but given y it is hard to compute any x at all such that $y = f(x)$. For long hash chains the straightforward solutions of either (i) storing $f_k^*(x)$ and reading it out in reverse or (ii) computing each element of $f_k^*(x)$ from scratch starting from x are clearly too inefficient.

3 Binary Pebbling

We introduce a framework that captures the essence of binary pebbling algorithms as follows. We define a pebble as an algorithm proceeding in a certain number of rounds, where the initial rounds are used to compute the hash chain in the forward direction given the seed value x , and the hash chain is output in reverse in the remaining rounds, one element at a time.

For $k \geq 0$, we define pebble $P_k(x)$ below as an algorithm that runs for $2^{k+1} - 1$ rounds in total, and outputs $f_k^*(x)$ in reverse in its last 2^k rounds. It is essential




Fig. 1. Binary pebble $P_k(x)$, where $y_i = f^{2^k-2^i}(x)$ for $i = k, \dots, 0$.

that we include the initial $2^k - 1$ rounds in which no outputs are produced as an integral part of pebble $P_k(x)$, as this allows us to define and analyze binary pebbling in a fully recursive manner. In fact, in terms of a given **schedule** $T_k = \{t_r\}_{r=1}^{2^k-1}$ with $\#T_k = 2^k - 1$, a binary pebble $P_k(x)$ is completely specified by the following recursive definition, see also Figures 1–2:

Rounds $[1, 2^k)$: Store $y_i = f^{2^k-2^i}(x)$ for $i = k, \dots, 0$ using t_r hashes in round r .

Round 2^k : Output y_0 .

Rounds $(2^k, 2^{k+1})$: Run $P_{i-1}(y_i)$ in parallel for $i = 1, \dots, k$.

We will refer to rounds $[1, 2^k)$ as the **initial stage** of pebble P_k and to rounds $[2^k, 2^{k+1})$ as its **output stage**. Running pebbles in parallel means that pebbles

take turns to execute for one round each, where the order in which this happens within a round is irrelevant.

The behavior of pebbles P_0 and P_1 is fixed since $T_0 = \{\}$ and $T_1 = \{1\}$, respectively. Pebble $P_0(x)$ runs for one round only, in which $y_0 = x$ is output, using no hashes at all. Similarly, $P_1(x)$ runs for three rounds, performing one hash in its first round to compute $y_1 = x$ and $y_0 = f(x)$, outputting $f(x)$ in its second round, and then running $P_0(y_1)$ in the third round, which will output x . More generally, the following theorem shows that correct behavior follows for any pebble P_k independent of the particular schedule T_k , and furthermore that the total number of hashes performed by P_k is fixed as well.

Theorem 1. *Pebble $P_k(x)$ produces $f_k^*(x)$ in reverse in its output stage, performing $k2^{k-1}$ hashes ($2^k - 1$ in its initial stage and $(k-2)2^{k-1} + 1$ in its output stage).*

Proof The proof is by induction on k . For $k = 0$, we have that $P_0(x)$ outputs $f_0^*(x) = x$ in its one and only round, using 0 hashes.

For $k \geq 1$, we see that $P_k(x)$ first outputs $y_0 = f^{2^k-1}(x)$ in round 2^k , which is the last element of $f_k^*(x)$. Next, pebbles $P_{i-1}(y_i)$ run in parallel for $i = 1, \dots, k$. The induction hypothesis yields that each $P_{i-1}(y_i)$ produces $f_{i-1}^*(f^{2^k-2^i}(x))$ in reverse in its last 2^{i-1} out of $2^i - 1$ rounds. Hence, in round $2^k + 1$, $P_0(y_1)$ outputs $\{y_1\} = f_0^*(f^{2^k-2}(x))$. In the next two rounds, $P_1(y_2)$ outputs $f_1^*(f^{2^k-4}(x))$ in reverse. And so on until finally $P_{k-1}(y_k)$ outputs $f_{k-1}^*(f^{2^k-1}(x))$ in reverse in the last 2^{k-1} rounds of $P_k(x)$. The total number of hashes performed by P_k is $2^k - 1 + \sum_{i=1}^k (i-1)2^{i-2} = k2^{k-1}$, using that P_{i-1} performs $(i-1)2^{i-2}$ hashes per the induction hypothesis. \square

Schedule T_k specifies the number of hashes for the initial stage of P_k . To analyze the **work** done by P_k in its output stage, we introduce sequence W_k of length $2^k - 1$ to denote the number of hashes performed by P_k in each of its last $2^k - 1$ rounds— noting that by definition no hashes are performed by P_k in round 2^k . The following recurrence relation for W_k will be used throughout our analysis.

Lemma 1. $W_0 = \{\}$, $W_k = T_{k-1} + W_{k-1} \parallel 0 \parallel W_{k-1}$.

Proof Pebble P_0 runs for 1 round only, so $W_0 = \{\}$. For $k \geq 1$, we see that in the last $2^k - 1$ rounds of pebble P_k , a pebble P_{k-1} runs in parallel to pebbles P_i for $i = 0, \dots, k-2$. In these rounds, pebble P_{k-1} performs $T_{k-1} \parallel 0 \parallel W_{k-1}$ hashes by definition, whereas pebbles P_i for $i = 0, \dots, k-2$ perform $W_{k-1} \parallel 0^{*2^{k-1}}$ hashes in total, as this matches the number of hashes for a pebble P_{k-1} in its last $2^{k-1} - 1$ rounds (consider the output stage in Figure 1). Hence, in total $W_k = T_{k-1} + W_{k-1} \parallel 0 \parallel W_{k-1}$ hashes. \square

We have the following lower bound for $\max(W_k)$, the maximum number of hashes performed by P_k in any round of its output stage. Interestingly, this lower bound holds for *any* schedule T_k . In Section 5 we will present an optimal schedule achieving the lower bound.

Theorem 2. $\max(W_k) \geq \lceil k/2 \rceil$, for $k \geq 2$.

Proof Let $k \geq 2$ and consider the average number of hashes per round during the first half of the output stage. From Theorem 1, Lemma 1, and $|T_{k-1}| = |W_{k-1}| = 2^{k-1} - 1$, we have

$$\max(W_k) \geq \frac{\#T_{k-1} + \#W_{k-1}}{|T_{k-1} + W_{k-1}|} = \frac{(k-1)2^{k-2}}{2^{k-1} - 1} > \frac{k-1}{2}.$$

Hence, $\max(W_k) \geq \lceil k/2 \rceil$. □

To analyze the **storage** needed by P_k we will count the number of hash values stored by P_k at the start of each round. We introduce sequence $S_k = \{s_r\}_{r=1}^{2^{k+1}-1}$ to denote the total storage used by P_k in each round. For instance, $s_1 = 1$ as P_k only stores x at the start, and $s_{2^k} = k + 1$ as P_k stores y_0, \dots, y_k at the start of round 2^k independent of schedule T_k .

4 Speed-1 and Speed-2 Binary Pebbles

In this section we analyze the performance of speed-1 pebbles and speed-2 pebbles. We use speed-1 pebbles to demonstrate our framework, whereas the analysis of speed-2 pebbles, which correspond to Jakobsson's original algorithm [Jak02], will be used in the analysis of our optimal pebbles in the next section.

Speed-1 pebbles are defined by setting $T_k = 1 * 2^{k-1}$, hence one hash evaluation in each initial round of P_k . To define speed-2 pebbles we set $T_0 = \{\}$ and $T_k = 0 * 2^{k-1} - 1 \parallel 2 * 2^{k-1} - 1 \parallel 1$ for $k \geq 1$, hence a speed-2 pebble is idle in the first part of the initial stage and then hashes twice in each round until the end of the initial stage. As can be seen from Theorem 4 below, the storage requirements are reduced by a factor of 2 for speed-2 pebbles over speed-1 pebbles.

Theorem 3. *Both speed-1 and speed-2 pebbles P_k use up to $\max(W_k) = k - 1$ hashes in any output round, for $k \geq 1$.*

Proof For a speed-1 pebble, Lemma 1 implies $\max(W_k) = k - 1$ for $k \geq 1$, as all elements of T_{k-1} are equal to 1.

For a speed-2 pebble we prove by induction on k that $\max(W_k) = k - 1$. This clearly holds for $k = 1, 2$. For $k \geq 3$, we have, using Lemma 1,

$$\begin{aligned} T_{k-1} &= 0 * 2^{k-2} - 1 \parallel 2 * 2^{k-2} - 1 \parallel 1 \\ W_{k-1} &= 0 \parallel T_{k-2} + W_{k-2} \parallel 0 \parallel W_{k-2}. \end{aligned}$$

Therefore,

$$\max(W_k) = \max(T_{k-1} + W_{k-1}) = \max(W_{k-1}, 2 + W_{k-2}),$$

noting that the last element of $W_{k-2} = 0$. Applying the induction hypothesis twice, we conclude $\max(W_k) = \max(k - 2, k - 1) = k - 1$. □

Lemma 2.

$$\begin{aligned} S_0 &= \{1\}, \\ S_k &= (1^{*2^k} \parallel S_{k-1}) + (0 \parallel 1^{*2^{k-1}-1} \parallel S_{k-1} \parallel 0^{*2^{k-1}}), \text{ for a speed-1 } P_k, \\ S_k &= (1^{*2^k} \parallel S_{k-1}) + (0^{*2^{k-1}} \parallel S_{k-1} \parallel 0^{*2^{k-1}}), \text{ for a speed-2 } P_k. \end{aligned}$$

Proof Pebble $P_0(x)$ only needs to store x during its one and only round, therefore $S_0 = \{1\}$. For $k \geq 1$, any pebble $P_k(x)$ also needs to store x throughout all of its rounds, where pebble $P_{k-1}(y_k) = P_{k-1}(x)$ takes over the storage of x during the output stage. This accounts for the term $1^{*2^k} \parallel S_{k-1}$.

In addition, a speed-1 pebble needs to store a hash value from round 2 until it reaches y_{k-1} in round 2^{k-1} . From thereon, the total additional storage corresponds to running a speed-1 $P_{k-1}(y_{k-1})$ pebble. This accounts for the term $0 \parallel 1^{*2^{k-1}-1} \parallel S_{k-1} \parallel 0^{*2^{k-1}}$.

A speed-2 pebble needs no additional storage during its first 2^{k-1} rounds. Then it needs to store an additional hash value from round $2^{k-1} + 1$ on. By taking $0^{*2^{k-1}} \parallel S_{k-1} \parallel 0^{*2^{k-1}}$ as additional term, we account for both the additional hash value stored by a speed-2 pebble during rounds $(2^{k-1}, 2^{k-1} + 2^{k-2}]$ and the storage corresponding to a speed-2 $P_{k-1}(y_{k-1})$ pebble, running from round $2^{k-1} + 1$. \square

Theorem 4. *A speed-1 pebble P_k uses up to $\max(S_k) = \max(k+1, 2k-2)$ storage, and a speed-2 pebble P_k uses up to $\max(S_k) = k+1$ storage.*

Proof Using that $s_{2^k} = k+1$, we write $S_k = A_k \parallel k+1 \parallel B_k$, where $|A_k| = |B_k| = 2^k - 1$.

For a speed-1 pebble P_k , it can easily be checked that $\max(S_k) = \max(k+1, 2k-2)$ holds for $k = 0, 1$. To prove this for $k \geq 2$, we note that it suffices to show $\max(A_k, B_k) = 2k-2$, as $\max(S_k) = \max(A_k, k+1, B_k)$. Lemma 2 implies

$$\begin{aligned} A_k &= 1 \parallel 2^{*2^{k-1}-1} \parallel 1 + A_{k-1} \\ B_k &= A_{k-1} + B_{k-1} \parallel k \parallel B_{k-1}, \end{aligned}$$

so we have that $\max(A_k, B_k) = \max(A_{k-1} + B_{k-1}, k) = \max(2k-2, k) = 2k-2$ follows if we can show $\max(A_k + B_k) = 2k$, for $k \geq 1$. We prove the latter by induction on k . For $k = 1$, it is clearly true as $A_1 = B_1 = \{1\}$. For $k \geq 2$, we see that $\max(A_k + B_k) = \max(2 + A_{k-1} + B_{k-1}, k+2) = \max(2k, k+2) = 2k$ follows from the induction hypothesis, also using that the first element of $A_{k-1} + B_{k-1}$ is equal to k .

For a speed-2 pebble P_k , we note that $\max(S_k) = k+1$ follows from the fact that $A_k + B_k = k+1$, which we show by induction on k . For $k = 0$, $A_k + B_k = k+1$ is vacuously true, as A_0, B_0 are empty sequences. For $k \geq 1$, we see from Lemma 2 that

$$\begin{aligned} A_k &= 1^{*2^{k-1}} \parallel 1 + A_{k-1} \\ B_k &= A_{k-1} + B_{k-1} \parallel k \parallel B_{k-1}. \end{aligned}$$

Thus, from the induction hypothesis we have $A_{k-1} + B_{k-1} = k$, hence $A_k + B_k = k+1$. \square

5 Optimal Binary Pebbles

In this section, we will reduce the maximum number of hashes per round from $k - 1$ for a speed-2 pebble P_k to $\lceil k/2 \rceil$ for an optimal pebble P_k , without increasing the storage requirements. We do so by letting our optimal pebbles P_k be idle for the first $2^{k-1} - 1$ rounds, just as speed-2 pebbles do. During rounds $[2^{k-1}, 2^k)$, an optimal pebble will work at varying speeds, roughly as follows: the average speeds in each quarter are 2, 1, 2, and 3 hashes per round, respectively. To streamline the presentation, we will at first allow “ $\frac{1}{2}$ hashes” in the definition of our optimal schedule. At the end of this section, we will show how to round the schedule to integer values without affecting optimality.

We define the optimal schedule T_k as follows:

$$T_0 = \{\}, \quad T_k = 0^{*2^{k-1}-1} \parallel U_k \parallel V_k,$$

where

$$\begin{aligned} U_1 &= \{1\}, & U_k &= \frac{1}{2} + U_{k-1} \parallel 1^{*\lceil 2^{k-3} \rceil}, \\ V_1 &= \{\}, & V_k &= \frac{1}{2} + U_{k-1} \parallel \frac{1}{2} + V_{k-1}. \end{aligned}$$

For example, $T_1 = \{1\}$, $T_2 = \{0, \frac{3}{2}, \frac{3}{2}\}$, and $T_3 = \{0, 0, 0, 2, 1, 2, 2\}$.

Optimality is proved in the next two theorems. Subsequently, we will argue that optimal schedule T_k is essentially unique.

Theorem 5. *An optimal pebble P_k uses up to $\max(W_k) = k/2$ hashes in any output round, for $k \geq 2$.*

Proof We use Lemma 1 without explicitly referring to it.

Since $\max(W_k) = \max(T_{k-1} + W_{k-1})$, we obtain $\max(W_k) = k/2$, if we prove by induction on k that

$$T_k + W_k = T_{k-1} + W_{k-1} \parallel \frac{k+1}{2}^{*2^{k-1}}.$$

This property clearly holds for $k = 1, 2$. For $k \geq 3$, the definition of T_k implies that the property is in fact equivalent to

$$(U_k \parallel V_k) + (0 \parallel W_{k-1}) = \frac{k+1}{2}^{*2^{k-1}}. \quad (1)$$

From the definition of U_k, V_k and the induction hypothesis for $T_{k-2} + W_{k-2}$ we obtain

$$\begin{aligned} U_k \parallel V_k &= \frac{1}{2} + U_{k-1} \parallel 1^{*2^{k-3}} \parallel \frac{1}{2} + (U_{k-1} \parallel V_{k-1}), \\ 0 \parallel W_{k-1} &= 0 \parallel T_{k-3} + W_{k-3} \parallel \frac{k-1}{2}^{*2^{k-3}} \parallel 0 \parallel W_{k-2}. \end{aligned}$$

Since $0 \parallel T_{k-3} + W_{k-3}$ is equal to the first half of $0 \parallel W_{k-2}$, we get from the induction hypothesis that indeed all elements of $(U_k \parallel V_k) + (0 \parallel W_{k-1})$ are equal to $\frac{k+1}{2}$. \square

Let $\text{len}(x) = \lceil \log_2(x+1) \rceil$ denote the bit length of nonnegative integer x . The next two lemmas give closed formulas for the optimal schedule T_k and its partial sums. Lemma 4 will be used to prove Theorem 6, but these formulas also provide the basis for our efficient in-place implementation of optimal binary pebbling.

Lemma 3. For optimal schedule $T_k = \{t_r\}_{r=1}^{2^k-1}$, we have for $2^{k-1} \leq r < 2^k$:

$$t_r = \frac{1}{2} \left(k + 1 - \text{len}((2r) \bmod 2^{\text{len}(2^k-r)}) \right).$$

Proof The proof is by induction on k . For $0 \leq k \leq 2$, the formula is easily checked. For $k \geq 3$, we distinguish two cases.

Case $2^{k-1} \leq r < 2^{k-1} + 2^{k-2}$. We first note that $(2r) \bmod 2^{\text{len}(2^k-r)} = 2r - 2^k$. If $r \geq 2^{k-1} + 2^{k-3}$, we have $t_r = 1$ by definition and we see the formula for t_r yields 1 as well as $\text{len}(2r - 2^k) = k - 1$. Otherwise $r < 2^{k-1} + 2^{k-3}$, hence we have $t_r = t_{r+2^{k-2}}$. So, this case reduces to the case below by noting that also $(2(r + 2^{k-2})) \bmod 2^{\text{len}(2^k-(r+2^{k-2}))} = 2r - 2^k$.

Case $2^{k-1} + 2^{k-2} \leq r < 2^k$. From the definition of the optimal schedule we see that in this case $t_r = \frac{1}{2} + t'_{r-2^{k-1}}$, where $T_{k-1} = \{t'_z\}_{z=1}^{2^{k-1}-1}$. From the induction hypothesis we get:

$$t'_{r-2^{k-1}} = \frac{1}{2} \left(k - \text{len}((2(r - 2^{k-1})) \bmod 2^{\text{len}(2^{k-1}-(r-2^{k-1}))}) \right).$$

Rewriting this formula for $t'_{r-2^{k-1}}$ we obtain

$$t_r = \frac{1}{2} + \frac{1}{2} \left(k - \text{len}((2r - 2^k) \bmod 2^{\text{len}(2^k-r)}) \right).$$

Noting that $\text{len}(2^k - r) \leq k$, we see that the formula holds for t_r as well. \square

Lemma 4. For optimal schedule $T_k = \{t_r\}_{r=1}^{2^k-1}$, we have for $0 \leq j \leq 2^{k-1}$:

$$\sum_{r=2^{k-j}}^{2^k-1} t_r = \frac{1}{2} \left(j(k-m) + (m+3-l)2^l - 2^m \right) - 1,$$

where $l = \text{len}(j)$ and $m = \text{len}(2^l - j)$.

Proof The proof is by induction on j . For $j = 0$, both sides are equal to 0.

For $1 \leq j \leq 2^{k-1}$, Lemma 3 implies that

$$t_{2^{k-j}} = \frac{1}{2} \left(k + 1 - \text{len}((-2j) \bmod 2^l) \right).$$

Combined with the induction hypothesis for $j - 1$ we obtain

$$\sum_{r=2^{k-j}}^{2^k-1} t_r = \frac{1}{2} \left(j(k-m') + m' + 1 - \text{len}((-2j) \bmod 2^l) + (m'+3-l')2^{l'} - 2^{m'} \right) - 1,$$

where $l' = \text{len}(j - 1)$ and $m' = \text{len}(2^{l'} - j + 1)$. We distinguish two cases.

Case $l' = l - 1$. This means that $j = 2^{l-1}$, and hence $m = l$ and $m' = 1$. We are done as both sides are equal to $\frac{1}{2}j(k+4-l) - 1$.

Case $l' = l$. This means that $2^{l-1} < j < 2^l$, hence $0 < 2^{l+1} - 2j < 2^l$. This implies $\text{len}((-2j) \bmod 2^l) = m + 1$, so we see that both sides are equal if $m = m'$. If $m' = m + 1$, we see that $2^l - j = 2^m - 1$ and that therefore both sides are equal as well. \square

Theorem 6. *An optimal pebble P_k uses up to $\max(S_k) = k + 1$ storage.*

Proof We prove that the storage requirements of an optimal pebble do not exceed the storage requirements of a speed-2 pebble, hence that $\max(S_k) = k + 1$ for an optimal pebble as well.

Consider the rounds in which a speed-2 pebble and an optimal pebble store the values $y_i = f^{2^k - 2^i}(x)$ for $i = k, \dots, 1$. We claim that an optimal pebble will never store y_i before a speed-2 pebble does. Clearly, a speed-2 pebble stores y_i in round $2^k - 2^{i-1}$ for $i = k, \dots, 1$. However, in round $2^k - 2^{i-1}$ an optimal pebble still has to compute at least as many hashes as a speed-2 pebble needs to reach y_0 :

$$\sum_{r=2^k - 2^{i-1}}^{2^k - 1} t_r = 2^{i-2}(k + 4 - i) - 1 \geq 2^i - 1,$$

using Lemma 4 for $j = 2^{i-1}$. □

The above analysis implies that the optimal schedule T_k is essentially unique. That is, if we require that the bounds stated in Theorems 5 and 6 must be met, optimal schedule T_k is basically the only solution within our framework for binary pebbling. Clearly, like for any schedule in our framework, we have $T_0 = \{\}$ and $T_1 = \{1\}$. For $k \geq 2$, we first note that optimal schedule T_k must start with $2^{k-1} - 1$ zeros in order to meet the bound of Theorem 6. More precisely, by writing $S_k = A_k \parallel k+1 \parallel B_k$ with $|A_k| = |B_k| = 2^k - 1$, as we also did in the proof of Theorem 4, and noting that $\max(B_k) = k$, we observe that this bound for B_k is met only if optimal pebble P_k does not perform any hashes in its first $2^{k-1} - 1$ rounds. Next, by reconsidering the proof of Theorem 5 we actually see that the number of hashes performed by an optimal pebble P_k in each of the last 2^{k-1} rounds of its initial stage is determined uniquely as well. A simple calculation shows that Eq. (1) necessarily holds for any optimal schedule: $\#U_k + \#V_k + \#W_{k-1} = (k + 1)2^{k-2}$ hashes to be performed in 2^{k-1} rounds requires exactly $(k + 1)/2$ hashes to be performed in each of these rounds. Hence, using induction on k , we see that U_k and V_k are determined uniquely by Eq. (1).

As a final step we will round the optimal schedule T_k to integer values, without affecting optimality. For example, we round $T_2 = \{0, \frac{3}{2}, \frac{3}{2}\}$ to $\{0, 1, 2\}$ or to $\{0, 2, 1\}$. In general, we make sure that if an element is rounded up then its neighbors are rounded down, and vice versa. The rounding also depends on the parity of k to alternate between rounding up and rounding down. Hence, we define the rounded optimal schedule by:

$$t_r = \left\lfloor \frac{1}{2} \left((k + r) \bmod 2 + k + 1 - \text{len}((2r) \bmod 2^{\text{len}(2^k - r)}) \right) \right\rfloor, \quad (2)$$

for $2^{k-1} \leq r < 2^k$. Accordingly, we see that optimal pebble P_k will use up to $\max(W_k) = \lceil k/2 \rceil$ hashes in any output round, matching the lower bound of Theorem 2.

6 Optimized Implementations

A hash chain is deployed as follows as part of an authentication mechanism like Lamport’s asymmetric identification scheme. Given a random seed value x , the initial stage of any type of binary pebble $P_k(x)$ is simply performed by iterating the hash function f and storing the values $y_i = f^{2^k-2^i}(x)$ for $i = k, \dots, 0$. The value of y_0 is then output, e.g., as part of the registration protocol for Lamport’s scheme. The other hash values y_1, \dots, y_k are stored for the remainder of the output stage, e.g., for use in later runs of Lamport’s identification protocol.

The initial stage is preferably executed inside the secure device that will later use the hash chain for identification. However, for lightweight devices such as smart cards, RFID tags, sensors, etc., the initial stage will typically be run on a more powerful device, after which the hash values y_1, \dots, y_k will be inserted in the lightweight device and the hash value y_0 can be used for registration.

To implement the output stage of a pebble P_k one needs to handle potentially many pebbles all running in parallel. The pseudocode in [Jak02,CJ02,YSEL09] suggests rather elaborate techniques for keeping track of the (state of) pebbles. On the contrary, we will show how to minimize storage and computational overhead by exploiting specific properties of Jakobsson’s speed-2 pebbles and our optimal pebbles. In particular, we present in-place hash chain reversal algorithms, where the entire state of these algorithms (apart from the hash values) is represented between rounds by a single k -bit counter *only*.

We introduce the following terminology to describe the state of a pebble P_k . This terminology applies to both speed-2 pebbles and optimal pebbles. Pebble P_k is said to be **idle** if it is in rounds $[1, 2^{k-1})$, **hashing** if it is in rounds $[2^{k-1}, 2^k]$, and **redundant** if it is in rounds $(2^k, 2^{k+1})$. An idle pebble performs no hashes at all, while a hashing pebble will perform at least one hash per round, except for round 2^k in which P_k outputs its y_0 value. The work for a redundant pebble P_k is taken over by its child pebbles P_0, \dots, P_{k-1} during its last $2^k - 1$ output rounds.

The following theorem provides the basis for our in-place algorithms by showing precisely how the state of all pebbles running in parallel during the output stage of a pebble P_k can be determined from the round number. Let $x_i \in \{0, 1\}$ denote the i th bit of nonnegative integer x , $0 \leq i < \text{len}(x)$.

Theorem 7. *For a speed-2 pebble or optimal pebble P_k in output round $2^{k+1} - c$, $1 \leq c \leq 2^k$, we have for every i , $0 \leq i \leq k$, exactly one non-redundant pebble P_i present if and only if bit $c_i = 1$, and if present, P_i is in round $2^i - (c \bmod 2^i)$.*

Proof The proof is by induction on c . For $c = 2^k$, only $c_k = 1$, which corresponds to pebble P_k being the only pebble around. Also, P_k is in its 2^k th round.

For $1 \leq c < 2^k$, write $c' = c+1$ and let $i' \geq 0$ be maximal such that $c' \bmod 2^{i'} = 0$. Hence $c'_{i'} = 1$. By the induction hypothesis for c' , pebble $P_{i'}$ is in its first output round $2^{i'}$. So, in the next round $P_{i'}$ becomes redundant, and is replaced by child pebbles $P_{i'-1}, \dots, P_0$ which will all be in their first round. As $c = c' - 1$, this corresponds to the fact that $c_{i'} = 0$ and $c_{i'-1} = \dots = c_0 = 1$, also noting that $2^i - (c \bmod 2^i) = 1$ for $i = i' - 1, \dots, 0$.

For $i > i'$, we have $c_i = c'_i$. All non-redundant pebbles in round $2^{k+1} - c'$ remain so in round $2^{k+1} - c$, and for these pebbles the round number becomes $2^i - (c' \bmod 2^i) + 1 = 2^i - (c \bmod 2^i)$, as required. \square

As a corollary, we see that a non-redundant pebble P_i is hashing precisely when $c_{i-1} = 0$, and P_i is idle otherwise, since for $c_i = 1$ we have that $c_{i-1} = 0$ if and only if $2^i - (c \bmod 2^i) \geq 2^{i-1}$. This holds also for $i = 0$ if we put $c_{-1} = 0$.

6.1 In-Place Speed-2 Binary Pebbles

We present an in-place implementation of a speed-2 pebble P_k for which the overall storage is limited to the space for k hash values and one k -bit counter c . As explained above, we will assume that hash values y_1, \dots, y_k are given as input and that y_0 has been output already. Thus, P_k has exactly $2^k - 1$ output rounds remaining. We use c to count down the output rounds.

The basis for our in-place algorithm is given by the next theorem.

Theorem 8. *For a speed-2 pebble P_k in output round $2^{k+1} - c$, $1 \leq c \leq 2^k$, each non-redundant pebble P_i present stores $e + 1$ hash values, where e is maximal such that $c_{i-1} = \dots = c_{i-e} = 0$ with $0 \leq e \leq i$.*

Proof From Theorem 7 it follows that non-redundant pebble P_i is in round $r = 2^i - (c \bmod 2^i)$. Since $0 \leq e \leq i$ is maximal such that $c_{i-1} = \dots = c_{i-e} = 0$, we have that $2^i - 2^{i-e} < r \leq 2^i - 2^{i-e-1}$. This implies that P_i stores $e + 1$ hash values in round r , as Lemma 2 says that for a speed-2 pebble P_i the storage requirements throughout its first 2^i rounds are given by sequence D_i , where $D_0 = \{1\}$ and $D_i = 1^{*2^{i-1}} \parallel 1 + D_{i-1}$. \square

Theorem 8 suggests an elegant approach to store the hash values for a speed-2 pebble P_k throughout its output stage. We use a single array z of length k to store all hash values as follows. Initially, $z[0] = y_1, \dots, z[k-1] = y_k$, and counter $c = 2^k - 1$. This corresponds to pebble P_k being at the start of its output stage, where it starts to run pebbles $P_i(y_{i+1})$ in parallel, for $i = 0, \dots, k-1$, each of these (non-redundant) pebbles P_i storing exactly one hash value in array z . In general, in output round $2^{k+1} - c$ of P_k , we let each non-redundant pebble P_i store its hash values in array z in segment $z[i-e, i]$ (corresponding to $c_{i-e} = 0, \dots, c_{i-1} = 0$ and $c_i = 1$). As a result, the non-redundant pebbles jointly occupy consecutive segments of array z , storing exactly $\text{len}(c)$ hash values in total.

Algorithm 1 describes precisely what pebble P_k does in round r , $2^k < r < 2^{k+1}$. Note that we set $c = 2^{k+1} - r$ at the start of round r . Based on Theorem 7, we process the bits of c as follows, using operations $\text{pop}_0(c)$ and $\text{pop}_1(c)$ to count and remove all trailing 0s and 1s from c , respectively.

Let $i' \geq 0$ be maximal such that $c \bmod 2^{i'} = 0$. Hence $c_{i'} = 1$. From Theorem 7, we see that pebble $P_{i'}$ is in its first output round $2^{i'}$, hence $P_{i'}$ becomes redundant in the next round, and each of its children will take over one hash value. The hash values $y_0, \dots, y_{i'}$ computed by $P_{i'}$ in its initial stage are stored in $z[0], \dots, z[i']$. So,

Algorithm 1 In-place speed-2 binary pebble $P_k(x)$.

Initially: array $z[0, k]$ where $z[i-1] = f^{2^k - 2^i}(x)$ for $i = 1, \dots, k$.

Round r , $2^k < r < 2^{k+1}$:

```
1: output  $z[0]$ 
2:  $c \leftarrow 2^{k+1} - r$ 
3:  $i \leftarrow \text{pop}_0(c)$ 
4:  $z[0, i] \leftarrow z[1, i]$ 
5:  $i \leftarrow i + 1$ ;  $c \leftarrow \lfloor c/2 \rfloor$ 
6:  $q \leftarrow i - 1$ 
7: while  $c \neq 0$  do
8:    $z[q] \leftarrow f(z[i])$ 
9:   if  $q \neq 0$  then  $z[q] \leftarrow f(z[q])$ 
10:   $i \leftarrow i + \text{pop}_0(c) + \text{pop}_1(c)$ 
11:   $q \leftarrow i$ 
```

we output $z[0] = y_0$ for $P_{i'}$ and move $y_1, \dots, y_{i'}$ to entries $z[0], \dots, z[i' - 1]$. This makes entry $z[i']$ available. We distinguish two cases.

Case $\text{len}(c - 1) = \text{len}(c) - 1$. In this case no new hash values need to be stored, and $z[i']$ will be unused from this round on.

Case $\text{len}(c - 1) = \text{len}(c)$. Let $i'' \geq i' + 1$ be maximal such that $c \bmod 2^{i''} = 2^{i'}$. Hence $c_{i''} = 1$. We claim that pebble $P_{i''}$ is the unique pebble that needs to store an additional hash value. Pebble $P_{i''}$ is in round $2^{i''} - (c \bmod 2^{i''}) = 2^{i''} - 2^{i'}$, so it is $2^{i'}$ rounds from the end of its initial stage. We store its additional hash value in $z[i']$.

This explains Algorithm 1. In the first iteration of the loop in lines 7–11, we have that $q = i'$ holds at line 8. Each hashing pebble performs two hashes, except when a pebble is at the end of its initial stage (corresponding to $q = 0$). Essentially no processing is done for idle pebbles, due to the use of operation $\text{pop}_1(c)$ in line 10.

6.2 In-Place Optimal Binary Pebbles

In this section we turn the algorithm for speed-2 pebbles into one for optimal pebbles by making three major changes. See Algorithm 2.

First, we make sure that the number of hashes performed by each hashing pebble P_i is in accordance with Eq. (2). The actual hashing by P_i is done in the loop in lines 14–18. Theorem 7 states that the round number for P_i is given by $2^i - (c \bmod 2^i)$, hence by $2^i - j$ if we set $j = (-r) \bmod 2^i$ in line 10. By ensuring that $l = \text{len}(j)$ and $m = \text{len}(2^l - j)$ holds as well, we have that the number of hashes as specified by Eq. (2) can be computed as $\lfloor (p + i + 1 - s)/2 \rfloor$, where $p = (i + j) \bmod 2$ and $s = (m + 1) \bmod (l + 1)$ (actually using that $\text{len}((2^l - 2j) \bmod 2^l) = \text{len}(2^{l+1} - 2j) \bmod (l + 1)$ holds for $j \geq 1$).

Second, we make sure that each hashing pebble P_i will store the correct hash values for y_i, \dots, y_0 . To this end, note that Lemma 4 tells precisely how many hashes pebble P_i still needs to compute at the start of round j . Thus we set h to this value (plus one) in line 12, and test in line 16 if the current hash value must be stored (that is, whether h is an integral power of 2).

Algorithm 2 In-place optimal binary pebble $P_k(x)$.

Initially: array $z[0, k]$ where $z[i-1] = f^{2^k - 2^i}(x)$ for $i = 1, \dots, k$.

Round r , $2^k < r < 2^{k+1}$:

```
1: output  $z[0]$ 
2:  $c \leftarrow 2^{k+1} - r$ 
3:  $i \leftarrow \text{pop}_0(c)$ 
4:  $z[0, i] \leftarrow z[1, i]$ 
5:  $i \leftarrow i + 1$ ;  $c \leftarrow \lfloor c/2 \rfloor$ 
6:  $m \leftarrow i$ ;  $s \leftarrow 0$ 
7: while  $c \neq 0$  do
8:    $l \leftarrow i$ 
9:    $i \leftarrow i + \text{pop}_0(c)$ 
10:   $j \leftarrow (-r) \bmod 2^i$ 
11:   $p \leftarrow (i + j) \bmod 2$ 
12:   $h \leftarrow \lfloor (p + j(i - m) + (m + 3 - l)2^l - 2^m)/2 \rfloor$ 
13:   $q \leftarrow \text{len}(h) - 1$ 
14:  for  $d \leftarrow 1$  to  $\lfloor (p + i + 1 - s)/2 \rfloor$  do
15:     $y \leftarrow z[q]$ 
16:    if  $h = 2^q$  then  $q \leftarrow q - 1$ 
17:     $z[q] \leftarrow f(y)$ 
18:     $h \leftarrow h - 1$ 
19:     $m \leftarrow i$ ;  $s \leftarrow m + 1$ 
20:     $i \leftarrow i + \text{pop}_1(c)$ 
```

Finally, we make sure that hashing pebble P_i will use the correct entries of array z . Since h records the number of hashes that P_i still needs to compute (plus one), it follows that the current hash value for P_i is stored in entry $z[q]$, where $q = \text{len}(h) - 1$. Hence, we set q to this value in line 13.

This explains the design of Algorithm 2. Note that only one bit length computation is used per hashing pebble (cf. line 13).

6.3 Optimal Binary Pebbles with Minimal Computational Overhead

Even though the computational overhead for our in-place implementation is small, it may still be relatively large if hash evaluations themselves take very little time. For instance, if the hash function is (partly) implemented in hardware. Using Intel's AES-NI instruction set one can implement a 128-bit hash function that takes a few cycles only (e.g., see [BÖS11], noting that for one-way hash chains no collision-resistance is needed such that one can use Matyas-Meyer-Oseas for which the key is fixed). Therefore, we also provide an implementation minimizing the computational overhead at the expense of some additional storage.

We will keep some state for each pebble, or rather for each *hashing* pebble only. Although an optimal pebble P_k will store up to k hash values at any time, we observe that no more than $\lceil k/2 \rceil$ hashing pebbles will be present at any time. As in our in-place algorithms we will thus avoid any storage (and processing) for idle pebbles, as can be seen from Algorithm 3.

Algorithm 3 Fast optimal binary pebble $P_k(x)$.

Initially: array $z[0, k]$ where $z[i-1] = f^{2^k - 2^i}(x)$ for $i = 1, \dots, k$; array $a[0, \lfloor k/2 \rfloor]$, $v = 0$.

Round r , $2^k < r < 2^{k+1}$:

```
1: output  $z[0]$ 
2:  $c \leftarrow 2^{k+1} - r$ 
3:  $i \leftarrow \text{pop}_0(c)$ 
4:  $z[0, i] \leftarrow z[1, i]$ 
5:  $i \leftarrow i + 1$ ;  $c \leftarrow \lfloor c/2 \rfloor$ 
6: if  $c$  odd then  $a[v] \leftarrow (i, 0)$ ;  $v \leftarrow v + 1$ 
7:  $u \leftarrow v$ 
8:  $w \leftarrow (r \bmod 2) + i + 1$ 
9: while  $c \neq 0$  do
10:    $w \leftarrow w + \text{pop}_0(c)$ 
11:    $u \leftarrow u - 1$ ;  $(g, g) \leftarrow a[u]$ 
12:   for  $d \leftarrow 1$  to  $\lfloor w/2 \rfloor$  do
13:      $y \leftarrow z[q]$ 
14:     if  $g = 0$  then  $q \leftarrow q - 1$ ;  $g = 2^q$ 
15:      $z[q] \leftarrow f(y)$ 
16:      $g \leftarrow g - 1$ 
17:   if  $q \neq 0$  then  $a[u] \leftarrow (q, g)$  else  $v \leftarrow v - 1$ 
18:    $w \leftarrow (w \bmod 2) + \text{pop}_1(c)$ 
```

A segment $a[0, v]$ of an array a of length $\lfloor k/2 \rfloor$ suffices to store the relevant hashing pebbles, where initially $v = 0$. In each round, at most one idle pebble P_i will become hashing, and if this happens pebble P_i is added to array a , cf. line 6. Later, once pebble P_i is done hashing, it will be removed again from array a , cf. line 17.

For each hashing pebble we store two values called q and g such that q matches the value of variable q in Algorithm 2 and g matches the value of $h - 2^q$ in Algorithm 2. Hence, we use g to count down to zero starting from the appropriate powers of 2, cf. line 14. Finally, variable w is introduced such that its value matches the value of $p + i + 1 - s$ in Algorithm 2. As a result, Algorithm 3 limits the computations for each hashing pebble to a few elementary operations only.

Note that Algorithm 3 is actually quite intuitive and remarkable at the same time. E.g., by focusing on variable w , one can easily see that the total number of hashes performed by pebble P_k in any output round does not exceed $\lceil k/2 \rceil$.

7 Extensions

In this section we briefly discuss three extensions, omitting details.

First, we show how to accommodate hash chains of arbitrary length n by generalizing the initialization of Algorithms 1–3 from $n = 2^k$ to any $n \geq 1$, without incurring any overhead. That is, given a seed value x , we will iterate the hash function f for exactly $n - 1$ times and output $f^{n-1}(x)$ (e.g., as part of the registration protocol for Lamport’s scheme). At the same time, we will store precisely those intermediate hash values in array z such that the state becomes equivalent to the state of binary pebble P_k at round $2^{k+1} - (n - 1)$, where $2^{k-1} < n \leq 2^k$.

We demonstrate this for in-place speed-2 pebbles, by extending Theorem 8 to the following theorem, which exactly describes the state of speed-2 pebble $P_k(x)$ at round $c = n - 1$.

Theorem 9. *For a speed-2 pebble $P_k(x)$ in output round $2^{k+1} - c$, $1 \leq c \leq 2^k$, each idle pebble P_i stores the hash value $f^{c-(c \bmod 2^i)-2^i}(x)$ and each hashing pebble P_i stores the following $e + 1$ hash values:*

$$f^{c-(c \bmod 2^i)-2^i}(x), \dots, f^{c-(c \bmod 2^i)-2^{i-e+1}}(x), f^{c-\max(1, 3(c \bmod 2^i))}(x),$$

where e is maximal such that $c_{i-1} = \dots = c_{i-e} = 0$ with $0 \leq e \leq i$.

Based on this theorem array z can be initialized simply by inspecting the bits of c (most-significant to least-significant) while iterating f . The remaining $n - 1$ output rounds are then executed exactly as in Algorithm 1. A slightly more complicated approach applies to our optimal pebbles, for which we omit the details.

Next, we show how to construct a *cascade* of pebbles $P_k(x), P_k(x'), P_k(x''), \dots$, for seed values x, x', x'', \dots , such that the initial stage of $P_k(x')$ is run in parallel to the output stage of $P_k(x)$ (more precisely, in parallel to rounds $(2^k, 2^{k+1})$ of $P_k(x)$), and so on. Hence, as soon as $P_k(x)$ is done, $P_k(x')$ will continue with producing hash chain $f_k^*(x')$ in reverse. In general, this combination corresponds exactly to running the output stage of a P_{k+1} pebble over and over again. Therefore, the maximum number of hashes in any round will be limited to $\max(W_{k+1})$. Using optimal pebbles P_k , we thus get no more than $\lceil (k + 1)/2 \rceil$ hashes per round. Moreover, we only need to store a maximum of $k + 1$ hash values at any time (conceptually, the value y_{k+1} that a P_{k+1} pebble would store is not present), hence essentially for free; this can even be reduced to $\lceil k/2 \rceil$ per round at the expense of increasing the maximum storage to $k + 2$ hash values.

To make such a cascade useful, the hash values $f^{2^k-1}(x'), f^{2^k-1}(x''), \dots$ need to be authenticated. A straightforward way is to run the registration protocol once $f^{2^k-1}(x')$ is output, and so on. A more refined way is to apply known techniques for “re-initializing” hash chains [Goy04,ZL05], hence using a one-time signature to authenticate $f^{2^k-1}(x')$, for which the public key is incorporated in the seed x .

In fact, this approach can be extended to eliminate initialization altogether, basically by starting with a P_{k+1} pebble in its first output round, where each y_i is assigned a seed value incorporating a one-time public key. Setting $y_i = f^{2^{k'}-2^i}(x)$ for $i = k', \dots, 0$ for some small k' , and using such seed values for the remaining y_i 's, this *bootstrapping* technique can be tuned for optimal performance.

Finally, we show how to eliminate the shifting done in line 4 of Algorithms 1–3, which forms a potential bottleneck as naively copying up to $k - 1$ hash values may be relatively expensive. A standard technique to avoid such excessive copying is to use an auxiliary array of “pointers” $d[0, k)$ such that $d[i]$ points to the entry in $z[0, k)$ that should actually be used, but this would break the in-place property. Fortunately, it turns out that permutation d can be characterized nicely as a function of the round number, thus allowing us to efficiently restore the in-place property, without copying even a single hash value throughout our algorithms.

8 Concluding Remarks

We have completely resolved the case of binary pebbling of hash chains by constructing an explicit optimal schedule. A major advantage of our optimal schedule is that it allows for very efficient *in-place* pebbling algorithms. This compares favorably with the greedy pebbling algorithms of [YSEL09], which require a substantial amount of storage beyond the hash values themselves. The pseudocode of Algorithms 1–3 is readily translated into efficient program code, applying further optimizations depending on the target platform.²

The security of one-way hash chains for use in authentication mechanisms such as Lamport’s asymmetric identification scheme does not depend on the collision resistance of the hash function. Therefore, it suffices to use 128-bit hash values—rather than 256-bit hash values, say. Using, for instance, the above mentioned Matyas-Meyer-Oseas construction one obtains a fast and simple one-way function $f : \{0, 1\}^{128} \rightarrow \{0, 1\}^{128}$ defined as $f(x) = \text{AES}_{\text{IV}}(x) \oplus x$, where IV is a 128-bit string used as fixed “key” for the AES block cipher.³ Consequently, even for *very long* hash chains of length 2^{32} , our in-place optimal pebbling algorithm will just store 516 bytes (32 hash values and one 32-bit counter) and perform at most 16 hashes per identification round. Similarly, *long* hash chains of length 2^{16} would allow devices capable only of lightweight cryptography to run 65535 rounds of identification (e.g., more than twice per hour over a period of three years), requiring only 258 bytes of storage and using at most 8 hashes per round.

We leave as an open problem whether binary pebbling yields the lowest space-time product. Reversal of a length- n hash chain using optimal binary pebbling requires $\log_2 n$ storage and $\frac{1}{2} \log_2 n$ time per round, yielding $\frac{1}{2} \log_2^2 n$ as space-time product. Coppersmith and Jakobsson [CJ02] derived a lower bound of approx. $\frac{1}{4} \log_2^2 n$ for the space-time product.⁴ Whether this lower bound is achievable is doubtful, because the lower bound is derived *without* taking into account that the maximum number of hashes during any round needs to be minimized. As a potential alternative we mention “Fibonacci” pebbling, considering hash chains of

² Sample code (in Python, Java, and C) available at www.win.tue.nl/~berry/pebbling/.

³ More precisely, function f should be one-way on its iterates [Lev85, Ped96].

⁴ Incidentally, this lower bound had been found already in a completely different context [GPRS96]. Without going into details, the problem studied in the area of algorithmic (automatic/computational) differentiation [GW08] is similar to the task for our pebble $P_k(x)$ of computing the hash chain $f_k^*(x)$ and outputting this sequence in reverse. The critical difference, however, is that in the context of algorithmic differentiation the goal is basically to minimize the *total time* for performing this task. This contrasts sharply with the goal in the cryptographic context, where we want to minimize the *maximum time spent in each output round* while performing this task. See Appendix A, where we show how to achieve both logarithmic *amortized* time and logarithmic space using a trivial kind of binary pebble. The latter result is comparable to what is achieved in [Gri92]; in fact, without the performance constraint unique for the cryptographic setting, as initiated by Jakobsson [Jak02, CJ02], it is even possible to attain the lower bound of [GPRS96]. Therefore, the solutions achieved in the area of algorithmic differentiation (and in related areas such as reversible computing [Per13], for that matter) do not carry over to the cryptographic setting.

length $n = F_k$, the k th Fibonacci number (e.g., for even k , storing $k/2$ elements at distances $F_{2i} - 1$, for $1 \leq i \leq k/2$, from the end of the hash chain).

As another direction for further research we suggest to revisit the problem of efficient Merkle tree traversal studied in [Szy04], which plays a central role in hash-based signature schemes [Mer87,Mer89]; in particular, it would be interesting to see whether algorithms for generating successive authentication paths can be done in-place. More generally, research into optimal (in-place) algorithms for hash-based signatures is of major interest both in the context of lightweight cryptography (e.g., see [PCTS02,MSS13]; more references in [YSEL09]) and in the context of post-quantum cryptography (e.g., see [BDE⁺13]).

Acknowledgments It is a pleasure to thank Niels de Vreede and Thijs Laarhoven for many discussions on variants of the problem, and to thank Niels especially for suggesting the bootstrapping approach (see Section 7). Moreover, the anonymous reviewers are gratefully acknowledged for their comments.

References

- BDE⁺13. J. Buchmann, E. Dahmen, S. Ereth, A. Hülsing, and M. Rückert. On the security of the Winternitz one-time signature scheme. *International Journal of Applied Cryptography*, 3(1):84–96, 2013.
- BÖS11. J.W. Bos, O. Özen, and M. Stam. Efficient hashing using the AES instruction set. In *Cryptographic Hardware and Embedded Systems (CHES 2011)*, volume 6917 of *Lecture Notes in Computer Science*, pages 507–522, Berlin, 2011. Springer-Verlag.
- CJ02. D. Coppersmith and M. Jakobsson. Almost optimal hash sequence traversal. In *Financial Cryptography 2002*, volume 2357 of *Lecture Notes in Computer Science*, pages 102–119, Berlin, 2002. Springer-Verlag.
- Goy04. V. Goyal. How to re-initialize a hash chain. eprint.iacr.org/2004/097, April 2004.
- GPRS96. J. Grimm, L. Potter, and N. Rostaing-Schmidt. Optimal time and minimum space-time product for reversing a certain class of programs. In M. Berz, C.H. Bischof, G. Corliss, and A. Griewank, editors, *Computational Differentiation Techniques, Applications, and Tools*, pages 95–106, Philadelphia, 1996. SIAM.
- Gri92. A. Griewank. Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation. *Optimization Methods and Software*, 1(1):35–54, 1992.
- GW08. A. Griewank and A. Walther. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. SIAM, Reading (MA), 2nd edition, 2008.
- Hal94. N. Haller. The S/KEY one-time password system. In *Proceedings of the Symposium on Network and Distributed System Security (NDSS)*, pages 151–157. Internet Society, February 1994. See also en.wikipedia.org/wiki/S/KEY.
- IR01. G. Itkis and L. Reyzin. Forward-secure signatures with optimal signing and verifying. In *Advances in Cryptology—CRYPTO ’01*, volume 2139 of *Lecture Notes in Computer Science*, pages 332–354, Berlin, 2001. Springer-Verlag.
- Jak02. M. Jakobsson. Fractal hash sequence representation and traversal. In *Proc. IEEE International Symposium on Information Theory (ISIT ’02)*, page 437. IEEE Press, 2002. Full version eprint.iacr.org/2002/001.
- Lam81. L. Lamport. Password authentication with insecure communication. *Communications of the ACM*, 24(11):770–772, 1981.
- Lev85. L. Levin. One-way function and pseudorandom generators. In *Proc. 17th Symposium on Theory of Computing (STOC ’85)*, pages 363–365, New York, 1985. A.C.M.
- Mer87. R. Merkle. A digital signature based on a conventional encryption function. In *Advances in Cryptology—CRYPTO ’87*, volume 293 of *Lecture Notes in Computer Science*, pages 369–378, Berlin, 1987. Springer-Verlag.

- Mer89. R. Merkle. A certified digital signature. In *Advances in Cryptology—CRYPTO '89*, volume 435 of *Lecture Notes in Computer Science*, pages 218–238, Berlin, 1989. Springer-Verlag.
- MSS13. N. Mourier, R. Stampf, and F. Strenzke. An implementation of the hash-chain signature scheme for wireless sensor networks. In *Lightweight Cryptography for Security and Privacy (LightSec 2013)*, volume 8162 of *Lecture Notes in Computer Science*, pages 68–80, Berlin, 2013. Springer-Verlag.
- PCTS02. A. Perrig, R. Canetti, J.D. Tygar, and D. Song. The TESLA broadcast authentication protocol. *RSA CryptoBytes*, 5(2):2–13, 2002.
- Ped96. T. P. Pedersen. Electronic payments of small amounts. In *Security Protocols*, volume 1189 of *Lecture Notes in Computer Science*, pages 59–68, Berlin, 1996. Springer-Verlag.
- Per13. K. Perumalla. *Introduction to Reversible Computing*. Chapman and Hall/CRC, 2013.
- Szy04. M. Szydło. Merkle tree traversal in log space and time. In *Advances in Cryptology—EUROCRYPT '04*, volume 3027 of *Lecture Notes in Computer Science*, pages 541–554, Berlin, 2004. Springer-Verlag.
- YSEL09. Dae Hyun Yum, Jae Woo Seo, Sungwook Eom, and Pil Joong Lee. Single-layer fractal hash chain traversal with almost optimal complexity. In *Topics in Cryptology – CT-RSA 2009*, volume 5476 of *Lecture Notes in Computer Science*, pages 325–339, Berlin, 2009. Springer-Verlag.
- ZL05. Y. Zhao and D. Li. An improved elegant method to re-initialize hash chains. eprint.iacr.org/2005/011, January 2005.

A Rushing Binary Pebbles

Below we show that it is easy to achieve logarithmic space and logarithmic *amortized* time per output round for binary pebbling algorithms. This contrasts sharply with the binary pebbling algorithms presented in this paper, which not only achieve logarithmic space but also logarithmic *worst case* time per output round.

In fact, logarithmic amortized time per output round is achieved by *any* binary pebble P_k as follows directly from Theorem 1: any pebble P_k performs $\#W_k = (k-2)2^{k-1} + 1$ hashes in total during its output stage consisting of 2^k rounds, hence the amortized number of hashes per output round is equal to $((k-2)2^{k-1} + 1)/2^k \approx k/2 - 1$. This holds for any schedule T_k satisfying $\#T_k = 2^k - 1$.

To achieve logarithmic space as well, we choose T_k such that the storage requirements for P_k are minimized. This can simply be done by postponing the evaluation of all hashes to the last round of the initial stage of P_k . Concretely, we define a **rushing** pebble P_k by setting $T_0 = \{\}$ and $T_k = 0^{*2^k-2} \parallel 2^k - 1$ for $k \geq 1$. Hence, a rushing pebble performs all $2^k - 1$ hashes in the last round of its initial stage; see also Figure 2.

As a consequence, a rushing pebble minimizes its storage requirements throughout, at the cost of some computationally very expensive rounds. This is summarized in the following results, which we state without proof.

Theorem 10. *A rushing pebble P_k uses up to $\max(W_k) = 2^{k-1} - 1$ hashes in any output round, for $k \geq 1$.*

Lemma 5. *For a rushing pebble P_k , we have:*

$$S_0 = \{1\}, \quad S_k = (0^{*2^k-1} \parallel 1 \parallel S_{k-1}) + (1^{*2^{k-1}} \parallel S_{k-1} \parallel 0^{*2^{k-1}}).$$

Theorem 11. *A rushing pebble P_k uses up to $\max(S_k) = k + 1$ storage.*

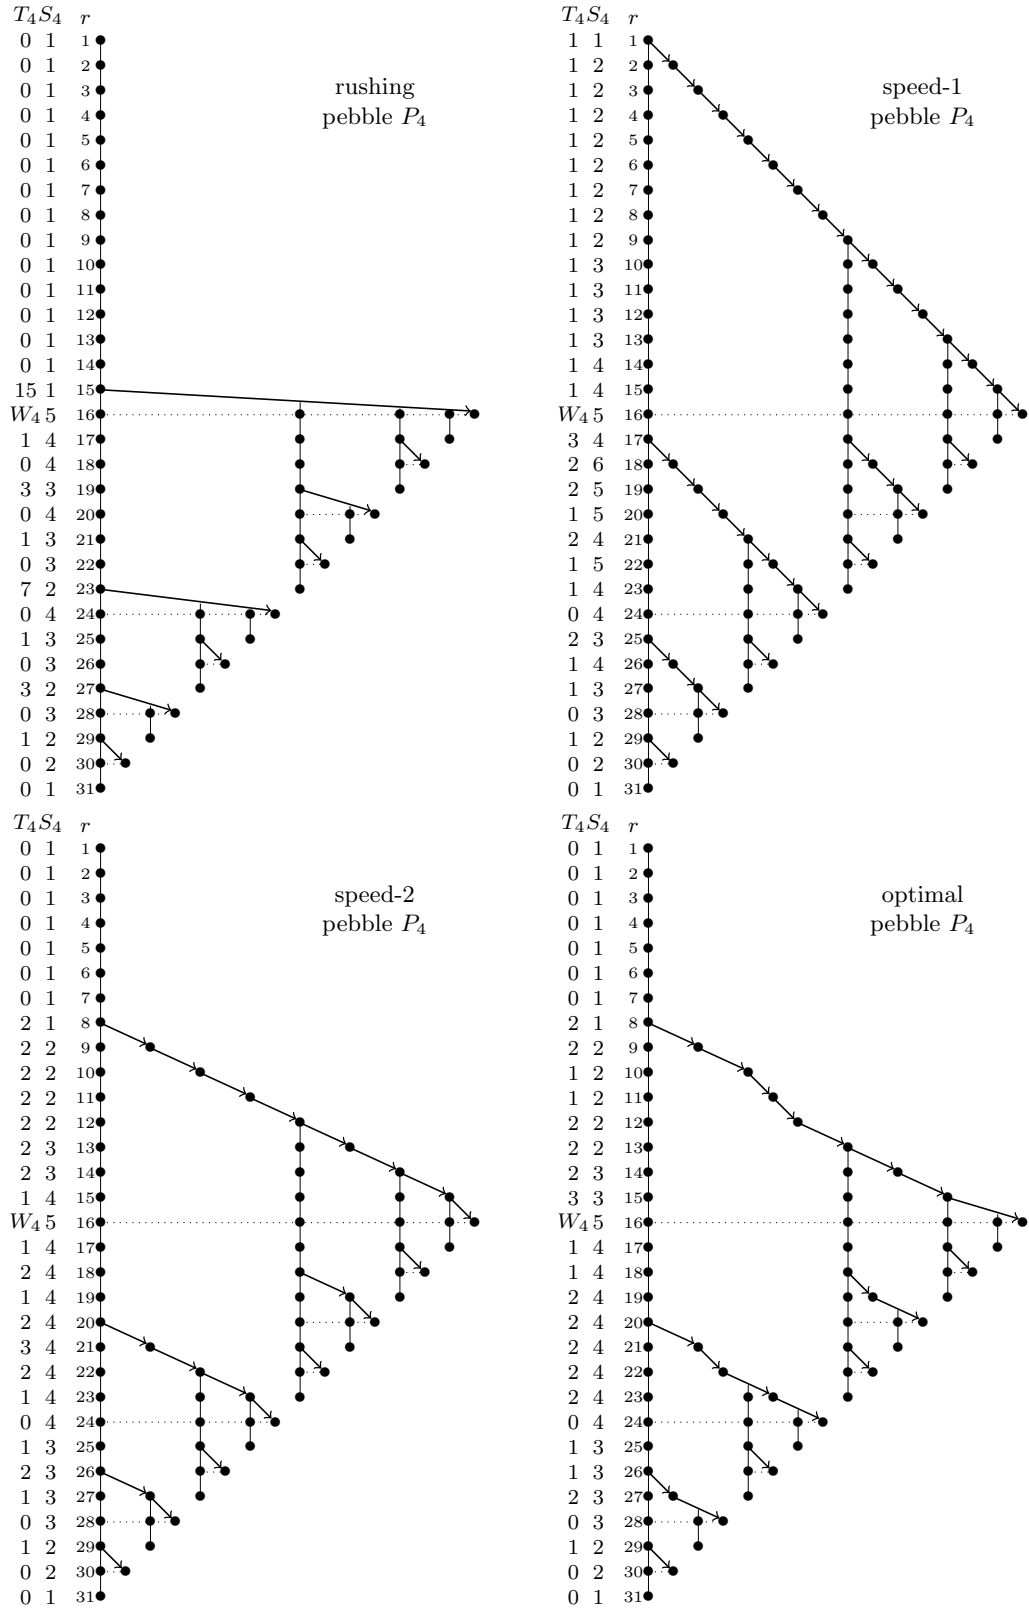


Fig. 2. Schedule T_4 resp. work W_4 and storage S_4 for binary pebbles P_4 in rounds $r = 1, \dots, 31$. Each bullet represents a stored value, arrows represent hashing, vertical lines represent copying.