

A Comparison of Cryptanalytic Tradeoff Algorithms

Jin Hong · Sunghwan Moon

Abstract The three major time memory tradeoff algorithms are compared in this paper. Specifically, the Hellman tradeoff algorithm, the distinguished point tradeoff method, and the rainbow table method, in their non-perfect table versions, are considered.

We show that, under parameters that are typically considered in theoretic discussions of the tradeoff algorithms, Hellman and distinguished point tradeoffs perform very close to each other and the rainbow table method performs somewhat better than the other two algorithms. Our method of comparison can easily be applied to other situations, where the conclusions could be different.

The analysis presented in this paper takes the effects of false alarms into account and also fully considers techniques for reducing storage, such as the ending point truncation method and index files.

Keywords time memory tradeoff · Hellman · distinguished point · rainbow table · random function

1 Introduction

There are numerous security systems in use today that rely on passwords. Access to many contents on the network requires one to login with a password and many file formats today have security features that restrict access to the file until the correct password is supplied. These usually employ a *password hash* system, which stores a one-way function image of the password in the system or file. Indeed, storing the password in its raw form within the file one wishes to set access control to would be meaningless.

A time memory tradeoff algorithm attempts to recover the password from the knowledge of the one-way function image, with the help of a table created through pre-computation.

Contact author: Jin Hong

J. Hong
Department of Mathematical Sciences and ISaC, Seoul National University, Seoul 151-747, Korea
E-mail: jinhong@snu.ac.kr

S. Moon
Department of Mathematics, Texas A&M University, College Station, TX 77843-3368, USA
E-mail: shmoon@math.tamu.edu

A realistic barrier to applying the tradeoff technique to any specific security system is the massive pre-computation required, before the actual attack can be mounted. The pre-computation cost is roughly proportional to the space of possible passwords and, since most users do not use long or random passwords, the tradeoff attacker is free to choose a smaller subset consisting of short or more likely passwords and decide to be satisfied with recovering only those passwords belonging to this subset. Then the pre-computation requirement no longer stands as a barrier to the tradeoff attack.

It has long been known that properly *salting* a password can remove any realistic threats from time memory tradeoff attacks. But, such measures are still not being taken by many proprietary systems, and some systems are known to be using both the newer salted and the older non-salted versions of the security system simultaneously to remain compatible with older systems. Hence, the time memory tradeoff technique still remains a powerful tool against these vulnerable password hash systems. Since human generated passwords will continue to be used for some time, one would like to understand the power and limitations of the tradeoff techniques.

In this work, we consider the three major tradeoff algorithms. These are the original tradeoff algorithm [8] devised by Hellman, the distinguished points method, attributed to Rivest in [4, p.100], and the rainbow table method [12], announced by Oechslin. There are perfect table versions of these methods and, even though they are more efficient in recovering passwords, they require more pre-computation to work. These are hence less practical and we shall not consider them in this paper. There are also multi-target versions of tradeoff algorithms [3, 7] which attracted attention as attacks on streamciphers, but we will not work with these either. The most practical application of the tradeoff technique today is with password hash systems and we will present the current work with this application in mind.

It has been shown [2] that the performance of tradeoff algorithms that are known today is asymptotically the best one can hope to find, among all algorithms belonging to a certain category. But the measure of being best used in this theory is only accurate modulo a small multiplicative factor. The relative performance of different tradeoff algorithms has not been accurately analyzed, and the choice of which algorithm to use can be a difficult one for someone new to the tradeoff technique. In practice, experience seems to be the major basis of the decisions, and researchers seem to have varied opinions on which algorithm performs better.

Comparison of tradeoff algorithms has been a controversial subject. There are two major obstacles to making a fair comparison of tradeoff algorithms. The first is that the online time of each algorithm is hard to predict accurately, due to the effects of events called false alarms. Some answers to this problem may now be found in [1, 9] for the Hellman and rainbow cases. The second obstacle concerns the number of bits that is required to store each table entry resulting from the pre-computation. There is a technique for reducing this number of bits called ending point truncation which has not been fully analyzed. Due to these obstacles, previous comparisons of tradeoff algorithms have mostly relied on rough arguments that emphasize a certain advantageous characteristic of one algorithm against another.

There is a natural occurring measure¹ of how efficiently a tradeoff algorithm can balance time against memory and this measure becomes accessible once the first obstacle mentioned above is resolved. In this work, we carefully note that this measure of tradeoff efficiency for different algorithms are expressed in different units and provide arguments for converting

¹ The optimal value of this measure is referred to as the tradeoff characteristic in [1], where it is used to compare the perfect version of the rainbow table method against other algorithms.

them to directly comparable units. This transition of units is intimately connected to the second obstacle mentioned above.

Apart from the above two obstacles that are due to our lack of knowledge, there is yet another problem of how to compare different tradeoff performances that are achievable only after different pre-computation efforts.

In this work, we clear both of the obstacles mentioned above and use it to provide fair comparisons between tradeoff algorithms. More precisely, we shall present a method to visualize what can be achieved by each algorithm in terms of pre-computation and tradeoff performance. This will be done in a unified way so that the range of choices possible with each algorithm can directly be compared against each other. Any potential user of the tradeoff technique can use this information to decide on which algorithm to use and which set of parameters to use with the algorithm. The judgement of which algorithm is more suitable depends on how the user relatively values the pre-computation cost and tradeoff efficiency, and hence cannot be done in an objective manner in most cases.

In presenting the above comparison method, we will mainly focus on a certain set of parameters that is typically considered during theoretic analysis of tradeoff algorithms. Under this parameter set, the Hellman and distinguished point methods is shown to perform very similarly to each other, while the rainbow table method shows better performance than the other two algorithms, under the additional assumption that the success rate requirement placed on the tradeoff algorithms is rather high. In this specific case, the comparison judgement will stand true for any reasonable way of valuing the pre-computation cost and tradeoff efficiency. Conclusions at other parameter ranges, which may be suitable for many situations, can be different.

Another contribution of this paper is in making the basis of tradeoff technique analysis theoretically more concrete. We discuss the use of the term *random function* and how it is related to the analysis of the tradeoff technique. We identify common arguments that are not mathematically correct in the strict sense and provide plausible justifications for still using such arguments in the analysis of tradeoff algorithms.

The rest of the paper is organized as follows. We briefly fix notation and terminologies in the next section. The use of random functions in the analysis of the tradeoff algorithms is discussed in Section 3. This is followed by a section clarifying the connection between the theory of tradeoff algorithms and its application to password hash systems. In Section 5, 6, and 7, we study the distinguished point, Hellman, and rainbow table tradeoff algorithms, in turn. For each algorithm, we present an accurate version of the tradeoff curve that does not ignore small multiplicative factors and also analyze applicable storage reduction techniques. Comparison of tradeoff performances under different parameter sets for the same algorithm is discussed in Section 8, and performance comparison between different algorithms is given in Section 9. Finally, the work is summarized and concluding remarks are given in Section 10.

2 Time Memory Tradeoff Algorithms

The details of the tradeoff algorithms will not be explained and we ask readers to refer to the original papers on the Hellman tradeoff [8] and the rainbow table method [12]. Here, we only list and fix the basic terminologies. Notation of this section will be used throughout the paper. From now on, we shall refer to the tradeoff algorithm that integrates the distinguished point idea into the Hellman tradeoffs as the DP tradeoff and refer to the rainbow table method simply as the rainbow tradeoff.

A *Hellman chain*, associated with a one-way function $F : \mathcal{N} \rightarrow \mathcal{N}$, is of the form

$$\mathbf{sp}_i = \mathbf{x}_{i,0} \xrightarrow{F} \mathbf{x}_{i,1} \xrightarrow{F} \cdots \xrightarrow{F} \mathbf{x}_{i,t} = \mathbf{ep}_i \quad (1 \leq i \leq m).$$

The *Hellman table* $\{(\mathbf{sp}_i, \mathbf{ep}_i)\}_{i=1}^m$, consisting of starting and ending point pairs, is sorted on the ending points to make table lookups easier. We have omitted the *reduction functions*, as this will only be mentioned briefly in Section 4.3. The complete set of m chains, consisting of m rows and $t + 1$ columns, is a *Hellman matrix*. The original Hellman tradeoff specified the matrix stopping rule $mt^2 = |\mathcal{N}|$ and considered the use of exactly t tables, but we shall allow more flexibility. We let the positive integer parameters m and t satisfy $mt^2 = H_{msr} |\mathcal{N}|$, with the positive constant $H_{msr} = \Theta(1)$, i.e., the constant H_{msr} specifying the matrix stopping rule is to be neither very large nor too close to zero. The number of tables is set to $l = H_{nt} t$, where $H_{nt} = \Theta(1)$, so that l and t are roughly of the same order.

A *DP chain* is a chain iteratively constructed by applying $F : \mathcal{N} \rightarrow \mathcal{N}$ until a *distinguishing point* (DP) is reached. Although not absolutely necessary, to simplify our later discussion, we shall assume that the starting points are always chosen among non-DPs. Hence, in a DP chain, every point before the ending point, including the starting point, is a non-DP. The effects of this choice is quite negligible and a rigorous treatment that allows starting points to be DPs will lose any differences it has with our treatment, during the many numerical approximations that lead to the final statement of results.

We assume the distinguishing property is defined in such a way that a random point of \mathcal{N} satisfies it with probability $\frac{1}{t}$. To detect chains that fall into infinite loops, a chain length bound of \hat{t} is fixed and any chain that fails to reach a DP by this length during either the pre-computation phase or the online phase is discarded. Chains are generated until each table contains approximately m chains. What we mean by approximately m will be made more explicit at the start of Section 5.

Even though some of our result statements will display its dependence on \hat{t} , we shall mainly be interested of the limiting case where \hat{t} is sufficiently large. When $\hat{t} \gg t$, the number of discarded chains is minimized, so that any additional pre-computation is more efficiently be transformed into higher success rate of the tradeoff algorithm. Since pre-computation cost is the main barrier to any large scale applications of the tradeoff technique, such a choice is only natural in practical applications.

The terms *DP table* and *DP matrix* will be used, even though the DP matrix, with chains of variable lengths, can no longer be visualized as a rectangle. The positive integer parameters m and t are chosen to satisfy $mt^2 = D_{msr} |\mathcal{N}|$, with the positive constant $D_{msr} = \Theta(1)$. The DP tradeoff will use $l = D_{nt} t$ tables with $D_{nt} = \Theta(1)$.

We similarly have the notions of *rainbow chain*

$$\mathbf{sp}_i = \mathbf{x}_{i,0} \xrightarrow{F_1} \mathbf{x}_{i,1} \xrightarrow{F_2} \cdots \xrightarrow{F_t} \mathbf{x}_{i,t} = \mathbf{ep}_i \quad (1 \leq i \leq m),$$

rainbow table, and *rainbow matrix*. For rainbow tradeoffs, it is usual to take $mt = R_{msr} |\mathcal{N}|$ with $R_{msr} = \Theta(1)$. A rainbow tradeoff will use l tables, where l is a small positive integer.

Our *inversion target* will always be written as $\mathbf{y} = F(\mathbf{x})$ throughout the paper. The input \mathbf{x} and output \mathbf{y} will be referred to as the *password* and *password hash*, respectively, even though they may not be related to any password system.

If the current end $F^k(\mathbf{y})$ of the Hellman tradeoff's *online chain*

$$\mathbf{y} \xrightarrow{F} F(\mathbf{y}) \xrightarrow{F} F^2(\mathbf{y}) \xrightarrow{F} \cdots \xrightarrow{F} F^k(\mathbf{y})$$

of length k matches an ending point \mathbf{ep}_i , we have an *alarm*. If an alarm involving \mathbf{ep}_i shows the property $F^{t-k-1}(\mathbf{sp}_i) \neq \mathbf{x}$, it is said to be a *false alarm*. Notice that $F^{t-k-1}(\mathbf{sp}_i) = \mathbf{x}$ is

not guaranteed by the weaker condition $F^{t-k}(\mathbf{sp}_i) = \mathbf{y}$. False alarms are caused by *merges* between the online chain and a *pre-computed chain*. This notion of false alarms is also used with DP and rainbow tradeoffs.

Throughout this paper, the one-way function $F : \mathcal{N} \rightarrow \mathcal{N}$ being considered will always act on a set \mathcal{N} of size N . The parameters m and t satisfying an appropriate $mt^2 = H_{msr}N$, $mt^2 = D_{msr}N$, or $mt = R_{msr}N$ are assumed to be reasonable in the sense that $1 \ll m, t \ll N$.

There are *perfect* table versions of the DP and rainbow tradeoffs. These are when chains with identical ending points are removed and regenerated. These cases are interesting and can be more efficient than the non-perfect versions during the online phase, but requires more pre-computation to arrive at the same success rate. As we want to focus on the practical applications of tradeoff algorithms, we shall not consider perfect tables.

3 Random Functions

The purpose of this section is to understand the usage of the term *random function* and to become familiar with arguments that involve random functions.

Throughout this section, \mathcal{A} and \mathcal{B} will be finite sets of respective sizes A and B . The set of all functions $F : \mathcal{A} \rightarrow \mathcal{B}$ will be denoted by $\mathcal{B}^{\mathcal{A}}$, where the exponent notation symbolizes the fact that each such function is an ordered A -tuple of elements from \mathcal{B} , i.e., that it belongs to the A -times cartesian product of \mathcal{B} .

3.1 Definition of a random function

Readers that are fully comfortable with the cryptographic usage of the term random function may skip this subsection. The content of this subsection is briefly explained in [6, Section 5.2], but since we were unable to find this in any formally published work, in the interest of making it more widely known, we shall write this down with some more details.

Let us begin our discussion with a simpler notion that we are very comfortable with. Consider the following sentence.

The size of a random element from the set $\{0, 1, 2, 3\}$ is expected to be $\frac{3}{2}$.

We all know that this is a short way of expressing the following more explicit statement.

If an element is chosen uniformly at random from the set $\{0, 1, 2, 3\}$, then we can expect its size to be $\frac{3}{2}$.

The term *random element* that appears in the first expression does not refer to a true element that belongs to the set in consideration. Instead, it indicates that a certain method for selecting an element be used. Whenever the term is used, it is followed by a claim to some expected value, and one is to understand that the probability distribution to be used in computing the claimed expected value is to be the uniform distribution.

Sometimes the specific element that has been chosen is referred to as *the* random element, but this usage of the term is a source of confusion. Once the process of selection from the set is complete, no randomness can be found in the resulting specific element.

Let us now turn to random functions. The following sentence presents a typical usage of the term random function.

The image size of a random function $F : \{0, 1\} \rightarrow \{0, 1\}$ is expected to be $\frac{3}{2}$.

Once again, we should take this simply as a short way of expressing the following more explicit statement.

Consider the set of all functions $F : \{0, 1\} \rightarrow \{0, 1\}$. If a function is chosen uniformly at random from this set of four functions, then we can expect the image size of the selected function to be $\frac{3}{2}$.

The term *random function* does not refer to a concrete function that belongs to the set of functions in consideration. Instead, the phrase specifies that a certain method for *selecting* a function be used. Once the selection has been made, no randomness remains in the resulting specific and fixed function. The term random function is always accompanied by a claim to an expected value, and the term is used to signify that the uniform distribution on the set of functions is to be used in computing the expectation.

So far, there seems to be no ambiguity concerning the notion of random functions. This has been a straightforward extension of our common use of the term random element or random number. We now turn to the following quote² from a textbook [11, Section 5.6], that explains random functions.

A random function $F : \mathcal{A} \rightarrow \mathcal{B}$ is a function which assigns independent and random values $F(x) \in \mathcal{B}$ to all arguments $x \in \mathcal{A}$.

This is the *definition* of a random function which cryptographers, especially those working on symmetric key cryptography, are more accustomed to. Whereas, in our previous discussion, we could not define a random function and only explain the usage of the term, the above quote seems to be defining a random function as a concrete object. In fact, at times, readers may have come across sentences that resemble the following.

If $F : \{0, 1\} \rightarrow \{0, 1\}$ is *the* random function, then its image size is expected to be $\frac{3}{2}$.

The random function referred to in this usage example seems to be the one that has just been defined through the textbook quote, rather than have anything to do with the previously mentioned selection from a set of functions. But a closer look reveals that the quoted definition does not define a true function. Instead, the definition provides a process by which a specific function may be *constructed* or defined. A function that is being constructed is not truly a function until it has fully been defined on all its inputs. Hence, at least to a strict mathematician, the quote from the textbook is self-contradictory. Furthermore, as soon as the random function is fully specified, thus earning the status of a true function, no randomness can be found within the function. One more thing to note is that the latest usage example makes a claim to a certain expected value, as was done in all our previous examples. Hence, once again, we are not left with the definition of a concrete object, but a usage example where the appearance of the term random function specifies certain actions to be taken in constructing a function.

We have explained two approaches to the term random function. In both approaches, there are no real entities that correspond to the term random function. Sometimes the selected or fully constructed function is referred to as the random function, but such usage is a source of confusion, as no randomness remains in the specific function. The current situation is strictly analogous to the situation with random elements, except that there is more room for confusion with random functions in that we intuitively tend to view a fully defined specific function as still looking rather random, if we can not find any simple rule that can be used to specify the output corresponding to each input. The term *random function* is not the

² We have given some notational changes.

name of a mathematical object and cannot be defined. Instead, its appearance only signals that a certain process for specifying a function be used in computing an expected value.

Let us distinguish between the two approaches to random functions by referring to them through the terms *selection* and *construction*. We will now show that the two approaches are equivalent. Fix any specific function $F_0 : \mathcal{A} \rightarrow \mathcal{B}$. If a function is constructed by assigning a randomly and independently chosen element of \mathcal{B} to each element of \mathcal{A} , then the probability for the randomly constructed function to become identical to F_0 is $(\frac{1}{B})^A$. On the other hand, if a function is selected uniformly at random from the set of functions $\mathcal{B}^{\mathcal{A}}$, since the size of $\mathcal{B}^{\mathcal{A}}$ is B^A , the probability that F_0 is chosen is $\frac{1}{B^A}$. Hence, the construction approach to specifying a function brings about the uniform distribution on $\mathcal{B}^{\mathcal{A}}$. We can conclude that the selection and construction approaches to random functions are identical in that they provide the same distribution on which to compute the claimed expected value.

Even though the selection approach and construction approach give the same distribution on $\mathcal{B}^{\mathcal{A}}$, it is often much easier to work with the construction approach when explicitly computing expectations. Roughly, the selection of many points leads to the selection of a single function, and one can compute an expectation defined over the distribution of functions by suitably combining many expectations defined over the distribution of points. The expectations defined over points are simpler and easier to deal with. In our subsequent uses of the random function notion, we shall not distinguish between the two approaches unless absolutely necessary, but will mainly use expressions that can be read more naturally with the construction approach in mind.

We briefly remark that a randomly constructed function is very close to a random oracle. The only possible difference is that, while one needs to make external queries for the function values when using random oracles, with the construction approach to random functions, the one working with the function can choose each of the function images randomly by himself.

3.2 Random function arguments

There will be many arguments given in this paper that use the notion of random functions and not all of them will be strictly correct in the mathematical sense. In this subsection, we identify and discuss one such logical gap that is not easily noticed and provide justifications for still using such an argument.

Before discussing proofs that use the notion of random functions, we briefly digress and state a technical lemma that will be used throughout this paper.

Lemma 1 *For positive integers A and B, we have*

$$\left| \exp\left(-\frac{A}{B}\right) - \left(1 - \frac{1}{B}\right)^A \right| < \left\{ \frac{1}{2} \frac{A}{B^2} + \frac{1}{(A+1)!} \left(\frac{A}{B}\right)^{A+1} \right\} \exp\left(\frac{A}{B}\right).$$

Hence, if A and B are large integers such that $A = O(B)$, then

$$\exp\left(-\frac{A}{B}\right) \approx \left(1 - \frac{1}{B}\right)^A$$

is an accurate approximation.

Proof We start by writing $\exp\left(-\frac{A}{B}\right)$ in its Taylor series form and fully expanding the term $\left(1 - \frac{1}{B}\right)^A$.

$$\begin{aligned} & \left| \exp\left(-\frac{A}{B}\right) - \left(1 - \frac{1}{B}\right)^A \right| \\ &= \left| \left\{ 1 - \frac{A}{B} + \frac{1}{2!} \left(\frac{A}{B}\right)^2 - \dots \right\} - \left\{ 1 - \binom{A}{1} \frac{1}{B} + \binom{A}{2} \frac{1}{B^2} - \dots + (-1)^A \binom{A}{A} \frac{1}{B^A} \right\} \right|. \end{aligned}$$

After noting that the beginning two pairs of terms cancel out, we collect corresponding pairs from the two sequences of terms and bound the above by

$$\left\{ \left| \frac{A^2}{2!} - \binom{A}{2} \right| \frac{1}{B^2} + \dots + \left| \frac{A^A}{A!} - \binom{A}{A} \right| \frac{1}{B^A} \right\} + \left\{ \frac{1}{(A+1)!} \left(\frac{A}{B}\right)^{A+1} + \dots \right\}. \quad (1)$$

Since we have

$$\begin{aligned} 0 &\leq \frac{A^k}{k!} - \binom{A}{k} = \frac{1}{k!} \{A^k - A(A-1)\dots(A-k+1)\} \\ &= \frac{1}{k!} \left\{ \frac{k(k-1)}{2} A^{k-1} - \dots + (-1)^k (k-1)! A \right\} \\ &\leq \frac{1}{k!} \frac{k(k-1)}{2} A^{k-1} = \frac{1}{2} \frac{A^{k-1}}{(k-2)!}, \end{aligned}$$

for every $2 \leq k \leq A$, the terms of (1) that appear inside the first set of braces is bounded by

$$\begin{aligned} & \frac{1}{2} \left\{ \frac{A}{0!} \frac{1}{B^2} + \frac{A^2}{1!} \frac{1}{B^3} + \frac{A^3}{2!} \frac{1}{B^4} + \dots + \frac{A^{A-1}}{(A-2)!} \frac{1}{B^A} \right\} \\ &= \frac{1}{2} \frac{A}{B^2} \left\{ 1 + \frac{1}{1!} \frac{A}{B} + \frac{1}{2!} \left(\frac{A}{B}\right)^2 + \dots + \frac{1}{(A-2)!} \left(\frac{A}{B}\right)^{A-2} \right\} \\ &\leq \frac{1}{2} \frac{A}{B^2} \exp\left(\frac{A}{B}\right). \end{aligned}$$

As for the second set of braces from (1), it is easy to see that

$$\frac{1}{(A+1)!} \left(\frac{A}{B}\right)^{A+1} \exp\left(\frac{A}{B}\right)$$

can serve as its very rough bound. It now suffices to gather the two bounds to arrive at the claim. \square

Note that the error in the approximation $\left(1 - \frac{1}{B}\right)^A \approx \exp\left(-\frac{A}{B}\right)$ is extremely small. For example, when $A = B$, the bound stated in the lemma is at most $\frac{e}{B}$. This approximation will be used so frequently in this paper, that we shall not even reference the above lemma when applying it.

Our first example proof that uses random functions is now given.

Lemma 2 *When the finite sets \mathcal{A} and \mathcal{B} are sufficiently large with $A \leq B$, the image size of a random function $F : \mathcal{A} \rightarrow \mathcal{B}$ is expected to be*

$$B \left\{ 1 - \left(1 - \frac{1}{B}\right)^A \right\} \approx B \left\{ 1 - \exp\left(-\frac{A}{B}\right) \right\}.$$

Proof The approximation follows from the condition $A \leq B$ and Lemma 1. We mentioned this since we are at the very first application of the approximation, but our subsequent use of the approximation will be less informative and even silent.

It suffices to show the following statement, which is the interpretation given by the construction approach to random functions.

If one constructs a function $F : \mathcal{A} \rightarrow \mathcal{B}$ by assigning independently and randomly chosen elements of \mathcal{B} to each input element of \mathcal{A} , then the image size of the resulting function is expected to be $B\{1 - (1 - \frac{1}{B})^A\}$.

Let us focus on the elements of \mathcal{B} that remain as non-image points after the function construction is complete. We want to compute the expected ratio of such elements among \mathcal{B} .

If necessary, we can enumerate all elements of the finite set \mathcal{A} , so that we may refer to each element as an i -th element. When a random point of \mathcal{B} is assigned to the first element of \mathcal{A} , the ratio of points among \mathcal{B} that remain as non-images will become $(1 - \frac{1}{B})$. After a random point of \mathcal{B} has been assigned to the second element of \mathcal{A} , we can expect $(1 - \frac{1}{B})^2$ of the points of \mathcal{B} to remain as non-image points.

Since each assignment is independent of all other assignments, we may conclude that when every element of \mathcal{A} has been made to map to some element of \mathcal{B} , the ratio of the non-image points among \mathcal{B} is expected to be $(1 - \frac{1}{B})^A$. Hence, the ratio of the image points among \mathcal{B} is expected to be $\{1 - (1 - \frac{1}{B})^A\}$, as stated by the lemma. \square

Given a function $F : \mathcal{N} \rightarrow \mathcal{N}$, we shall write $F^k = F \circ \dots \circ F$ for the k -times iterated composition of the function F . The proof given above that involved random functions contains no logical gaps. In the rest of this section, we shall try to illuminate the hidden difficulty involved in proving the next lemma and also try to convince the readers that this logical gap may safely be ignored.

Lemma 3 *Let $F : \mathcal{N} \rightarrow \mathcal{N}$ be the random function on a finite set of size N . Given any nonnegative $m_0 \leq N$, define m_k through the recursive relation*

$$\frac{m_{i+1}}{N} = 1 - \exp\left(-\frac{m_i}{N}\right) \quad (i = 0, \dots, k-1).$$

If $\mathcal{M} \subset \mathcal{N}$ is of size m_0 , then the iterated image size $|F^k(\mathcal{M})|$ is expected to have the asymptotic form m_k , as N is sent to infinity.

The special case of this lemma for when the input set \mathcal{M} is the complete domain \mathcal{N} may be found in [5, 11]. The case when \mathcal{M} is not the complete domain is used in [12] to state the success probability of a non-maximal rainbow table. The work [12] includes a proof that does not mention random functions, and while the core argument given here will be the same, we focus on whether the core argument implies the above lemma.

We start our discussion with the single iteration case.

Lemma 4 *Let $F : \mathcal{N} \rightarrow \mathcal{N}$ be the random function on a finite set of size N . If $\mathcal{M} \subset \mathcal{N}$ is of size m_0 , then the size of $F(\mathcal{M})$ is expected to be*

$$m_1 = N\left\{1 - \left(1 - \frac{1}{N}\right)^{m_0}\right\}.$$

Proof It is clear that in dealing with the construction approach to random functions, the order in which elements of the domain are assigned elements of the range does not affect the expected value being computed. Hence, we can choose to give assignments to all elements of \mathcal{M} , before assigning elements to other points of the domain. It is also possible to stop the assignments when we are finished with \mathcal{M} , since the rest cannot affect the size of $F(\mathcal{M})$. With this observation, it now suffices to reread proof of Lemma 2 with the replacements $\mathcal{A} \leftarrow \mathcal{M}$, $A \leftarrow m_0$, $\mathcal{B} \leftarrow \mathcal{N}$, and $B \leftarrow N$. \square

We emphasize that this lemma for the simplest case contains no approximation or even any mentioning of an asymptotic form. The value written is the exact expected value. To see if this result concerning a single step can be iterated, we consider the following statement.

Lemma 5 (Incorrect) *Let $F : \mathcal{N} \rightarrow \mathcal{N}$ be the random function acting on a finite set of size N . If $\mathcal{M} \subset \mathcal{N}$ is of size m_0 , then the size of $F^2(\mathcal{M})$ is expected to be m_2 , where*

$$m_1 = N \left\{ 1 - \left(1 - \frac{1}{N} \right)^{m_0} \right\} \quad \text{and} \quad m_2 = N \left\{ 1 - \left(1 - \frac{1}{N} \right)^{m_1} \right\}.$$

As with the single iteration case, given by Lemma 4, we would expect this twice iterated version to hold exactly, i.e, contain no approximation, but we can *disprove* this statement with an explicit counterexample.

The set of all functions $F : \{0, 1\} \rightarrow \{0, 1\}$ can be visualized as follows.



When the input set \mathcal{M} is a single point, the image size expectation is clearly 1. This is in agreement with the value

$$2 \left\{ 1 - \left(1 - \frac{1}{2} \right)^1 \right\} = 1,$$

computed according to Lemma 4. When the input set is the complete domain $\{0, 1\}$, the image size expectation is

$$E_F [|F(\{0, 1\})|] = \frac{1}{4} \cdot 1 + \frac{1}{4} \cdot 2 + \frac{1}{4} \cdot 2 + \frac{1}{4} \cdot 1 = \frac{3}{2},$$

and this is also identical to the value

$$E_F [|F(\{0, 1\})|] = 2 \left\{ 1 - \left(1 - \frac{1}{2} \right)^2 \right\} = \frac{3}{2},$$

computed according to Lemma 4. Hence, as we have already proved, Lemma 4 holds exactly true for the $\mathcal{N} = \{0, 1\}$ case, regardless of the input set size.

Now, the four functions $F^2 = F \circ F$ can be visualized as follows.



Hence, when the input set \mathcal{M} is taken to be the complete domain, the twice iterated image size expectation is

$$E_F [|F^2(\{0, 1\})|] = \frac{2}{4} \cdot 1 + \frac{2}{4} \cdot 2 = \frac{3}{2}. \quad (2)$$

The corresponding value claimed by Lemma 5 is

$$2 \left\{ 1 - \left(1 - \frac{1}{2} \right)^{2 \left\{ 1 - \left(1 - \frac{1}{2} \right)^2 \right\}} \right\} = 2 \left\{ 1 - \left(1 - \frac{1}{2} \right)^{\frac{3}{2}} \right\} \approx 1.293. \quad (3)$$

The two values are clearly not equal.

Since Lemma 4 always holds exactly true, in particular, since it holds exactly true in the environment of the counterexample, we can conclude that Lemma 5 does not logical follow directly from Lemma 4. In particular, we cannot hope to claim the correctness of Lemma 3, our current goal, directly from the correctness of the single step result Lemma 4, at least without providing additional arguments.

There is still the possibility that Lemma 5 is asymptotically true as N is sent to infinity, but the focus of our argument here is that multiple iteration statement is not a direct consequence of the single iteration statement.


A natural attempt at fixing the current situation would be to relax the strict correlation between the two functions that are being composed. This is considered next.

Lemma 6 (Incorrect) *Let $F : \mathcal{N} \rightarrow \mathcal{N}$ and $G : \mathcal{N} \rightarrow \mathcal{N}$ be two independent random functions on a finite set of size N . If $\mathcal{M} \subset \mathcal{N}$ is of size m_0 , then the size of $G(F(\mathcal{M}))$ is expected to be m_2 , where*

$$m_1 = N \left\{ 1 - \left(1 - \frac{1}{N} \right)^{m_0} \right\} \quad \text{and} \quad m_2 = N \left\{ 1 - \left(1 - \frac{1}{N} \right)^{m_1} \right\}.$$

This second version for the double iteration case seems structurally much simpler than our previous attempt, given as Lemma 5. One might be tempted to say that this version is a trivial consequence of the single step result, stated by Lemma 4, but, once again, we can provide a counterexample.

We return to the situation of $F : \{0, 1\} \rightarrow \{0, 1\}$. The set of all possible double iterations of the four mapping can be visualized as follows.



When the input set \mathcal{M} is the complete domain $\{0, 1\}$, after separately counting the number of functions with image sizes one and two, the expected image size can be computed as

$$E_{F,G} \left[|G(F(\{0, 1\}))| \right] = \frac{12}{16} \cdot 1 + \frac{4}{16} \cdot 2 = \frac{5}{4}. \quad (4)$$

The corresponding value, as claimed by Lemma 6, is as was previously computed in (3). Since the two values disagree, it is clear that Lemma 6 cannot be true, at least in the strict sense. Once again, we can conclude that Lemma 3, even when relaxed to allow independent random functions at each iteration, is not a direct logical consequence of the single iteration result Lemma 4.

We have disproved the double iteration case Lemma 5 and even the seemingly simpler Lemma 6 with explicit counterexamples. On the other hand, we know from experience that Lemma 3 works quite well in accurately predicting the behavior of iterations done with specific functions. Let us attempt to prove the incorrect Lemma 6 directly so as to locate the source of this apparent contradiction.

The language pertaining to the selection approach to random functions will be used briefly. We start our examination of Lemma 6 by rewriting the conclusion we want to obtain as

$$E_{F,G \in \mathcal{N}^{\mathcal{N}}} \left[|G(F(\mathcal{M}))| \right] = N \left\{ 1 - \left(1 - \frac{1}{N} \right)^{E_{F \in \mathcal{N}^{\mathcal{N}}} [|F(\mathcal{M})|]} \right\}. \quad (5)$$

Here, the exponent on the right hand side given in terms of expectation is justified through Lemma 4. The left hand size expectation can be written according to its definition as follows.

$$\begin{aligned} E_{F,G \in \mathcal{N}^{\mathcal{N}}} [|G(F(\mathcal{M}))|] &= \frac{1}{(\mathbb{A}^{\mathbb{A}})^2} \sum_{F,G \in \mathcal{N}^{\mathcal{N}}} |G(F(\mathcal{M}))| \\ &= \frac{1}{\mathbb{A}^{\mathbb{A}}} \sum_{F \in \mathcal{N}^{\mathcal{N}}} \frac{1}{\mathbb{A}^{\mathbb{A}}} \sum_{G \in \mathcal{N}^{\mathcal{N}}} |G(F(\mathcal{M}))|. \end{aligned}$$

Now, for every fixed $F \in \mathcal{N}^{\mathcal{N}}$, the inner summation computes the image size that is expected, when a random function G is applied to the fixed input set $F(\mathcal{M})$. Hence, the above is equal to

$$\begin{aligned} \frac{1}{\mathbb{A}^{\mathbb{A}}} \sum_{F \in \mathcal{N}^{\mathcal{N}}} E_{G \in \mathcal{N}^{\mathcal{N}}} |G(F(\mathcal{M}))| &= \frac{1}{\mathbb{A}^{\mathbb{A}}} \sum_{F \in \mathcal{N}^{\mathcal{N}}} \mathbb{N} \left\{ 1 - \left(1 - \frac{1}{\mathbb{N}} \right)^{|F(\mathcal{M})|} \right\} \\ &= E_{F \in \mathcal{N}^{\mathcal{N}}} \left[\mathbb{N} \left\{ 1 - \left(1 - \frac{1}{\mathbb{N}} \right)^{|F(\mathcal{M})|} \right\} \right] = \mathbb{N} \left\{ 1 - E_{F \in \mathcal{N}^{\mathcal{N}}} \left[\left(1 - \frac{1}{\mathbb{N}} \right)^{|F(\mathcal{M})|} \right] \right\}. \end{aligned} \quad (6)$$

A moment of thought shows that this final form cannot be equal to the right hand side of (5). For example, given a few numbers, we know very well that the average of their inverses is not equal to the inverse of their average. The expectation operator does not commute with any nonlinear operator, and exponentiation with the base set to $\left(1 - \frac{1}{\mathbb{N}}\right)$ is certainly a nonlinear operator. Hence, we come to the conclusion that Lemma 6 cannot be true.

Our failure to connect the end of (6) to the right hand side of (5) resulted from the fact that the expectation operator does not commute with the exponentiation operator in general. But, there is one situation when the expectation operator does commute with a nonlinear operator. Namely, this is when the multiple numeric inputs being considered are all identical. For example, if we are given a few numbers, but the numbers are all the same, then the average of their inverses is trivially the inverse of their average. As an extension of this exceptional behavior, we can expect an approximate commutativity of operators when the condition of inputs being identical is slightly relaxed. That is, it is reasonable to believe that if the inputs being considered are very close to each other, then two values computed using the two orderings of the operators will be close to each other.

Let us briefly consider the binomial distribution $B(n, p)$ corresponding to n trials with probability of success p for each trial. We know that its standard deviation is $\sqrt{np(1-p)}$, so that the clustering of multiple test values around the mean value np is expected to intensify with the increase of n , at least when the distribution of test values are viewed relative to np , the expected data value. Hence, it is reasonable to believe that, with any natural distribution, the clustering of data around the expected value will become tighter, as the space that is being considered is enlarged.

In the current situation, the initial inputs to the expectation and nonlinear operators under consideration are the multiple image sizes $|F(\mathcal{M})|$, where the function F is made to run over all elements of $\mathcal{N}^{\mathcal{N}}$. These are certainly not equal to each other, but it is not unrealistic to believe that they will tend to cluster together as \mathbb{N} is increased.

Let us summarize the above discussion. When the set size \mathbb{N} is large, it is reasonable to expect the image sizes $\{|F(\mathcal{M})|\}_F$ to mostly cluster around the average $E_F[|F(\mathcal{M})|]$. Hence, we are justified in computing the expected image size of $\{G(F(\mathcal{M}))\}_{F,G}$ as if the random function G was simply given an input set of size $E_F[|F(\mathcal{M})|]$. Although Lemma 6 is not strictly true, we can believe it to be approximately true when \mathbb{N} is large.

Recall that our current goal is to justify Lemma 3. Since we know that it cannot directly be inferred from the truth of Lemma 4, the single step statement, we turn to providing a plausible justification for its use.

We have already mentioned that it was reasonable to believe the image sizes of a single application of F to be clustered when N is large. Let us take this belief to the extreme and make the unrealistic assumption that the relation of Lemma 4, i.e.,

$$\text{image size} = N \left\{ 1 - \left(1 - \frac{1}{N} \right)^{\text{input size}} \right\}, \quad (7)$$

concerning the input and image sizes of a single iteration, holds exactly true, not just on average, but for every choice of the function F and input set \mathcal{M} . Under this unrealistic assumption, the iterative relation of Lemma 3 would be true for each function F , over any number of iterations. This automatically implies that the iterative relations are true when averaged over all functions. In this argument, we are not requiring for the functions that are being composed to be independent from each other.

The use of Lemma 3 as an approximation will now be justified, if relaxing our unrealistic belief does not cause trouble. We first consider the discouraging extreme situations. When the function F is bijective, so that (7) is far from accurate for F , it is clear that this error will persist throughout all iterations. That is, the error is not dampened through the multiple iterations. At the other extreme is the case when the function F is the constant function. The situation given by the iteration of this function is not reflected by (7) in any way. Hence the the iterative relations appearing in Lemma 3 does not hold true for individual functions that are extremely far away from our unrealistic assumption. But it should be noted that the number of these extreme functions are very small compared to the number of all functions.

For the vast majority of the functions, relation (7) may not be exactly true for every input set \mathcal{M} , but will be such that (7) overestimates the image size for some input sets and also underestimates for other input sets. Now, suppose that for some F and \mathcal{M} , the size of $F(\mathcal{M})$ is much smaller than what is estimated by (7). Note that it is unlikely for the points of the image set $F(\mathcal{M})$ to be somehow related to the points of \mathcal{M} . In other words, there is no reason to believe application of F to the image set $F(\mathcal{M})$ will produce a second image that is also smaller than would be estimated by (7), on a set of size $|F(\mathcal{M})|$. This lack of relationship between a subset of \mathcal{N} viewed as an image set and the same subset viewed as an input set to the next iteration assures us that a sequential buildup of any local abnormal behavior is unlikely to happen. Hence, even if (7) is only approximately true and no longer strictly true for every input to a certain specific function, the relations of Lemma 3 will remain approximately true for the same specific function.

We have shown that the validity of a statement concerning a single step of a random function does not directly imply the validity of the same statement under multiple iterations. There are many works in the literature, especially those concerning time memory tradeoffs, that ignore the logical gap that we have discussed. These arguments are typically written in the language of classical occupancy problems, but are in fact random function arguments that ignore behavior under iterations. In the opposite direction, we have argued that results for a single iteration of a random function may be generalized to multiple iterations when the space that is being dealt with is very large and the statements are taken as good approximations.

The intension of this subsection was not in testing the validity of Lemma 3. In fact, at least for the case when \mathcal{M} is the full domain, although inaccessible to the current author, a full proof is provided by [5]. When the validity of Lemma 3 is firmly taken for granted, our long discussion makes the following claim quite plausible.

Lemma 7 (Belief) Let $F : \mathcal{N} \rightarrow \mathcal{N}$ be any usual function acting on a finite set of size N . If $\mathcal{M} \subset \mathcal{N}$ is of size m_0 , then the size of the iterated image $F^k(\mathcal{M})$ will be approximately m_k , where m_k is defined recursively through

$$\frac{m_{i+1}}{N} = 1 - \exp\left(-\frac{m_i}{N}\right) \quad (i = 0, \dots, k-1).$$

Here, a usual function is one that has been chosen or constructed without making any intentional decisions that are likely to affect its iterated image sizes in a predictable way.

In short, we are claiming that Lemma 3 is likely to hold true quite accurately for most functions and not just on average. We emphasize that this is a *belief*, which we shall freely use, rather than a mathematically proven fact. In fact Lemma 3 is a statement concerning an average behavior, and, even if we had provided a strictly mathematically correct proof, it cannot imply anything about the behavior of each individual function.

Recall that much of the discussion in this subsection was about whether we may interchange the application order of the expectation operator and a nonlinear operator. We now have reasons to believe that the order of application may be changed without introducing big errors when the space being dealt with is large. Throughout this paper, our *random function arguments* will carelessly disregard the orders in which the operators are applied.

4 Applying Time Memory Tradeoff to Password Hashes

One usually states the objective of a tradeoff algorithm as the inversion of a one-way function. A closer look reveals that there are two versions of the inversion problem and we will explain how one of these corresponds to the application of tradeoff technique on password systems. Issues concerning the use of random functions in the theoretic analysis of tradeoff algorithms are also discussed in this section.

4.1 Password hash

Let us briefly explain how the security features of many file formats that rely on passwords for access control work in its very basic form.

The designer of the system chooses and fixes a one-way function H . This one-way function is a part of the file format specification and is usually considered to be public, and can always be extracted from the related software even if it was not originally made public. When the owner of a file following this format wants access control to be applied to the file, the user supplies a password \mathbf{x} . An encryption key is derived from the password, and the main content of the file is replaced by its encryption under this key. The image $\mathbf{y} = H(\mathbf{x})$ of the user password, under the one-way function specified for the file format, is recorded within the file. Finally, any record of the encryption key and the raw password supplied by the user is destroyed.

Later, when authentication is required for file access, the supporting software asks for a password. The one-way function image $H(\mathbf{x}')$ of the newly supplied password \mathbf{x}' is computed by the software and is compared with the corresponding information \mathbf{y} stored within the file. If a perfect match $\mathbf{y} = H(\mathbf{x}')$ is found, the main body of the file is decrypted using the key derived from the password and access to the decrypted content is granted. Note that the one-way function image \mathbf{y} of the correct password is stored within the file without any protection and is accessible to anyone that has obtained the file.

User authentication procedure for system login works in much the same way. At the time of initial user registration to the system, the one-way function image of a password supplied by the user is recorded in a file that is stored within the system. In this case, access to one-way function images may be harder, but this information is often sent over the network in the clear to allow for logins to a group of computers after user registration at a central server.

As we have stated earlier, in this work, we shall refer to the one-way function image as the *password hash* and the input as the *password*, regardless of whether the one-way function to be attacked through the tradeoff technique is related to a password authentication system.

4.2 Uniqueness of the pre-image to a password hash

Our first question is whether a password hash uniquely determines the password.

Proposition 8 *Let $H : \mathcal{P} \rightarrow \mathcal{H}$ be a random function. Given any $\mathbf{x} \in \mathcal{P}$, the number of inputs that H maps to $H(\mathbf{x})$ is expected to be $1 + \frac{|\mathcal{P}|-1}{|\mathcal{H}|}$.*

Proof When using the construction approach to a random function, we are free to choose the order in which function value assignments are made to each domain element. So let us first assign a randomly chosen value of \mathcal{H} to $H(\mathbf{x})$ and then define all other function values.

The probability for any one of the later assignments to strike $H(\mathbf{x})$, which is an explicitly fixed value in \mathcal{P} , is $\frac{1}{|\mathcal{H}|}$. Each later assignment is independent of all other assignments, and we can expect the number of later assignments to $H(\mathbf{x})$ to be $\frac{|\mathcal{P}|-1}{|\mathcal{H}|}$. \square

Readers should not misinterpret the above proposition as giving the pre-image size of a random $y \in \mathcal{H}$ under a random H . For a random function H , the distribution on \mathcal{H} produced by $H(x)$, as password x runs over the password set \mathcal{P} randomly, is the uniform distribution, and every $y \in \mathcal{H}$ is expected to have $\frac{|\mathcal{P}|}{|\mathcal{H}|}$ -many pre-images, rather than $1 + \frac{|\mathcal{P}|-1}{|\mathcal{H}|}$. This is not in contradiction with the proposition, as it essentially deals with the distribution on \mathcal{H} produced from random inputs by the specific H that has been constructed, and this is different from the uniform distribution on \mathcal{H} . Those points of \mathcal{H} that lie outside $H(\mathcal{P})$, for the specifically constructed H , do not have any chance of appearing.

One can also ask for the pre-image size of a random password hash $\mathbf{y} \in H(\mathcal{P})$. Note that this question can only be asked after the random function H has fully been constructed. The corresponding answer will depend on the size of $H(\mathcal{P})$, but should be close to

$$\frac{|\mathcal{P}|}{E(|H(\mathcal{P})|)} \approx \frac{1}{1 - \frac{1}{e}} \approx 1.582.$$

Once again, this question is not related to the content of the above proposition. The current question deals with the uniform distribution on $H(\mathcal{P})$, which is different from the distribution on $H(\mathcal{P})$ given by the fully specified H . Those points with larger pre-image sets will have a larger probability of appearing than those with smaller pre-image sets.

Consider an application of the tradeoff technique to a blockcipher whose key length is identical to the block length. In such a case, one is working with $|\mathcal{P}| = |\mathcal{H}|$ and Proposition 8 states that there will be approximately two keys, on average, that map to a given target ciphertext. This is probably larger than what many would have naively expected. Of course, in practice, one usually assumes the use of a second ciphertext to almost uniquely identify

the key. In fact, if one interprets the key to two-ciphertexts mapping as a new one-way function, then Proposition 8 claims that the key is almost always uniquely determined from the two ciphertexts.

Let us next discuss what Proposition 8 implies for systems that rely on passwords for access control. These systems are usually designed so that the space \mathcal{H} of potential hash values is much larger than the space \mathcal{P} of admissible passwords. A typical password hash would be a bit string of at least 128 bits in length and the number of alphanumeric passwords consisting of ten characters is only $62^{10} \approx 2^{59.5}$. In such a case, Proposition 8 shows that a password hash $H(\mathbf{x})$, produced from a password \mathbf{x} , will almost always identify \mathbf{x} uniquely.

Furthermore, in practice, the set of all passwords admissible by the security system is not very important. Note that human chosen passwords are most certainly not uniformly distributed within the complete admissible password space. The tradeoff attacker first fixes a subset \mathcal{P} of all possible passwords and decides to be satisfied with recovering passwords that only lie within this subset. The size of this subset is determined by the computational power that the attacker can allocate to the pre-computation phase and should preferably cover the passwords that are most likely to be used. Under such a setting the password hash set \mathcal{H} is immensely larger than the set of passwords \mathcal{P} that is being considered and hence the password hash determines the password uniquely.

For the remainder of this paper, we assume that the target system for the application of the tradeoff technique is such that $|\mathcal{P}| \ll |\mathcal{H}|$, implying that the password hash uniquely determines the password.

4.3 The reduction function

The tradeoff technique requires the one-way function to be iterated. Since the range of the one-way function is usually larger than the domain, iteration is achieved by utilizing a *reduction function* $R: \mathcal{H} \rightarrow \mathcal{P}$. The role of the reduction function is to let a password hash be interpreted as another password. As any theoretic treatment of the tradeoff technique assumes $R \circ H$ to be a random function, let us check whether this is appropriate.

Proposition 9 *Let $|\mathcal{P}|$ be a divisor of $|\mathcal{H}|$, so that $\frac{|\mathcal{H}|}{|\mathcal{P}|}$ is an integer. Let $R: \mathcal{H} \rightarrow \mathcal{P}$ be any fixed function that is pre-image uniform in the sense that it is exactly $\frac{|\mathcal{H}|}{|\mathcal{P}|}$ -to-1. If $H: \mathcal{P} \rightarrow \mathcal{H}$ is a random function, then $R \circ H: \mathcal{P} \rightarrow \mathcal{P}$ is a random function.*

Proof In more precise terms, we want to show that the distribution on $\mathcal{P}^{\mathcal{P}}$, produced from the uniform distribution on $\mathcal{H}^{\mathcal{P}}$, through the mapping $H \mapsto R \circ H$, is the uniform distribution.

Let $F_0: \mathcal{P} \rightarrow \mathcal{P}$ be any specific function. It suffices to show that, after random construction of a function $H: \mathcal{P} \rightarrow \mathcal{H}$, we will find $R \circ H = F_0$ with probability $\frac{1}{|\mathcal{P}|^{|\mathcal{P}|}}$. It is clear that $\{R^{-1}(z)\}_{z \in \mathcal{P}}$ is a partition of \mathcal{H} into cells of size $\frac{|\mathcal{H}|}{|\mathcal{P}|}$. The event $F_0 = R \circ H$ will happen if and only if the value assigned as $H(x)$ belongs to the cell $R^{-1}(F_0(x))$, for every $x \in \mathcal{P}$. Since the size of $R^{-1}(F_0(x))$ is always $\frac{|\mathcal{H}|}{|\mathcal{P}|}$, and since the assignment to $H(x)$ is independent and random for every x , the probability of arriving at $F_0 = R \circ H$ is

$$\left(\frac{|\mathcal{H}|/|\mathcal{P}|}{|\mathcal{H}|}\right)^{|\mathcal{P}|} = \frac{1}{|\mathcal{P}|^{|\mathcal{P}|}},$$

as claimed. \square

Every application of the time memory tradeoff technique on a security system involves a specific one-way function $H : \mathcal{P} \rightarrow \mathcal{H}$ and there is no strictly logical reason to believe that the specific H will display the properties expected of a random function. Hence we first need to discuss if predicting the behavior of an explicit tradeoff implementation with arguments concerning random functions can be justified.

There can be two ways to resolve this problem. The first is to appeal to our intuition. When one ignores his knowledge of the inner working of the given specific function, it will seem as if the function is returning independently and randomly generated values to each given input. Hence, viewed from the outside, it looks as if the specific function is the random function of the construction sense. The second argument, which seems slightly more plausible, is that the one-way function used in the security system is in fact a function that has been selected from the pool of all functions. As we discussed earlier in Section 3.2, when the spaces involved are sufficiently large, unless we had chosen the one-way function in an unusual way, any property exhibited by a specific function will be close to the property averaged over all functions.

We have thus partly justified the use of random functions in place of specific one-way functions $H : \mathcal{P} \rightarrow \mathcal{H}$ when analyzing the behavior of time memory tradeoffs. What we have shown through Proposition 9 is that if we may treat the specific one-way function H as a random function, then the same can be done with the function $R \circ H : \mathcal{P} \rightarrow \mathcal{P}$. Hence, throughout this paper, while analyzing the behavior of time memory tradeoffs, we shall work with a random function $F : \mathcal{N} \rightarrow \mathcal{N}$.

The discussion of reduction functions now brings us to another logical gap that frequently appears in random function arguments, that is specific to the analysis of time memory tradeoffs. It is often the case that multiple tables are used in tradeoffs and any analysis of tradeoff properties will assume these tables to be independent. Although a different reduction function or a family of reduction functions is used with each table, it is not true that the tables are independent. In the language of the construction approach to random function, assignments made during computation of an earlier table will prevent later assignments to be made independently.

Suppose we try to analysis a time memory tradeoff properties that involve multiple tables under our simple assumption that $F : \mathcal{N} \rightarrow \mathcal{N}$ is a random function. Then this will require averaging over functions of a value that combines multiple figures that come from different tables and these multiple figures will be correlated to each other. This will be very complex and difficult to handle. Once again, since the domain and range spaces that are being considered are usually very large, we assume the interchange of the expectation operator and any nonlinear operator is possible and present analysis that averages over functions before combining the figures to arrive at the final expectation. This is equivalent to assuming that the random function used in computing multiple tables was not a single function, but that multiple independently constructed random functions were used for different tables.

4.4 Two versions of the inversion problem

We have already mentioned that we shall work in the situation $H : \mathcal{P} \rightarrow \mathcal{H}$ with sets of sizes $|\mathcal{P}| \ll |\mathcal{H}|$, so that a password hash almost always determines a unique password. We also know that any analysis of time memory tradeoff behavior is done with a random function $F : \mathcal{N} \rightarrow \mathcal{N}$, whose image does not uniquely determine the input. The two functions are related through the correspondence $H \mapsto F = R \circ H$.

Given $\mathbf{y} = H(\mathbf{x})$, the unique password \mathbf{x} is obtained through the tradeoff algorithm, which uses F , as follows. The tradeoff algorithm is given $R(\mathbf{y})$ to process and returns inputs $x \in \mathcal{P}$, such that $F(x) = R(\mathbf{y})$. This relation may be restated as $R(H(x)) = R(H(\mathbf{x}))$ and does not necessarily imply $x = \mathbf{x}$. Hence one tests whether the candidate password x satisfies $H(x) = H(\mathbf{x})$, outside of the normal tradeoff algorithm. Since an output of H uniquely determines the input, fulfillment of this test implies $x = \mathbf{x}$ and recovery of the correct password.

As discussed in the previous subsection, the number of inputs to F satisfying $F(x) = R(\mathbf{y})$ will only be two on average, and the number of such tests done outside the tradeoff algorithm will be very small. Hence, the cost of such tests may be ignored during complexity analysis.

During a tradeoff algorithm analysis, one usually does not mention anything about H or R , the source of the inversion problem, and simply assume the inversion target $\mathbf{y} = F(\mathbf{x})$ is given, for some function $F : \mathcal{N} \rightarrow \mathcal{N}$. In this work, the goal of the tradeoff algorithm will be to find *the* correct password \mathbf{x} , rather than *any* password, corresponding to the given \mathbf{y} . The *any* version may be useful when working to find the pre-image of a true hash function, but the *the* version is suitable when looking for the correct password to an access control mechanism.

Since it is logically impossible to distinguish between the many pre-images with only the \mathbf{y} information, our analysis will focus on whether \mathbf{x} is among the possibly multiple F -pre-images to \mathbf{y} , returned by the tradeoff algorithm. The determination of whether each returned value is *the* correct password is assumed to be done outside the tradeoff algorithm.

The small difference of looking for *the* pre-image rather than *any* pre-image implies that the tradeoff algorithm will succeed under different circumstances. The *the* version succeeds if and only if the correct password \mathbf{x} had appeared as an *input* to the one-way function F during the pre-computation phase, i.e., if \mathbf{x} is among the pre-computation matrix entries excluding the ending points. On the other hand, the *any* version succeeds if and only if the image $\mathbf{y} = F(\mathbf{x})$ had appeared as the function *output* during the pre-computation phase, i.e., if \mathbf{y} is among the pre-computation matrix entries excluding the starting points. The two approaches will show differences in properties such as success probability and online running time. Still, it should not be too difficult to tweak all the *the* version results given in this paper to work for the *any* version.

5 DP Tradeoff

An analysis of the DP tradeoff will be given in this section. We shall present a formula for computing the probability of success for the algorithm and also provide a tradeoff curve which takes the effects of false alarms into account. We also discuss the number of bits required to efficiently store the starting and ending point pairs.

If a chain is generated with a random function, with the chain length bound set to \hat{t} , the probability of not obtaining a DP chain will be $(1 - \frac{1}{\hat{t}})^{\hat{t}} \approx e^{-\hat{t}/t}$. Rather than successively generating more chains until we have m chains, we shall generate each table from

$$m_0 = \frac{m}{1 - e^{-\hat{t}/t}} \quad (8)$$

distinct starting points. Then we can expect to collect m chains that terminate at DPs. And in the limiting $\hat{t} \gg t$ case, which is of more practical interest, the two approaches will almost be the same.

All of our tradeoff algorithm analyses are done under the assumption that the one-way function is the random function.

5.1 Probability of success

Let us discuss how to compute the probability of success for a DP tradeoff under a given set of parameters. We shall first present general formulas connecting pre-computation and probability of success and then show how to compute these for specific parameter. Our first lemma is quite trivial.

Lemma 10 *The number of one-way function invocations required in either creating a DP chain or stopping at the \hat{t} -th iteration without having reached a DP is expected to be*

$$t(1 - e^{-\hat{t}/t}).$$

Proof It suffices to add the probabilities of having to compute the successive iterations. Since the next iteration is computed if and only if a DP has not yet been reached, the expected one-way function invocation count is

$$\sum_{i=1}^{\hat{t}} \left(1 - \frac{1}{t}\right)^{i-1} = t \left\{1 - \left(1 - \frac{1}{t}\right)^{\hat{t}}\right\},$$

which we can approximate to $t\{1 - \exp(-\frac{\hat{t}}{t})\}$. \square

In the above proof, we have implicitly assumed the one-way function to be a random function and computed the probability for the first i assignments to be non-DPs. A more exact analysis would additionally consider the possibility of the next assignment producing a previously assigned value. We have not done so as this seems to be a good approximation.

Clearly, the success rate of a tradeoff algorithm is intimately connected to the amount of pre-computation. So, let us first present a way to write down the pre-computation.

Proposition 11 *The pre-computation phase of the DP tradeoff is expected to require $D_{pc}N$ one-way function invocations, where the pre-computation coefficient is*

$$D_{pc} = D_{msr}D_{nt}.$$

Proof We know from Lemma 10 that each attempt at a DP chain creation is expected to require $t(1 - e^{-\hat{t}/t})$ one-way function invocations. On the other hand, the creation of a single DP table starts with $m_0 = \frac{m}{1 - e^{-\hat{t}/t}}$ chains. Together, this implies that the creation of a single DP table is expected to consume mt one-way function invocations, regardless of the \hat{t} value. Hence the total pre-computation requirement may be written as $mtl = mt^2 \frac{l}{t} = D_{msr}D_{nt}N$. \square

Note that the pre-computation cost of the DP tradeoff does not depend on the chain length bound \hat{t} .

The *coverage rate* D_{cr} of a DP table, is defined to be the expected number of distinct nodes that appear among the DP chains as inputs to the one-way function, divided by mt . Since we are taking starting points to be non-DPs, all of the nodes that are counted will be non-DPs. Note that the mentioned expectation is an average over the choice of one-way functions, in addition to the choice of starting points. In other words, the coverage rate is a certain expected value for the random function. Our next statement reduces the search for success rate to the computation of the coverage rate.

Proposition 12 *The success probability of the DP tradeoff is*

$$D_{ps} = 1 - e^{-D_{cr}D_{pc}}.$$

Proof If we are given $\mathbf{y} = F(\mathbf{x})$ as the inversion target, the DP tradeoff will succeed in recovering the correct answer \mathbf{x} , if and only if \mathbf{x} had appeared as one of the inputs to the one-way function during the creation of the DP table. As discussed earlier, this is not equivalent to asking for the appearance of \mathbf{y} among the output values. The objective of recovering *the correct*, rather than *any* inverse, corresponds to finding \mathbf{x} among the one-way function inputs.

By definition of the coverage rate, a single DP matrix is expected to contain $D_{cr}mt$ distinct nodes that were used as inputs to the one-way function. Hence the processing of a single table will fail in returning the correct answer with probability $(1 - \frac{D_{cr}mt}{N})$. The success probability of the complete DP tradeoff process is given by

$$D_{ps} = 1 - \left(1 - \frac{D_{cr}mt}{N}\right)^l \approx 1 - \exp\left(-D_{cr}\frac{mtl}{N}\right) = 1 - e^{-D_{cr}D_{pc}}.$$

As discussed at the end of Section 4.3, we confide that our treatment here of separate tables as being independent is not strictly correct. \square

If the creator of the inversion target $\mathbf{y} = F(\mathbf{x})$ has a sufficiently large storage capability, it may be possible for him to compute and collect the complete set of one-way function inputs used during the DP table creation, and choose a password \mathbf{x} that does not belong to this set. The DP tradeoff will always fail under such challenges. The above proof cannot capture this situation since the one-way function F was taken to be a random function while defining D_{cr} .

In practice, this implies that, for our analysis to be correct, the hidden answer \mathbf{x} needs to be chosen without reference to the DP tradeoff table. Note that this is not as strong a requirement as asking for the choice of \mathbf{x} to be random. The choice only needs to be unrelated to the structure of the DP matrices.

Within this subsection, all chains belonging to the DP matrix will be seen as having been aligned at the starting points, rather than at the ending points, and the starting point column will be referred to as the 0-th column.

The above expression for probability of success can only be put to use if we know how to compute the coverage rate. Our computation of the coverage rate will be done in two steps. Of the m_0 chains generated, only m will be DP chains, but we disregard this in the first step and count the number of new nodes added by each column of the complete matrix. Sum of these values gives us the total number of all distinct input entries. In the second step, we will count the number of nodes that belonged to chains not ending at DPs and subtract these from the total count.

Let us write m_j for the number of new nodes added by the j -th column. The m_0 value, stated by (8), conforms to this notation.

Lemma 13 *The number of new nodes added by each column satisfies the recurrence relation*

$$m_j = N \left\{ 1 - \exp\left(-\frac{m_{j-1}}{N}\right) \right\} \left(1 - \frac{1}{t}\right) \left\{ 1 - \frac{\sum_{i=0}^{j-1} m_i}{N(1-1/t)} \right\}.$$

Proof Suppose a node positioned in the $(j-1)$ -th column is old, in the sense that it has appeared in one of the 0-th through $(j-2)$ -th columns. Application of random function to this node will not result in a random element of \mathcal{N} , but a node that had appeared in one of the 1-st through $(j-1)$ -th columns. Hence when counting new nodes of the j -th column we need only consider the nodes of the j -th column that are assigned as images to new nodes of the $(j-1)$ -th column. Recalling Lemma 4, we write this as the $N\{1 - \exp(-\frac{m_{j-1}}{N})\}$ part appearing in the claimed equation.

Of the distinct entries that have appeared in the j -th column, that are not automatically old, we want to filter out the DPs. The previous count is restricted to the non-DPs by multiplying the $(1 - \frac{1}{t})$ factor.

Still, not all of these non-DPs are new nodes. Those that have appeared in previous columns are removed by multiplying $\{1 - \frac{\sum_i m_i}{N(1-1/t)}\}$. Notice that we have $N(1 - \frac{1}{t})$, rather than N , in the denominator, as we are dealing only with non-DPs at this point. \square

The next two lemmas are technical computation results. We first turn the recursive formula for m_j into a difference equation concerning a certain sum of m_j .

Lemma 14 *Let $\mu_i = \frac{m_i}{N(1-1/t)}$ and $\sigma_j = \sum_{i=0}^{j-1} \mu_i$. Then, σ_j satisfies the recursive formula*

$$\sigma_{j+1} - \sigma_j = \frac{m_0}{N} - \frac{1}{t}\sigma_j - \frac{1}{2}\sigma_j^2 \quad \text{with} \quad \sigma_0 = 0,$$

which is accurate up to modulo $O(\frac{1}{t^3})$.

Proof It is straightforward to rewrite the recursive formula of Lemma 13 in terms of the notation μ_j .

$$\mu_j = \left\{1 - \exp\left(-\left(1 - \frac{1}{t}\right)\mu_{j-1}\right)\right\} \left(1 - \sum_{i=0}^{j-1} \mu_i\right).$$

This may be rewritten once again as

$$\exp\left(-\left(1 - \frac{1}{t}\right)\mu_{j-1}\right) = 1 - \frac{\mu_j}{1 - \sigma_j} = \frac{1 - \sigma_{j+1}}{1 - \sigma_j}.$$

Now, by taking products of both sides over $j = 1, \dots, k$, we can find

$$\exp\left(-\left(1 - \frac{1}{t}\right)\sigma_k\right) = \frac{1 - \sigma_{k+1}}{1 - \sigma_1}.$$

We have thus arrived at a relation involving only the σ_k notation.

By expanding the exponential function in its Taylor series, we obtain

$$\sigma_{k+1} = 1 - (1 - \sigma_1) \left\{1 - \left(1 - \frac{1}{t}\right)\sigma_k + \frac{1}{2}\left(1 - \frac{1}{t}\right)^2 \sigma_k^2 - \dots\right\},$$

and we can modify the above into the difference equation

$$\sigma_{k+1} - \sigma_k = \sigma_1 - \left(\sigma_1 + \frac{1}{t} - \frac{\sigma_1}{t}\right)\sigma_k - \frac{1}{2}(1 - \sigma_1)\left(1 - \frac{1}{t}\right)^2 \sigma_k^2 + \dots.$$

Noting that the left hand side $\sigma_{k+1} - \sigma_k = \mu_k$ is of order $O(\frac{m}{N}) = O(\frac{1}{t^2})$, we remove every term on the right hand side of $O(\frac{1}{t^3})$ order. This may easily be done after noting that $\sigma_1 = \mu_0$ is $O(\frac{1}{t^2})$ and that σ_k is $O(\frac{mk}{N})$, which is at most $O(\frac{1}{t})$. The simplified equation is now

$$\sigma_{k+1} - \sigma_k = \mu_0 - \frac{1}{t}\sigma_k - \frac{1}{2}\sigma_k^2 + O\left(\frac{1}{t^3}\right).$$

It is clear that the initial condition $\sigma_1 = \mu_0$ may be replaced by $\sigma_0 = 0$, under this recursive formula. As a final tweak, we subtract $\frac{m_0}{N(t-1)}$, which is of $O(\frac{1}{t^3})$ order, from the constant term $\mu_0 = \frac{m_0}{N(1-1/t)} = \frac{m_0}{N} \left(1 + \frac{1}{t-1}\right)$, to arrive at the claimed formula. \square

Now that we have a difference equation, we can obtain σ_k through an application of the Euler method.

Lemma 15 *For each non-negative integer k , we have*

$$m_k \approx N(\sigma(k+1) - \sigma(k)),$$

where

$$\sigma(k) = \frac{\boxminus^2 - 1}{t} \frac{\exp(\boxminus \frac{k}{t}) - 1}{(\boxminus + 1) \exp(\boxminus \frac{k}{t}) + (\boxminus - 1)} \quad \text{with} \quad \boxminus = \sqrt{1 + \frac{2D_{msr}}{1 - e^{-t/t}}}.$$

Proof Let function $\sigma : \mathbf{R} \rightarrow \mathbf{R}$ be the unique solution to the differential equation

$$\frac{d}{dk} \sigma = \frac{m_0}{N} - \frac{1}{t} \sigma - \frac{1}{2} \sigma^2 \quad \text{and} \quad \sigma(0) = 0. \quad (9)$$

If one defines the sequence $\{\sigma_k\}_{k \geq 0}$ through the corresponding difference equation

$$\sigma_{k+1} - \sigma_k = \frac{m_0}{N} - \frac{1}{t} \sigma_k - \frac{1}{2} \sigma_k^2 \quad \text{and} \quad \sigma_0 = 0, \quad (10)$$

the the Euler method tells us that $\sigma(k)$, the evaluation of the function σ at the non-negative integer k , may be approximated by the sequence value σ_k . We may turn this the other way around to present approximate values of σ_k through the function evaluations $\sigma(k)$.

The unique solution to differential equation (9) is

$$\sigma(k) = \frac{2m_0 t}{N} \frac{\exp\left(\sqrt{1 + \frac{2m_0 t^2}{N}} \frac{k}{t}\right) - 1}{\left(\sqrt{1 + \frac{2m_0 t^2}{N}} + 1\right) \exp\left(\sqrt{1 + \frac{2m_0 t^2}{N}} \frac{k}{t}\right) + \left(\sqrt{1 + \frac{2m_0 t^2}{N}} - 1\right)}.$$

The form of $\sigma(k)$ stated by this lemma is obtained when $m_0 = \frac{m}{1 - e^{-t/t}}$ and $mt^2 = D_{msr} N$ are substituted.

Since the definition of σ_k given by (10) is identical to the approximate recursive relation of Lemma 14, we have

$$\sigma(k) \approx \sigma_k \approx \sum_{i=0}^{k-1} \mu_i, \quad \text{where} \quad \mu_i = \frac{m_i}{N(1 - 1/t)}.$$

This allows us to write

$$m_k \approx N\left(1 - \frac{1}{t}\right) (\sigma(k+1) - \sigma(k)) \approx N(\sigma(k+1) - \sigma(k)),$$

where the $\frac{1}{t}$ term removal is justifiable, as it is of strictly smaller order. \square

This completes the first step of the coverage rate computation. The coverage rate corresponds to the number of distinct non-DP nodes contained in just the DP chains, but the currently computed m_k includes all points contained in even the non-DP chains. We need to account for these nodes belonging to non-DP chain nodes. This is the second step to finding the coverage rate.

Proposition 16 *The coverage rate of a single DP table is expected to be*

$$D_{cr} = \frac{2}{e^{\hat{t}/t} - 1} \int_0^{\hat{t}/t} \frac{\exp(\boxtimes u) - 1}{(\boxtimes + 1)\exp(\boxtimes u) + (\boxtimes - 1)} \exp(u) du,$$

where $\boxtimes = \sqrt{1 + \frac{2D_{msr}}{1 - e^{-\hat{t}/t}}}$.

Proof To count the number of distinct non-DPs belonging to all DP chains, we need to subtract the number of all new points belonging to non-DP chains from $\sum_{i=0}^{\hat{t}-1} m_i$. Before doing this, we first need to consider whether these points may not also appear within the DP chains as new points. It is clear that any new node belonging to a non-DP chain cannot have appeared in a column previous to its position, as the node is assumed to be new. Furthermore, such a node cannot appear within the DP chains in the same column or any future columns, since it would then reach a DP before the chain length bound is exceeded. Hence new nodes belonging to non-DP chains do not appear within any DP chains, and we may safely remove all these new points without worrying about their possible contribution to DP chain coverage.

Now, let us count how many points belong to non-DP chains, each column at a time. We start with the 0-th column. Among all m_0 chains, even though we do not know ahead of time which ones they would turn out to be, there will be $m_0(1 - \frac{1}{t})^{\hat{t}}$ chains that do not reach a DP even after \hat{t} more iterations. Hence $m_0(1 - \frac{1}{t})^{\hat{t}}$ nodes among the m_0 nodes belonging to the 0-th column need to be removed from the count of new nodes. As for the 1-st column, we had focused on m_1 chains, but $m_1(1 - \frac{1}{t})^{\hat{t}-1}$ nodes among these will not reach a DP before exceeding chain length bound, and need to be removed. The general term is now clear.

The coverage rate of a single DP table can now be stated as

$$\frac{1}{mt} \sum_{k=0}^{\hat{t}-1} m_k \left\{ 1 - \left(1 - \frac{1}{t}\right)^{\hat{t}-k} \right\}.$$

Using Lemma 15, we can approximate this to

$$\begin{aligned} & \frac{1}{mt} \sum_{k=0}^{\hat{t}-1} \mathbf{N}(\sigma(k+1) - \sigma(k)) \left\{ 1 - \left(1 - \frac{1}{t}\right)^{\hat{t}-k} \right\} \\ &= \frac{\mathbf{N}}{mt} \frac{1}{t} \sigma(\hat{t}) + \frac{\mathbf{N}}{mt} \sum_{k=0}^{\hat{t}-1} \sigma(k) \left(1 - \frac{1}{t}\right)^{\hat{t}-k} \frac{1}{t} \\ &\approx \frac{\sigma(\hat{t})}{D_{msr}} + \frac{t}{D_{msr}} \exp\left(-\frac{\hat{t}}{t}\right) \sum_{k=0}^{\hat{t}-1} \sigma(k) \exp\left(\frac{k}{t}\right) \frac{1}{t} \end{aligned}$$

Since the coverage rate is of $O(1)$ order and the first term $\frac{\sigma(\hat{t})}{D_{msr}}$ is of $O(\frac{1}{t})$ order, we simply discard the first term, and the summation term can be approximated by the integration

$$\frac{t}{D_{msr}} e^{-\hat{t}/t} \int_0^{\hat{t}/t} \sigma(tu) \exp(u) du,$$

when $\frac{1}{t}$ is small. The claimed formula follows after substitution of $\sigma(tu)$, as given by Lemma 15, and some simplifications. \square

We state the $\hat{t} \gg t$ case separately for later use.

Proposition 17 *The expected coverage rate of a single DP table is approximately*

$$D_{cr} = \frac{2}{\sqrt{1 + 2D_{msr} + 1}},$$

when the chain length bound \hat{t} is sufficiently large.

Proof When the chain length bound \hat{t} is sufficiently large, almost all of the $m_0 \approx m$ chains that are generated will terminate with a DP, and hence the coverage rate may be computed as $\frac{1}{m\hat{t}} \sum_{i=0}^{\hat{t}-1} m_i$.

Based on Lemma 15, we may write

$$D_{cr} = \lim_{\hat{t} \rightarrow \infty} \frac{\sum_{i=0}^{\hat{t}-1} m_i}{m\hat{t}} \approx \lim_{\hat{t} \rightarrow \infty} \frac{N\sigma(\hat{t})}{m\hat{t}} = \lim_{\hat{t} \rightarrow \infty} \frac{2}{1 - e^{-\hat{t}/t}} \frac{e^{\boxtimes \hat{t}/t} - 1}{(\boxtimes + 1)e^{\boxtimes \hat{t}/t} + (\boxtimes - 1)},$$

where $\boxtimes = \sqrt{1 + \frac{2D_{msr}}{1 - e^{-\hat{t}/t}}}$. The limit can be simplified to what is claimed. \square

5.2 Time memory tradeoff curve

Our next goal is to summarize the ability of the DP tradeoff algorithm to balance storage size and online time in a single tradeoff equation.

We now visualize the chains of the DP matrix as having been aligned at the ending points. The online iterations for the processing of a single DP table are counted starting from the 1-st iteration. That is, checking whether the given inversion target $\mathbf{y} = F(\mathbf{x})$ is among the DPs in the DP table is referred to as the 1-st iteration.

Our first task is to find the probability for merges to occur between DP chains.

Lemma 18 *Fix a random function $F : \mathcal{N} \rightarrow \mathcal{N}$ and suppose that we are given a pre-computed DP chain of length $j \leq \hat{t}$, generated with F from a random non-DP starting point. If a second chain is generated with F from a random starting point, the probability for it to become a DP chain of length i and merge with the given pre-computed chain is*

$$\frac{t}{N} \left\{ \exp\left(\frac{\min\{i, j\}}{t}\right) - 1 \right\} \exp\left(-\frac{i}{t}\right).$$

Proof Within this proof, let us refer to the event of the second chain becoming a DP chain of length i and merging with the pre-computed chain simply as *the event*.

We first restrict ourselves to the $i \leq j$ case and fix notation for the two chains as follows.

$$\begin{array}{ccccccccccc} \mathbf{x}_0 & \rightarrow & \cdots & \rightarrow & \mathbf{x}_{j-i} & \rightarrow & \mathbf{x}_{j-i+1} & \rightarrow & \mathbf{x}_{j-i+2} & \rightarrow & \cdots & \rightarrow & \mathbf{x}_{j-1} & \rightarrow & \mathbf{x}_j \\ & & & & & & \mathbf{y}_0 & \rightarrow & \mathbf{y}_1 & \rightarrow & \mathbf{y}_2 & \rightarrow & \cdots & \rightarrow & \mathbf{y}_{i-1} & \rightarrow & \mathbf{y}_i \end{array}$$

The nodes $\mathbf{x}_0, \dots, \mathbf{x}_{j-1}$ are non-DPs and \mathbf{x}_j is a DP.

Let us consider all possible scenarios by which the event can occur. If the randomly chosen starting point \mathbf{y}_0 happens to be equal to \mathbf{x}_{j-i} , then the second chain will follow the first chain and the event surely will occur. On the other hand, if either \mathbf{y}_0 is one of the points $\mathbf{x}_0, \dots, \mathbf{x}_{j-i-1}, \mathbf{x}_{j-i+1}, \dots, \mathbf{x}_{j-1}$, or is a DP, then the event cannot occur. In the remaining case, i.e. when \mathbf{y}_0 is neither a DP nor any one of the points $\mathbf{x}_0, \dots, \mathbf{x}_{j-1}$, then the possibility of the event occurring remains. Furthermore, in this last case, we may freely set $F(\mathbf{y}_0)$ to a randomly chosen point of \mathcal{N} .

The above argument may now be repeated. If the randomly chosen $\mathbf{y}_1 = F(\mathbf{y}_0)$ is equal to \mathbf{x}_{j-i+1} , then the event occurs. If \mathbf{y}_1 is either a DP or one of the points $\mathbf{x}_0, \dots, \mathbf{x}_{j-i}, \mathbf{x}_{j-i+2}, \dots, \mathbf{x}_{j-1}$, then the event cannot occur. And if \mathbf{y}_1 is neither a DP nor one of the points $\mathbf{x}_0, \dots, \mathbf{x}_{j-1}$, then the event occurrence is yet undecided and we are free to define $\mathbf{y}_2 = F(\mathbf{y}_1)$ to a random point of \mathcal{N} .

Hence, when $i \leq j$, the probability for the event to occur may be written as

$$\frac{1}{N} + \left(1 - \frac{1}{t} - \frac{j}{N}\right) \frac{1}{N} + \left(1 - \frac{1}{t} - \frac{j}{N}\right)^2 \frac{1}{N} + \dots + \left(1 - \frac{1}{t} - \frac{j}{N}\right)^i \frac{1}{N},$$

which is equal to

$$\frac{1}{N} \frac{1 - \left(1 - \frac{1}{t} - \frac{j}{N}\right)^{i+1}}{1 - \left(1 - \frac{1}{t} - \frac{j}{N}\right)}.$$

Noting that $\frac{j}{N} \ll \frac{1}{t}$ and using $\left(1 - \frac{1}{t}\right)^{i+1} \approx \left(1 - \frac{1}{t}\right)^i \approx \exp\left(-\frac{i}{t}\right)$, we can approximate this as

$$\frac{t}{N} \left\{1 - \exp\left(-\frac{i}{t}\right)\right\}.$$

We can similarly work with the $i \geq j$ case. The event can occur only if the beginning random choices $\mathbf{y}_0, \dots, \mathbf{y}_{i-j-1}$ are made among non-DPs that are different from $\mathbf{x}_0, \dots, \mathbf{x}_{j-1}$. The probability for the event to occur is

$$\left(1 - \frac{1}{t} - \frac{j}{N}\right)^{i-j} \frac{1}{N} + \left(1 - \frac{1}{t} - \frac{j}{N}\right)^{i-j+1} \frac{1}{N} + \dots + \left(1 - \frac{1}{t} - \frac{j}{N}\right)^i \frac{1}{N},$$

which is approximately

$$\frac{t}{N} \left\{\exp\left(-\frac{i-j}{t}\right) - \exp\left(-\frac{i}{t}\right)\right\}.$$

The results for the cases $i \leq j$ and $i \geq j$ can now be combined and stated as claimed. \square

With the probability of alarms in our hands, we can compute the cost induced by false alarms.

Lemma 19 *The number of extra one-way function invocations induced by alarms is expected to be*

$$t \frac{D_{msr}}{1 - e^{-\hat{t}/t}} \left\{2 - 8e^{-\hat{t}/2t} + \left(5 + 3(\hat{t}/t) - \frac{1}{2}(\hat{t}/t)^2\right) e^{-\hat{t}/t} + e^{-2\hat{t}/t}\right\},$$

for each DP table.

Proof When the chains are generated from $m_0 = \frac{m}{1 - e^{-\hat{t}/t}}$ non-DP starting points, one can expect to collect

$$\frac{m}{1 - e^{-\hat{t}/t}} \left(1 - \frac{1}{t}\right)^{j-1} \frac{1}{t} \approx \frac{\frac{m}{t}}{1 - e^{-\hat{t}/t}} \exp\left(-\frac{j}{t}\right) \quad (11)$$

DP chains of length j . The probability of collision between the online chain and any one of these DP chains of length j , at the i -th iteration of the online phase, is given by Lemma 18. Here, the 1-st iteration deals with an online chain of length one, rather than zero, that starts at the unknown password and ends at the inversion target. The third component is the work required at each collision. If we take advantage of the fact that there is a chain length bound,

the maximum number of iterations required to deal with a collision between a pre-computed chain of length j and an online chain of length i is $\min\{\hat{t} - i + 1, j\}$.

Multiplying the three components and summing over all possible indices i and j , the expected number of iteration can be expressed as

$$\sum_{i=1}^{\hat{t}} \sum_{j=1}^{\hat{t}} \frac{\frac{m}{t}}{1 - e^{-\hat{t}/t}} \exp\left(-\frac{j}{t}\right) \cdot \frac{t}{N} \left\{ \exp\left(\frac{\min\{i, j\}}{t}\right) - 1 \right\} \exp\left(-\frac{i}{t}\right) \cdot \min\{\hat{t} - i + 1, j\}.$$

Replacing $\frac{i}{t}$ with u and $\frac{j}{t}$ with v , the above can be approximated by the integration

$$\frac{\frac{m^2}{N} t}{1 - e^{-\hat{t}/t}} \int_0^{\hat{t}/t} \int_0^{\hat{t}/t} \exp(-u) \exp(-v) \left\{ \exp(\min\{u, v\}) - 1 \right\} \min\left\{\frac{\hat{t}}{t} - u, v\right\} dv du,$$

when $\frac{1}{t}$ is small. The claimed value appears when this definite integral is computed. \square

Finally, we write the tradeoff curve for the DP tradeoff in a way that takes the extra cost of false alarms into account.

Theorem 20 *The time memory tradeoff curve for the DP tradeoff is $TM^2 = D_{tc} N^2$, where the tradeoff coefficient is*

$$D_{tc} = \left\{ (2D_{msr} + 1) - \frac{8D_{msr}}{e^{\hat{t}/2t}} + \frac{(5 + \frac{3\hat{t}}{t} - \frac{\hat{t}^2}{2t^2})D_{msr} - 2}{e^{\hat{t}/t}} + \frac{D_{msr} + 1}{e^{2\hat{t}/t}} \right\} \frac{1 - e^{-D_{cr}D_{pc}}}{1 - e^{-\hat{t}/t}} \frac{D_{pc}^2}{D_{cr}D_{msr}}.$$

Proof The i -th DP table is processed if and only if all previous tables do not return the correct inverse. The probability of such a failure is $(1 - \frac{D_{cr}mt}{N})^{i-1}$. The time required in processing a single table is the sum of hash invocation counts given by Lemma 10 and Lemma 19. Hence the expected total running time of the DP tradeoff may be written as

$$T = \sum_{i=1}^l \left(1 - \frac{D_{cr}mt}{N}\right)^{i-1} \left\{ (1 - e^{-\hat{t}/t}) + \frac{D_{msr}}{1 - e^{-\hat{t}/t}} \left(2 - \frac{8}{e^{\hat{t}/2t}} + \frac{5 + \frac{3\hat{t}}{t} - \frac{\hat{t}^2}{2t^2}}{e^{\hat{t}/t}} + \frac{1}{e^{2\hat{t}/t}}\right) \right\} t.$$

The summation index here appears only in the first multiplicative factor, and we can easily check that

$$\begin{aligned} \sum_{i=1}^l \left(1 - \frac{D_{cr}mt}{N}\right)^{i-1} &= \frac{N}{D_{cr}mt} \left\{ 1 - \left(1 - \frac{D_{cr}mt}{N}\right)^l \right\} \\ &\approx \frac{t}{D_{cr}D_{msr}} \left\{ 1 - \exp\left(-D_{cr} \frac{mtl}{N}\right) \right\} = \frac{1 - e^{-D_{cr}D_{pc}}}{D_{cr}D_{msr}} t. \end{aligned}$$

The running time can now be rewritten as

$$T = \frac{1 - e^{-D_{cr}D_{pc}}}{D_{cr}D_{msr}} \left\{ (1 - e^{-\hat{t}/t}) + \frac{D_{msr}}{1 - e^{-\hat{t}/t}} \left(2 - \frac{8}{e^{\hat{t}/2t}} + \frac{5 + \frac{3\hat{t}}{t} - \frac{\hat{t}^2}{2t^2}}{e^{\hat{t}/t}} + \frac{1}{e^{2\hat{t}/t}}\right) \right\} t^2.$$

Since the storage is $M = ml$, we have

$$\begin{aligned} TM^2 &= \frac{1 - e^{-D_{cr}D_{pc}}}{D_{cr}D_{msr}} \left\{ (1 - e^{-\hat{t}/t}) + \frac{D_{msr}}{1 - e^{-\hat{t}/t}} \left(2 - \frac{8}{e^{\hat{t}/2t}} + \frac{5 + \frac{3\hat{t}}{t} - \frac{\hat{t}^2}{2t^2}}{e^{\hat{t}/t}} + \frac{1}{e^{2\hat{t}/t}}\right) \right\} (mtl)^2 \\ &= \left\{ (1 - e^{-\hat{t}/t}) + \frac{D_{msr}}{1 - e^{-\hat{t}/t}} \left(2 - \frac{8}{e^{\hat{t}/2t}} + \frac{5 + \frac{3\hat{t}}{t} - \frac{\hat{t}^2}{2t^2}}{e^{\hat{t}/t}} + \frac{1}{e^{2\hat{t}/t}}\right) \right\} \frac{1 - e^{-D_{cr}D_{pc}}}{D_{cr}D_{msr}} D_{pc}^2 N^2, \end{aligned}$$

which is equal to what is claimed. \square

The following statement is an immediate consequence of the above theorem.

Corollary 21 *The time memory tradeoff curve for the DP tradeoff is $M^2T = D_{tc}N^2$ with*

$$D_{tc} = (2D_{msr} + 1)(1 - e^{-D_{cr}D_{pc}}) \frac{D_{pc}^2}{D_{cr}D_{msr}},$$

when the chain length bound \hat{t} is sufficiently large.

5.3 Efficient use of storage

The storage size M appearing in the tradeoff curves is the total number of starting and ending point pairs that need to be stored in the tradeoff tables. In practice, it is important to know the actual storage size, or the number of bits, required. Each starting and ending point pair can surely be stored in $2 \log N$ bits, but there are techniques that allow more efficient storage.

Below, we shall assume a suitable method of enumerating the elements of \mathcal{N} has been fixed and treat elements of \mathcal{N} as $\log N$ -bit integers. This enumeration is trivial when \mathcal{N} is the set of all bit strings of certain length, but may require a small amount of work when \mathcal{N} is given as the set of character passwords with certain complicated structures.

The most widely known technique for storage reduction, other than the trivial fact that the distinguished part of the ending point need not be recorded in the DP table, concerns the choice of starting points. One fixes the starting points³ to the integers 0 through $m - 1$, instead of using randomly chosen points. Then the storage of each starting point will require $\log m$ bits, rather than $\log N$ bits. If one wishes to use different starting points for each separate table, one can concatenate the table number to the above mentioned integers before using them as starting points. Note that any part that is common within each table need not be stored separately for each starting point. In view of what random functions are, employing such a deterministic way of choosing starting points will not cause any change in the behavior of the tradeoff.

The second method of reducing storage is called the hash table technique. To facilitate fast table lookups, the tradeoff tables are usually sorted on the ending point before being written to storage. Instead of performing the sorting operation, one fixes a hash function⁴ that maps elements of \mathcal{N} to $\log m$ -bit strings and records each starting and ending point pair at the position in the table addressed by the hash value of the ending point. There are various ways to deal with collision of addresses. Note that, since the address contains m bits of information, roughly m bits of the ending point can be removed with minimal loss of information. One advantage of this method, in addition to reducing storage, is that, whereas a lookup to a sorted table requires time logarithmic in table size, the hash table technique allows constant time table lookups.

Let us also explain a simplified form of the hash table technique, commonly referred to as the index file method. After sorting the table on the ending points, one focuses on the most significant $\{(\log m) - \varepsilon\}$ bits of each ending point, where ε is a small positive integer. For each integer $0 \leq i < \frac{m}{2^\varepsilon}$, we can expect to find 2^ε consecutive entries in the sorted table that have the first $\{(\log m) - \varepsilon\}$ bits of the ending point equal to integer i . One can now remove $\{(\log m) - \varepsilon\}$ bits of each ending point and provide an index file that points to the starting

³ To be more exact, this should be 0 through $m_0 - 1$, but we will be content with the approximate value.

⁴ The same term is used, but this is not necessarily a cryptographic hash function. At the extreme, even truncation to $\log m$ bits may be used.




Fig. 1 Index file technique (The sorted list on the LHS is transformed to the RHS list, which contains two less bits per entry.)

positions for each i value without losing any information. The number of entries contained in the index file is only $\frac{m}{2^r}$ and hence the additional storage required by its appearance can be ignored. An example is illustrated in Figure 1. In practice, the index file often contains the number of entries corresponding to each index value rather than the full physical addresses.

The two methods considered so far reduce storage requirement without any loss of information concerning each starting and ending point pair. This is not so with the third and final method we explain. In this method one simply truncates a part of the ending point before storage. After each iteration of the online phase, the current end of the online chain is truncated and compared with the truncated ending points of the table. The table lookups may now return a match even when a merge between the online chain and a pre-computation chain did not happen. Still, since we were already expecting false alarms, no new measure needs to be devised to deal with such a new type of false alarms. Aggressive ending point truncation will cause more frequent false alarms, hence it should be balanced with its effect in reducing storage.

The word truncation may give the impression that such a method is applicable only when the space \mathcal{N} consists of bit strings. On spaces that look different, any surjective map that is pre-image uniform, in the sense that the number of pre-images for each element in the range is identical, can serve as the truncation operation. In practice, password hashes are usually bit strings and one does not apply the reduction function at the end of any chain. In fact, DPs are usually defined using the password hashes, rather than the passwords produced through the reduction function.

The ending point truncation method seems to be known among many cryptographers, but no guideline as to how much of the ending point can be truncated can be found. Let us analyze the effects of ending point truncation on the running time of the online phase. Below, we assume that the ending points are truncated in such a way that two random points of \mathcal{N} , when truncated in the specified manner, will have probability $\frac{1}{r}$ of agreeing with each other. We shall express such a situation as having $\frac{1}{r}$ probability of truncated match. For example, if $\log t$ bits from the ending points were truncated with $D_{msr} = 1$, so that $(\log m + \log t)$ bits remain, then the truncated matches would happen with probability $\frac{1}{mt}$. When truncating ending DPs, one should truncate the random-looking part, rather than the distinguished part. Removal of the distinguished part can always be undone, and does not cause any loss of ending point information.

Lemma 22 *Let us assume the use of ending point truncation with the truncated match probability set to $\frac{1}{r}$. The number of extra one-way function invocations induced by false alarms related to ending point truncation is expected to be*

$$t \frac{1 - 2(\hat{t}/t) e^{-\hat{t}/t} - e^{-2\hat{t}/t}}{1 - e^{-\hat{t}/t}} \frac{mt}{r},$$

for each DP table.

Proof Consider a random function $F : \mathcal{N} \rightarrow \mathcal{N}$ and suppose that the first chain, generated with F and a random non-DP starting point, produced a DP chain of length $j \leq \hat{t}$. Now, suppose a second chain is generated with F from a random non-DP starting point. Let us compute the probability for the second chain to become a DP chain of length i and not merge with the first chain, but have the same truncated ending point as the first chain.

The first i nodes of the second chain must be chosen among non-DPs that are different from the j pre-ending points of the first chain. The i -th node chosen, when truncated, needs to agree with the truncated ending point of the first chain. Note that this agreement already requires the final point to be a DP. Thus the probability we aimed to write can be expressed as

$$\left(1 - \frac{1}{t} - \frac{j}{N}\right)^i \frac{1}{r} \approx \exp\left(-\frac{i}{t}\right) \frac{1}{r}. \quad (12)$$

Now, we can combine the number of DP chains of length j , as given by (11), together with the probability of non-merging truncated collision with such a chain, as given by (12), to write the cost of truncation related false alarms as

$$\sum_{i=1}^{\hat{t}} \sum_{j=1}^{\hat{t}} \frac{m}{1 - e^{-\hat{t}/t}} \exp\left(-\frac{j}{t}\right) \cdot \exp\left(-\frac{i}{t}\right) \frac{1}{r} \cdot \min\{\hat{t} - i + 1, j\}.$$

It now suffices to simplify this expression. Replacing $\frac{i}{t}$ with u and $\frac{j}{t}$ with v , the above can be approximated by the definite integral

$$\frac{mt^2}{1 - e^{-\hat{t}/t}} \frac{1}{r} \int_0^{\hat{t}/t} \int_0^{\hat{t}/t} \exp(-u) \exp(-v) \min\left\{\frac{\hat{t}}{t} - u, v\right\} dv du,$$

when $\frac{1}{t}$ is small. We arrive at the claimed value when this is explicitly computed. \square

Combining Lemma 10, 19, and 22, we know that the online processing of a single DP table requires

$$t (1 - e^{-\hat{t}/t}) \quad (13)$$

$$+ t \frac{D_{msr}}{1 - e^{-\hat{t}/t}} \left\{ 2 - 8e^{-\hat{t}/2t} + (5 + 3(\hat{t}/t) - \frac{1}{2}(\hat{t}/t)^2) e^{-\hat{t}/t} + e^{-2\hat{t}/t} \right\} \quad (14)$$

$$+ t \frac{1}{1 - e^{-\hat{t}/t}} \left\{ 1 - 2(\hat{t}/t) e^{-\hat{t}/t} - e^{-2\hat{t}/t} \right\} \frac{mt}{r}, \quad (15)$$

invocations of the one-way function. The first two terms (13) and (14) are both of order $\Theta(t)$. The total cost will remain of the same order if the third term (15), addressing the cost of false alarms related to ending point truncation, is of the same order. One can see that this is equivalent to requiring $\frac{mt}{r} = \Theta(1)$. Now, observe that when $\frac{mt}{r} = \Theta(1)$, increasing the truncated matching probability $\frac{1}{r}$ by a factor of two will have a significant undesirable effect on the total online running time, while the reduction in storage achieved by such an increase is only by a single bit per table entry. Hence we would like the term (15) to become a very small fraction of the sum of the terms (13) and (14). In other words, for each set of parameters t , m , and \hat{t} , it is advisable to choose r as small as possible, i.e., truncate as much as possible, under the condition that

$$\frac{mt}{r} \ll \frac{(1 - e^{-\hat{t}/t})^2 + D_{msr} \left\{ 2 - 8e^{-\hat{t}/2t} + (5 + 3(\hat{t}/t) - \frac{1}{2}(\hat{t}/t)^2) e^{-\hat{t}/t} + e^{-2\hat{t}/t} \right\}}{1 - 2(\hat{t}/t) e^{-\hat{t}/t} - e^{-2\hat{t}/t}}.$$

All our further discussion will assume such an approach has been taken so that Theorem 20, in particular, remains valid.

Let us summarize the discussion of this subsection. We shall be very rough and ignore all constants of order $\Theta(1)$. Sequential use of starting points allows each starting point to be recorded in approximately $\log m$ bits. One can truncate close to $\log t$ bits from each ending point with minimal effect on the online running time. Of the remaining $\log m + \log t$ bits of the ending point, we do not need to store the $\log t$ bits that is fixed through the distinguishing property. Furthermore, the hash table or index file technique allows us to remove almost $\log m$ more bits without any loss of information. In all, $\log m$ bits are required to store each starting point and a very small number of bits are required to store each ending point.

6 Hellman Tradeoff

In this section, we summarize facts about the Hellman tradeoff. Even though none of these have appeared in the literature in the current form, each fact, other than that concerning optimal use of storage, is either very easy to prove or is a straightforward consequence of the recent works [9, 10].

All of our tradeoff algorithm analyses are done under the assumption that the one-way function is the random function.

Our first two facts are rather easily obtainable from the definitions. The pre-computation effort of the Hellman tradeoff can be expressed as follows.

Proposition 23 *The pre-computation phase of the Hellman tradeoff requires $H_{pc}N$ one-way function invocations, where the pre-computation coefficient is*

$$H_{pc} = H_{msr}H_{mt}.$$

Proof Since a single Hellman table consists of m chains of length t , its creation requires mt one-way function invocations. The total pre-computation cost is $mtl = mt^2 \frac{l}{t} = H_{msr}H_{mt}N$. \square

The proof for the next statement is almost identical to that for Proposition 12, and we shall be very brief.

Proposition 24 *The success probability of the Hellman tradeoff is*

$$H_{ps} = 1 - e^{-H_{cr}H_{pc}}.$$

Proof The success probability expected after completely processing all l tables is

$$1 - \left(1 - \frac{H_{cr}mt}{N}\right)^l \approx 1 - \exp\left(-H_{cr} \frac{mtl}{N}\right) = 1 - e^{-H_{cr}H_{pc}}.$$

This argument ignores the fact that the multiple tables are not independent from each other for each randomly chosen function F . \square

We next show how to compute the coverage rate, so that the above expression for probability of success can be put to use.

Proposition 25 *The coverage rate of a single Hellman table is expected to be*

$$H_{cr} = \frac{\sqrt{2}}{\sqrt{H_{msr}}} \frac{e^{\sqrt{2H_{msr}}} - 1}{e^{\sqrt{2H_{msr}}} + 1}.$$

Proof The coverage rate of a single Hellman table is given in [10] as

$$H_{cr} = \frac{2N}{mt} \left(1 - \frac{m}{2N}\right) \frac{e^{\sqrt{mt^2/2N}} - e^{-\sqrt{mt^2/2N}}}{(\sqrt{2N/m} + 1)e^{\sqrt{mt^2/2N}} + (\sqrt{2N/m} - 1)e^{-\sqrt{mt^2/2N}}}.$$

Since $\frac{m}{2N} \ll 1$, we can approximate this by ignoring the second multiplicative factor. When the matrix stopping rule $mt^2 = H_{msr}N$ is applied, the approximation becomes

$$H_{cr} = \frac{2t}{H_{msr}} \frac{e^{\sqrt{2H_{msr}}} - 1}{t\sqrt{2/H_{msr}}(e^{\sqrt{2H_{msr}}} + 1) + (e^{\sqrt{2H_{msr}}} - 1)},$$

after some rearrangements. Noting that $(e^{\sqrt{2H_{msr}}} - 1) \ll t\sqrt{2/H_{msr}}(e^{\sqrt{2H_{msr}}} + 1)$, we approximate this once more to

$$H_{cr} = \frac{2t}{H_{msr}} \frac{e^{\sqrt{2H_{msr}}} - 1}{t\sqrt{2/H_{msr}}(e^{\sqrt{2H_{msr}}} + 1)},$$

which is equal to the claimed coverage rate. \square

The performance of the Hellman tradeoff is compactly expressed by the following time memory tradeoff curve.

Theorem 26 *The time memory tradeoff curve for the Hellman tradeoff is $M^2T = H_{tc}N^2$, where the tradeoff coefficient is*

$$H_{tc} = \left(1 + \frac{H_{msr}}{6}\right) (1 - e^{-H_{cr}H_{pc}}) \frac{H_{pc}^2}{H_{cr}H_{msr}}.$$

Proof The i -th Hellman table is processed if and only if all previous tables have failed in returning the correct inverse. The probability of such a failure is $(1 - \frac{H_{cr}mt}{N})^{i-1}$. The number of one-way function invocations required in processing a single Hellman table may be found in [9], and is $(1 + \frac{H_{msr}}{6})t$, with the cost of resolving alarms taken into account. Hence the expected total running time of the Hellman tradeoff may be written as

$$T = \sum_{i=1}^l \left(1 - \frac{H_{cr}mt}{N}\right)^{i-1} \left(1 + \frac{H_{msr}}{6}\right) t.$$

The summation index here appears only in the first multiplicative factor, and we can easily check that

$$\begin{aligned} \sum_{i=1}^l \left(1 - \frac{H_{cr}mt}{N}\right)^{i-1} &= \frac{N}{H_{cr}mt} \left\{1 - \left(1 - \frac{H_{cr}mt}{N}\right)^l\right\} \\ &\approx \frac{t}{H_{cr}H_{msr}} \left\{1 - \exp\left(-H_{cr} \frac{mt}{N}\right)\right\} = \frac{1 - e^{-H_{cr}H_{pc}}}{H_{cr}H_{msr}} t. \end{aligned}$$

The running time can now be rewritten as

$$T = \frac{1 - e^{-H_{cr}H_{pc}}}{H_{cr}H_{msr}} \left(1 + \frac{H_{msr}}{6}\right) t^2. \quad (16)$$

Since the storage is $M = ml$, we have

$$M^2T = \frac{1 - e^{-H_{cr}H_{pc}}}{H_{cr}H_{msr}} \left(1 + \frac{H_{msr}}{6}\right) (mt)^2 = \frac{1 - e^{-H_{cr}H_{pc}}}{H_{cr}H_{msr}} \left(1 + \frac{H_{msr}}{6}\right) H_{pc}^2 N^2,$$

which is equal to what is claimed. \square

Before continuing, we note that the time T , computed in the proof as (16), counts the number of one-way function computations, and includes the efforts for resolving false alarms. Since the number of table lookups will be smaller, we make this more explicit.

Lemma 27 *The full online processing of the Hellman tradeoff, that use the parameters m , t , and l , is expected to require*

$$t^2 \frac{1 - e^{-H_{cr}H_{pc}}}{H_{cr}H_{msr}}$$

lookups to the Hellman tables.

Proof The i -th Hellman table is processed if and only if all previous tables have failed in returning the correct inverse and processing of each table requires t table lookups. Hence, the expected total number of table lookups is

$$\sum_{i=1}^l \left(1 - \frac{H_{cr}mt}{N}\right)^{i-1} t = \frac{1 - e^{-H_{cr}H_{pc}}}{H_{cr}H_{msr}} t^2,$$

as claimed. \square

We have so far secured access to the pre-computation cost, the success probability, and the tradeoff performance of the Hellman tradeoff. It remains to discuss the efficient use of storage. The three approaches to storage reduction, discussed in Section 5.3, remain valid for Hellman tradeoffs and an analysis of the ending point truncation method is given below. The concept of probability of truncated match, explained for the DP tradeoff, will be carried over to the Hellman tradeoff.

Lemma 28 *Let us assume the use of ending point truncation with the truncated match probability set to $\frac{1}{r}$. The number of extra one-way function invocations induced by truncation related alarms is expected to be*

$$t \frac{mt}{2r},$$

for each Hellman table.

Proof Fix a random function $F : \mathcal{N} \rightarrow \mathcal{N}$ and suppose that we are given a pre-computed chain of length t , generated with F from a random starting point. Now consider a second chain generated with F from a random starting point. The probability for it to produce an alarm related to truncation, i.e., a truncated ending point match without a merge with the first chain, on the i -th iteration, is

$$\left(1 - \frac{1}{N}\right)^i \left(\frac{1}{r} - \frac{1}{N}\right) \approx \left(1 - \frac{i}{N}\right) \left(\frac{1}{r} - \frac{1}{N}\right) \approx \frac{1}{r}.$$

This is because the first i nodes of the second chain must be chosen among nodes that are different from the t pre-ending points of the first chain.

Taking account of all m pre-computed chains, the cost induced by the truncation related alarms can now be written as

$$\sum_{i=1}^t \frac{m}{r} \cdot (t - i + 1) \approx \frac{mt^2}{r} \sum_{i=1}^t \left(1 - \frac{i}{t}\right) \frac{1}{t}.$$

When $\frac{1}{t}$ is small, by replacing $\frac{i}{t}$ with u , the above can be approximated with the definite integral

$$\frac{mt^2}{r} \int_0^1 (1 - u) du,$$

which computes to $\frac{mt^2}{2r}$, as claimed. \square

Combining this with what we saw during the proof of Theorem 26, the total online time required to deal with a single Hellman table can be stated as

$$t + t \frac{H_{msr}}{6} + t \frac{mt}{2r}.$$

The first two terms are of order $\Theta(t)$. After reviewing the arguments concerning ending point truncation for the DP tradeoff, we see that it is advisable to use ending point truncation that truncates as much as possible, while satisfying the condition

$$\frac{mt}{2r} \ll 1 + \frac{H_{msr}}{6}.$$

Let us summarize the number of bits required to store each starting and ending point pair. We shall ignore all constants of $\Theta(1)$ order and be very rough. Each starting point requires $\log m$ bits. Ending points may be truncated so that slightly more than $\log m + \log t$ bits remain without visible side-effects on the online running time. The hash table method allows almost $\log m$ additional bits to be saved per ending point without any loss of information. In all, $\log m$ bits are required for each starting point and slightly more than $\log t$ bits are required for each ending point.

7 Rainbow Tradeoff

In this section, we summarize facts about the rainbow tradeoff. The contents appearing in the first half of this section are either very easy to prove or trivial extensions to ideas that have appeared before.

There are two ways of ordering the online phase processing of multiple rainbow tables. In the sequential approach, one fully processes one table before moving onto the next table. In the simultaneous approach, one searches through the same k -th column of all rainbow matrices before moving onto the next column. The simultaneous approach is more efficient in terms of the expected number of one-way function invocations. In practice, handling multiple tables simultaneously may increase the average table lookup time.

In this work, we assume that the l rainbow tables are processed with the simultaneous approach. The 1-st iteration refers to the searching of $\mathbf{y} = F(\mathbf{x})$ among the ending points of all l tables. The k -th iteration will require $(k-1) \cdot l$ invocations of the one-way function and l lookups to different tables. Columns of the rainbow matrices are numbered from the 0-th, containing the starting points, to the t -th, containing the ending points.

Our first claim is an easy consequence of the relation $mt = R_{msr}N$ that defines the notation R_{msr} .

Proposition 29 *The pre-computation phase of the Rainbow tradeoff requires $R_{pc}N$ one-way function invocations, where the pre-computation coefficient is*

$$R_{pc} = R_{msr}l.$$

Contents of the following lemma for $l = 1$ were already used in [9], but let us rewrite it here for easy reference. The first statement of this lemma is a trivial extension of a similar statement appearing in [12]. As the proof of the first statement, given in the appendix of [12], makes no mentioning of random functions, we rewrite it here within the framework

explained in Section 3 of the current paper. The two proofs are essentially the same at the core.⁵

Lemma 30 *The probability for the first k iterations of the online phase to fail is*

$$\prod_{i=1}^k \left(1 - \frac{m_{t-i}}{N}\right)^l,$$

where $m_0 = m$ and $\frac{m_{t+1}}{N} = 1 - \exp\left(-\frac{m_t}{N}\right)$. This product may be approximated by

$$\left\{ \frac{2N + m(t-k-1)}{2N + m(t-1)} \frac{2N + m(t-k-2)}{2N + m(t-2)} \right\}^l.$$

Proof The number of distinct nodes expected in each rainbow matrix column is given by the stated m_i . As fully discussed in Section 3.2, there are logical gaps that lead to this claim, but its use as a good approximation can still be justified. In passing, we remark that the different reduction functions used at each column do not remove the logical gap and they do not even provide independence of random function construction between columns.

We can now suppose that a specific one-way function F has been given, and that the rainbow matrix constructed from F contains m_i distinct nodes in the i -th column, for each $0 \leq i \leq t$. Let $\mathbf{y} = F(\mathbf{x})$ be the inversion target. The i -th iteration of the online phase succeeds if and only if the hidden answer \mathbf{x} is located within the $(t-i)$ -th column. Assuming that \mathbf{x} was chosen without reference to the rainbow matrix, the i -th iteration fails with probability $\left(1 - \frac{m_{t-i}}{N}\right)$, and all k iterations will fail with the stated probability. Once again, we have ignored the interdependence between columns.

The second statement is based on the approximation

$$\frac{m_i}{N} \approx \frac{1}{N/m + i/2}.$$

This is a very small generalization of Theorem 1 from [1], which treats the $m = N$ case. The proof there can easily be modified to fit the current statement. After rewriting this as

$$1 - \frac{m_{t-i}}{N} \approx \frac{2N + m(t-i-2)}{2N + m(t-i)},$$

the sequential cancelations within the product become visible, and we are left with the claimed approximation. \square

We can arrive at the next claim by substituting $k = t$ into the above lemma and appropriately approximating the outcome.

Proposition 31 *The success probability of the rainbow tradeoff is*

$$R_{ps} = 1 - \left(\frac{2}{2 + R_{msr}}\right)^{2l}.$$

Performance of the rainbow tradeoff is compactly expressed by the following theorem.

⁵ Simplification of the approximation given by Lemma 30, for the special case of $m = N$ and $l = 1$, results in the relation $\prod_{i=c-1}^{t-1} \left(1 - \frac{m_i}{N}\right) \approx \frac{c(c-1)}{t(t+1)}$. We acknowledge that this relation was used multiple times within [1] and that the authors of the paper are likely to have been aware of the statement given here.

Theorem 32 *The time memory tradeoff curve for the rainbow tradeoff is $M^2T = R_{tc}N^2$, where the tradeoff coefficient is*

$$R_{tc} = \frac{l^3}{(2l+1)(2l+2)(2l+3)} \left(\begin{array}{c} \{(2l-1) + (2l+1)R_{msr}\}(2 + R_{msr})^2 \\ -4\{(2l-1) + l(2l+3)R_{msr}\} \left(\frac{2}{2 + R_{msr}}\right)^{2l} \end{array} \right).$$

Proof Substituting $k = i - 1$ into Lemma 30, we know that the i -th iteration is processed with probability

$$\begin{aligned} & \left\{ \frac{2N + m(t-i)}{2N + m(t-1)} \frac{2N + m(t-i-1)}{2N + m(t-2)} \right\}^l \\ & \approx \left\{ \left(1 - \frac{m(i-1)}{2N + m(t-1)}\right) \left(1 - \frac{m(i-1)}{2N + m(t-2)}\right) \right\}^l \approx \left(1 - \frac{R_{msr} \frac{i}{t}}{2 + R_{msr}}\right)^{2l}. \end{aligned}$$

The probability of alarm associated with a single chain in a single rainbow matrix at the i -th iteration may be inferred from [9] to be $\frac{i+1}{N}$. Hence, the expected total running time of the rainbow tradeoff, with false alarms associated with all m rows taken into account, may be written as

$$\begin{aligned} T &= \sum_{i=1}^t l \left\{ (i-1) + (t-i+1) \frac{m(i+1)}{N} \right\} \left(1 - \frac{R_{msr} \frac{i}{t}}{2 + R_{msr}}\right)^{2l} \\ &\approx t^2 l \sum_{i=1}^t \left\{ \frac{i}{t} + \left(1 - \frac{i}{t}\right) R_{msr} \frac{i}{t} \right\} \left(1 - \frac{R_{msr} \frac{i}{t}}{2 + R_{msr}}\right)^{2l} \frac{1}{t}. \end{aligned}$$

This may be approximated by the definite integral

$$T = t^2 l \int_0^1 u \{1 + R_{msr}(1-u)\} \left(1 - \frac{R_{msr} u}{2 + R_{msr}}\right)^{2l} du,$$

which computes to

$$T = t^2 l \frac{\left[\begin{array}{c} \{(2l-1) + (2l+1)R_{msr}\}(2 + R_{msr})^2 \\ -4\{(2l-1) + l(2l+3)R_{msr}\} \left(\frac{2}{2 + R_{msr}}\right)^{2l} \end{array} \right]}{(2l+1)(2l+2)(2l+3)R_{msr}^2} \quad (17)$$

It now suffices to combine this with the storage size $M = ml$ and simplify to arrive at the claim. \square

It should be noted that the time T appearing in the above tradeoff curve gives the count of one-way function invocations and ignores table lookups.

Lemma 33 *The full online processing of the rainbow tradeoff is expected to require*

$$t l \frac{2 + R_{msr} - 2 \left(\frac{2}{2 + R_{msr}}\right)^{2l}}{(2l+1)R_{msr}}$$

lookups to the rainbow tables.

Proof At the start of proof to Theorem 32, we saw that the i -th iteration is processed with approximate probability

$$\left(1 - \frac{\mathbb{R}_{msr} \frac{i}{t}}{2 + \mathbb{R}_{msr}}\right)^{2l}.$$

Since each iteration requires l table lookups, it suffices to compute

$$\sum_{i=1}^t l \left(1 - \frac{\mathbb{R}_{msr} \frac{i}{t}}{2 + \mathbb{R}_{msr}}\right)^{2l} \approx tl \int_0^1 \left(1 - \frac{\mathbb{R}_{msr} u}{2 + \mathbb{R}_{msr}}\right)^{2l} du,$$

to arrive at the expected number of table lookups. \square

We now turn to the issue of efficient use of storage. The three approaches to storage reduction, discussed in Section 5.3, remain valid for rainbow tradeoffs and an analysis of the ending point truncation method is given below. The concept of probability of truncated match, used for the DP and Hellman tradeoffs, will also be carried over to the rainbow tradeoff.

Lemma 34 *Let us assume the use of ending point truncation with the truncated match probability set to $\frac{1}{r}$. The number of extra one-way function invocations induced by false alarms related to ending point truncation is expected to be*

$$t^2 l \frac{m}{r} \frac{-4 + 4l\mathbb{R}_{msr} + (2l+1)\mathbb{R}_{msr}^2 + 4\left(\frac{2}{2+\mathbb{R}_{msr}}\right)^{2l}}{(2l+1)(2l+2)\mathbb{R}_{msr}^2}$$

Proof For exactly the same reason as given in the proof of Lemma 28, the probability for a randomly generated second chain to produce a truncation induced alarm without merging with the first chain is

$$\left(1 - \frac{1}{\mathbb{N}}\right)^i \left(\frac{1}{r} - \frac{1}{\mathbb{N}}\right) \approx \left(1 - \frac{i}{\mathbb{N}}\right) \left(\frac{1}{r} - \frac{1}{\mathbb{N}}\right) \approx \frac{1}{r}.$$

After recalling the probability for the i -th iteration to be processed, and taking account of all the ml pre-computed chains, the expected online cost can be written as

$$\sum_{i=1}^t (t-i+1) \cdot \frac{ml}{r} \cdot \left(1 - \frac{\mathbb{R}_{msr} \frac{i}{t}}{2 + \mathbb{R}_{msr}}\right)^{2l}.$$

Replacing $\frac{i}{t}$ with u , the above can be approximated by the definite integral

$$\frac{mt^2 l}{r} \int_0^1 (1-u) \left(1 - \frac{\mathbb{R}_{msr} u}{2 + \mathbb{R}_{msr}}\right)^{2l} du,$$

when $\frac{1}{t}$ is small, and the claimed value appears when this is computed. \square

After reviewing the arguments concerning ending point truncation for the DP and Hellman tradeoffs, we can compare the value given by this lemma against the previously computed main online time (17) to conclude that it is advisable to use ending point truncation that truncates as much as possible, while satisfying the condition

$$\frac{m}{r} \ll \frac{\{(2l-1) + (2l+1)\mathbb{R}_{msr}\}(2 + \mathbb{R}_{msr})^2 - 4\{(2l-1) + l(2l+3)\mathbb{R}_{msr}\}\left(\frac{2}{2+\mathbb{R}_{msr}}\right)^{2l}}{(2l+3)\{-4 + 4l\mathbb{R}_{msr} + (2l+1)\mathbb{R}_{msr}^2 + 4\left(\frac{2}{2+\mathbb{R}_{msr}}\right)^{2l}\}}.$$

Note that the right hand side is of $\Theta(1)$ order so that the total online time remains of $\Theta(t^2)$ order when ending point truncation satisfies $\frac{m}{r} = O(1)$.

Let us summarize the number of bits required to store each starting and ending point pair. We shall ignore all constants of $\Theta(1)$ order and be very rough. Each starting point requires $\log m$ bits. Ending points may be truncated so that slightly more than $\log m$ bits remain without visible side-effects on the online running time. The hash table method allows most of the remaining $\log m$ bits to be removed from the ending point without any loss of information. In all, $\log m$ bits are required for each starting point and only a very small number of bits are required for each ending point.

8 Optimal Tradeoff Parameters

In this section, we shall find optimal parameter sets for each of the three tradeoff algorithms.

Let us present our initial arguments in terms of the Hellman tradeoff. The balance between time and memory achievable by the Hellman tradeoff is expressed by the curve $M^2T = H_{tc}N^2$. It is clear that the Hellman algorithm at parameters that bring about a smaller tradeoff coefficient H_{tc} will require less resources to run. In other words, tradeoff coefficient H_{tc} is a measure of the tradeoff efficiency, with a smaller value representing better tradeoff performance.

The tradeoff coefficient H_{tc} is determined by the parameters m , t , and l . It should first be noticed that a better tradeoff coefficient should always be achievable, if one decides to sacrifice the success probability for finding the correct password. Hence, any comparison between two Hellman tradeoff coefficients, achievable through two different sets of parameters, should be done under the condition that they produce the same success probability.

Arguments similar to the above may be made for the DP and rainbow tradeoffs. Hence, for each of the three algorithms, we shall work to find the smallest tradeoff coefficient achievable under a fixed requirement on the success rate. This is not yet a comparison between different algorithms, but only a study of optimal tradeoff coefficient for each separate algorithm. Such an analysis may certainly seem interesting in view of optimal usage of tradeoff algorithms, but can be of limited value in practice. Parameters achieving better tradeoff performance may require more pre-computation, and with large scale implementations of the tradeoff technique, lowering the pre-computation cost may be much more valuable than achieving better tradeoff performance. Our purpose of locating the optimal tradeoff parameters is for its use in the next section, where we compare between different algorithms.

8.1 DP tradeoff

The parameter set that achieves optimal DP tradeoff performance, under a fixed requirement on the probability of success, is given below.

Proposition 35 *Let $0 < D_{ps} < 1$ be any fixed value. The DP tradeoff, under any set of parameters m , t , l , and \hat{t} , that are subject to the relations*

$$mt^2 = 1.26453N, \quad l = 1.28007 \ln\left(\frac{1}{1 - D_{ps}}\right)t, \quad \text{and} \quad \hat{t} = 2.59169t,$$

attains the given value D_{ps} as its probability of success, and exhibits tradeoff performance corresponding to

$$D_{tc} = 5.49370 D_{ps} \left\{ \ln(1 - D_{ps}) \right\}^2,$$

as the four parameters are varied. Under any such choice of parameters, the number of one-way function invocations required for the pre-computation phase is

$$D_{pc} = 1.61869 \ln\left(\frac{1}{1 - D_{ps}}\right),$$

in multiples of N .

The three relations restricting the parameter choices give optimal parameters in the sense that no choice of m , t , l , and \hat{t} can lead to a tradeoff coefficient smaller than the above while achieving D_{ps} as its probability of success.

Proof Proposition 11 and 12 state the probability of success for DP tradeoffs as

$$D_{ps} = 1 - e^{-D_{cr} D_{pc}} = 1 - e^{-D_{cr} D_{msr} D_{nt}}.$$

Recalling the definition $D_{nt} = \frac{l}{t}$, this relation may equivalently be stated as

$$l = \frac{1}{D_{cr} D_{msr}} \ln\left(\frac{1}{1 - D_{ps}}\right) t. \quad (18)$$

Now, referencing Proposition 16, we know that the DP coverage rate $D_{cr} = D_{cr}[D_{msr}, \hat{t}/t]$ may be treated as a function of the two variables D_{msr} and $\frac{\hat{t}}{t}$. Hence, given any m , t , \hat{t} , and D_{ps} , if we set $D_{msr} = \frac{mt^2}{N}$ and $D_{cr} = D_{cr}[D_{msr}, \hat{t}/t]$, and also fix l through relation⁶ (18), then the DP tradeoff with these parameters will always achieve success probability of D_{ps} .

Keeping in mind that we may freely choose m , t , and \hat{t} , and still obtain any requested success probability, we now work to minimize the DP tradeoff coefficient

$$D_{tc} = \left\{ (2D_{msr} + 1) - \frac{8D_{msr}}{e^{\hat{t}/2t}} + \frac{(5 + \frac{3\hat{t}}{t} - \frac{\hat{t}^2}{2t^2})D_{msr} - 2}{e^{\hat{t}/t}} + \frac{D_{msr} + 1}{e^{2\hat{t}/t}} \right\} \frac{1 - e^{-D_{cr} D_{pc}}}{1 - e^{-\hat{t}/t}} \frac{D_{pc}^2}{D_{cr} D_{msr}},$$

as given by Theorem 20. After some regrouping of variables, we can rewrite this as

$$\begin{aligned} D_{tc} &= D_{imp} \left[D_{msr}, \frac{\hat{t}}{t} \right] \cdot (1 - e^{-D_{cr} D_{pc}}) (-D_{cr} D_{pc})^2 \\ &= D_{imp} \left[D_{msr}, \frac{\hat{t}}{t} \right] \cdot D_{ps} \{ \ln(1 - D_{ps}) \}^2, \end{aligned}$$

where

$$\begin{aligned} D_{imp} \left[D_{msr}, \frac{\hat{t}}{t} \right] &= \left\{ (2D_{msr} + 1) - \frac{8D_{msr}}{e^{\hat{t}/2t}} + \frac{(5 + \frac{3\hat{t}}{t} - \frac{\hat{t}^2}{2t^2})D_{msr} - 2}{e^{\hat{t}/t}} + \frac{D_{msr} + 1}{e^{2\hat{t}/t}} \right\} \\ &\times \frac{1}{(1 - e^{-\hat{t}/t})} \frac{1}{D_{cr}^3 \left[D_{msr}, \frac{\hat{t}}{t} \right]} \frac{1}{D_{msr}}. \end{aligned} \quad (19)$$

It is clear that, when the probability of success requirement is fixed, minimizing D_{tc} is equivalent to finding the minimum of $D_{imp}[D_{msr}, \hat{t}/t]$. Note that, even though $D_{msr} = \frac{mt^2}{N}$ and \hat{t}/t share the parameter t , since we are free to set m , t , and \hat{t} to any value, there are enough degrees of freedom, and we may treat D_{msr} and \hat{t}/t as independent variables when looking for the minimum of $D_{imp}[D_{msr}, \hat{t}/t]$.

⁶ Note that l must be set to an integer. Since the RHS of (18) is rather large, the error to the success probability, introduced by taking the nearest integer to the RHS value, will be very small.




Fig. 2 Tradeoff coefficient for DP tradeoff at fixed probability of success ($D_{imp} = \frac{D_{tc}}{D_{ps} \{\ln(1-D_{ps})\}^2}$)

After substituting $D_{cr}[D_{msr}, \hat{l}/t]$, as given by Proposition 16, into the right hand side of (19), we can use numerical methods to find its minimum. One discovers that the minimum value of $D_{imp} = 5.49370$ is obtained at $D_{msr} = 1.25453$ and $\hat{l}/t = 2.59169$. The claimed relation between l and t follows from (18), after evaluation of $\frac{1}{D_{cr}[D_{msr}, \hat{l}/t]D_{msr}}$ at $D_{msr} = 1.25453$ and $\hat{l}/t = 2.59169$. The final claim concerning the pre-computation cost follows from an evaluation based on Proposition 11. \square

The parameter set achieving minimum tradeoff coefficient for the DP tradeoff is visible through Figure 2. It plots $D_{imp} = \frac{D_{tc}}{D_{ps} \{\ln(1-D_{ps})\}^2}$, which is given by (19), as a function of variables D_{msr} and \hat{l}/t .

The tradeoff curve, as given by this proposition, allows us to say more about the tradeoff than the previously known rough curve of $M^2T \approx N^2$. Suppose that, for some fixed set of parameters, the success rate of the DP tradeoff is not too small, and suppose that one wishes to increase the success rate, to the extent that the failure rate becomes the square of its current value. Then, for optimal choice of parameters, the D_{ps} factor will change little and the $\{\ln(1-D_{ps})\}^2$ factor will increase by a factor of four. Hence, one must allow an increase in the online time by a factor of four or use twice the current storage. The proposition also shows that one must endure twice the pre-computation cost to achieve this aim. Of course, the simplest way of doing this would be to increase the number of tables by twice, while keeping all other parameters the same.

While the above result gives the parameters that achieves the optimal tradeoff performance, in practical applications, pre-computation is very costly and one is more likely to choose a sufficiently large \hat{l} , so as not to discard any of the pre-computed results.

Proposition 36 *Let $0 < D_{ps} < 1$ be any fixed value. When the use of $\hat{l} \gg t$ is assumed, the DP tradeoff, under any set of parameters m , t , and l , that are subject to the relations*

$$mt^2 = 0.562047N \quad \text{and} \quad l = 2.18614 \ln\left(\frac{1}{1-D_{ps}}\right)t,$$

attains the given value D_{ps} as its probability of success, and exhibits tradeoff performance corresponding to

$$D_{tc} = 7.01057 D_{ps} \{\ln(1-D_{ps})\}^2,$$




Fig. 3 Tradeoff coefficients for DP tradeoff with $\hat{t} \gg t$ and Hellman tradeoff at fixed probability of success

as the three parameters are varied. Under any such choice of parameters, the number of one-way function invocations required for the pre-computation phase is

$$D_{pc} = 1.22871 \ln\left(\frac{1}{1 - D_{ps}}\right),$$

in multiples of N .

The two relations restricting the parameter choices give optimal parameters in the sense that, as long as $\hat{t} \gg t$ is assumed, no choice of m , t , and l can lead to a tradeoff coefficient smaller than the above while achieving D_{ps} as its probability of success.

Proof The proof is almost identical to that of Proposition 35. The only difference is that we refer to Proposition 17 to view D_{cr} as a function of D_{msr} and obtain the tradeoff coefficient from Corollary 21. Through some regrouping of terms we can write

$$D_{tc} = \left(2 + \frac{1}{D_{msr}}\right) \frac{1}{D_{cr}^3} (1 - e^{-D_{cr} D_{pc}}) (D_{cr}^2 D_{pc}^2)$$

and by substituting D_{cr} into the appropriate part of D_{tc} , we obtain

$$D_{tc} = \left(2 + \frac{1}{D_{msr}}\right) \left(\frac{\sqrt{1 + 2D_{msr}} + 1}{2}\right)^3 D_{ps} \{\ln(1 - D_{ps})\}^2. \quad (20)$$

It suffices to minimize

$$D_{tmp}[D_{msr}] := \frac{D_{tc}}{D_{ps} \{\ln(1 - D_{ps})\}^2} = \left(2 + \frac{1}{D_{msr}}\right) \left(\frac{\sqrt{1 + 2D_{msr}} + 1}{2}\right)^3,$$

seen a function of the single variable D_{msr} . \square

In comparison to the previous optimal parameters that utilizes \hat{t} as a free variable, this version shows less efficient tradeoff performance, but requires less pre-computation. The behavior of the DP tradeoff coefficient with $\hat{t} \gg t$, under a fixed requirement for success rate is given as the left hand side graph of Figure 3. The point of minimum tradeoff coefficient is marked, together with the position corresponding to $D_{msr} = 1$.

8.2 Hellman tradeoff

We now turn to the Hellman tradeoffs. This is very similar to the DP tradeoff case that uses a sufficiently large \hat{t} .

Proposition 37 *Let $0 < H_{ps} < 1$ be any fixed value. The Hellman tradeoff, under any set of parameters m , t , and l , that are subject to the relations*

$$mt^2 = 2.25433 N \quad \text{and} \quad l = 0.598941 \ln\left(\frac{1}{1-H_{ps}}\right)t,$$

attains the given H_{ps} as its probability of success, and exhibits the tradeoff performance corresponding to

$$H_{tc} = 1.50217 H_{ps} \{\ln(1-H_{ps})\}^2,$$

as the three parameters are varied. Under any such choice of parameters, the number of one-way function invocations required for the pre-computation phase is

$$H_{pc} = 1.35021 \ln\left(\frac{1}{1-H_{ps}}\right)$$

in multiples of N .

The two relations restricting the parameter choices give optimal parameters in the sense that no choice of m , t , and l can lead to a tradeoff coefficient smaller than the above while achieving H_{ps} as its probability of success.

Proof Since the proof is similar to those of Proposition 35 and 36 we shall be more compact. Based on Proposition 23 and 24, we may claim the relation $l = \frac{1}{H_{cr} H_{msr}} \ln\left(\frac{1}{1-H_{ps}}\right)t$. Reference to Proposition 25 shows that the Hellman coverage rate $H_{cr} = H_{cr}[H_{msr}]$ may be seen as a function of $H_{msr} = \frac{mt^2}{N}$. Hence, given any m , t , and H_{ps} , we can set l to an appropriate value with which the Hellman tradeoff achieves success probability of H_{ps} .

We now work to minimize the Hellman tradeoff coefficient. The content of Theorem 26 may be rewritten as

$$H_{tc} = \left(\frac{1}{H_{msr}} + \frac{1}{6}\right) \frac{1}{H_{cr}^3} H_{ps} \{\ln(1-H_{ps})\}^2$$

and substitution of H_{cr} results in

$$H_{tc} = \left(\frac{1}{H_{msr}} + \frac{1}{6}\right) \left(\frac{\sqrt{H_{msr}}}{\sqrt{2}} \frac{e^{\sqrt{2H_{msr}}+1}}{e^{\sqrt{2H_{msr}}-1}}\right)^3 H_{ps} \{\ln(1-H_{ps})\}^2. \quad (21)$$

For a fixed success probability, it suffices to minimize

$$H_{tmp}[H_{msr}] := \frac{H_{tc}}{H_{ps} \{\ln(1-H_{ps})\}^2} = \left(\frac{1}{H_{msr}} + \frac{1}{6}\right) \left(\frac{\sqrt{H_{msr}}}{\sqrt{2}} \frac{e^{\sqrt{2H_{msr}}+1}}{e^{\sqrt{2H_{msr}}-1}}\right)^3, \quad (22)$$

which is a function of the single variable H_{msr} .

One can use numeric methods to identify the minimum value $H_{tmp} = 1.50217$, which appears at $H_{msr} = 2.25433$. The two remaining constants appearing in the proposition may now be obtained through appropriate evaluations. \square

The original Hellman tradeoff, which is set to use $mt^2 = N$ and $l = t$ attains a success probability of 57.68% and the tradeoff curve $M^2T = 0.7797N^2$, when the cost of false alarms are taken into account. In comparison, the choice of $mt^2 = 2.2543N^2$ and $l = 0.5160t$, recommended by Proposition 37, gives $M^2T = 0.6409N^2$, while achieving the same success rate. This is visible through the right hand side graph of Figure 3, where the two dots mark the two parameter choices we have discussed.

The price paid for this better tradeoff performance is the moderate increase in pre-computation from N to $1.1630N$. Indeed, after combining Proposition 24 and 25 into

$$H_{pc} = -\frac{\ln(1 - H_{ps})}{H_{cr}[H_{msr}]} = \ln\left(\frac{1}{1 - H_{ps}}\right) \frac{\sqrt{H_{msr}}}{\sqrt{2}} \frac{e^{\sqrt{2H_{msr}}} + 1}{e^{\sqrt{2H_{msr}}} - 1}, \quad (23)$$

one can check that the pre-computation $H_{pc}[H_{msr}]$ required under any fixed probability of success is an increasing function of H_{msr} . Hence, while any point that is situated to the left of the minimal point may not be optimal in view of tradeoff performance, it corresponds to less pre-computation. Depending on the available computational resources, one may choose to lower pre-computation cost rather than increase the tradeoff efficiency. On the other hand, increasing H_{msr} beyond the minimizing value 2.25433 will have bad effects on both the pre-computation and the tradeoff performance and should be avoided.

Let us briefly return to the DP tradeoff that only utilizes sufficiently large \hat{t} . By combining Proposition 12 and 17, we can write

$$D_{pc} = -\frac{\ln(1 - D_{ps})}{D_{cr}[D_{msr}]} = \ln\left(\frac{1}{1 - D_{ps}}\right) \frac{\sqrt{1 + 2D_{msr}} + 1}{2}, \quad (24)$$

and, as with the Hellman tradeoff, confirm that D_{pc} is an increasing function of D_{msr} . Since we know from Proposition 36 that the best performance is achieved at $D_{msr} = 0.562047$, the choice of $D_{msr} < 0.562047$ may be reasonable in view of lowering pre-computation cost, but using $D_{msr} > 0.562047$ should be avoided. In particular, the use of $D_{msr} = 1$ cannot be justified.

8.3 Rainbow tradeoff

The analyses of optimal parameters for the DP and Hellman tradeoffs were very similar. Rainbow tradeoff does not allow the same approach because we have less control over the parameter l . The number of tables l used with DP and Hellman tradeoffs are quite large and we had treated l as if it were a continuous variable. In the rainbow tradeoff case, the table count is usually a small integer and we must keep in mind that it takes only discrete values.

Let us start with a fixed number of tables l . For any given requirement on the success rate, we can rewrite Proposition 31 as

$$R_{msr} = 2 \left\{ \left(\frac{1}{1 - R_{ps}} \right)^{\frac{1}{2l}} - 1 \right\} \quad (25)$$

and understand this as a lower bound on R_{msr} that can be used. It is clear that increasing R_{msr} under a fixed l will increase the pre-computation cost $R_{msr}lN$. One can also work with the tradeoff coefficient R_{tc} , as provided by Theorem 32, to confirm that increasing R_{msr} under a fixed l will reduce the tradeoff efficiency. Hence, under any fixed l , the exact value of R_{msr} , suggested by (25), should be used to achieve the required success rate.




Fig. 4 Tradeoff coefficient of rainbow tradeoff as a function of success rate requirement at small number tables ($l = 1$: dotted, $l = 2$: dashed, $l = 3$: solid)

We can now treat R_{msr} as a function of the success rate requirement R_{ps} , for any fixed l . After substituting R_{msr} , as given by (25), into the tradeoff coefficient of Theorem 32, one can rewrite it as

$$R_{tc} = \frac{4l^3}{(2l+1)(2l+2)(2l+3)} \times \left[\begin{aligned} & \left\{ -(2l+3) + 2(2l+1) \left(\frac{1}{1-R_{ps}} \right)^{\frac{1}{2l}} \right\} \left(\frac{1}{1-R_{ps}} \right)^{\frac{1}{l}} \\ & + \left\{ (2l+1)^2 - 2l(2l+3) \left(\frac{1}{1-R_{ps}} \right)^{\frac{1}{2l}} \right\} (1-R_{ps}) \end{aligned} \right]. \quad (26)$$

For each fixed l , this is a function of the single variable R_{ps} . A plot of this is given as Figure 4 for table counts $l = 1, 2$, and 3 . The right hand side box is a magnified partial view of the left hand side box in logarithmic scale.

Recalling that a smaller tradeoff coefficient implies better tradeoff performance, one can clearly read from the figure that the use of $l = 1$ is optimal when the requirement for success rate is very low and that the use of successively higher number of tables becomes optimal as the success rate requirement is made more stringent. We have numerically solved for the explicit probabilities at which the transition to the next table count should be made and have recorded this in Table 1.

Let us briefly explain the content of the table with examples. Suppose one aims to achieve the success probability of 99.9% with the rainbow tradeoff. Since 0.999 sits between 0.998775 and 0.999314, it is optimal to use ten tables. If one is requested to set the probability of failure to $\frac{1}{2^7}$, we locate -7 between -6.17353 and -7.08171 and conclude that six tables would be optimal. To understand the other three columns of the table, let us focus on the row that sits between $l = 1$ and $l = 2$. The use of a single table with $R_{msr} = 1.87905$, or the use two tables at $R_{msr} = 0.785335$ will both result in the optimal tradeoff coefficient of $R_{tc} = 1.48026 = 2^{0.565848}$ and success rate 73.4166%.

Note that any given success rate requirement R_{ps} makes a certain number of tables l as optimal, and the l value fixes R_{msr} through (25). Since the tradeoff coefficient of Theorem 32 is already determined by l and R_{msr} , and since the relation (25) guarantees R_{ps} success rate, any parameter set satisfying the mentioned restriction will be optimal in view of the tradeoff coefficient. Let us gather what we have discussed in a proposition.

Table 1 Range of success probability requirements for which each table count l is optimal

l	R_{ps}	$\log_2(1 - R_{ps})$	$\log_2 R_{rc}$	$R_{msr}[R_{ps}, l \uparrow]$	$R_{msr}[R_{ps}, l \downarrow]$
1	0	0	0		
2	0.734166	-1.91140	0.565848	1.87905	0.785335
3	0.886651	-3.14116	2.08082	1.44688	0.874929
4	0.946562	-4.22600	2.88968	1.25878	0.884357
5	0.973305	-5.22729	3.41666	1.14577	0.873341
6	0.986146	-6.17353	3.79818	1.06812	0.856920
7	0.992618	-7.08171	4.09387	1.01079	0.839893
8	0.995992	-7.96295	4.33425	0.966542	0.823891
9	0.997795	-8.82486	4.53663	0.931326	0.809415
10	0.998775	-9.67274	4.71157	0.902658	0.796529
11	0.999314	-10.5104	4.86585	0.878902	0.785129
12	0.999614	-11.3404	5.00406	0.858929	0.775059
13	0.999782	-12.1649	5.12941	0.841927	0.766150
14	0.999877	-12.9850	5.24421	0.827299	0.758246
15	0.999930	-13.8020	5.35019	0.814594	0.751208
15	0.999960	-14.6163	5.44869	0.803466	0.744914

Proposition 38 Let $0 < R_{ps} < 1$ be any given fixed value. Locate the table count l from Table 1 that corresponds to the given R_{ps} and compute

$$R_{msr} = 2 \left\{ \left(\frac{1}{1 - R_{ps}} \right)^{\frac{1}{2l}} - 1 \right\}.$$

Then the rainbow tradeoff that uses the located l and any parameters m and t satisfying the relation

$$mt = R_{msr} \mathbf{N}$$

attains the given value R_{ps} as its probability of success. The tradeoff performance corresponding to

$$R_{rc} = \frac{l^3}{(2l+1)(2l+2)(2l+3)} \left(\frac{\{(2l-1) + (2l+1)R_{msr}\}(2 + R_{msr})^2}{-4\{(2l-1) + l(2l+3)R_{msr}\}(1 - R_{ps})} \right),$$

can be observed as m and t are varied under the restriction. With any such choice of parameters, the number of one-way function invocations required for the pre-computation phase is

$$R_{pc} = R_{msr} l,$$

in multiples of \mathbf{N} .

The choice of l through Table 1 and the single relation concerning m and t lead to optimal parameters in the sense that no choice of m , t , and l can result in a tradeoff coefficient smaller than the above while achieving R_{ps} as its probability of success.

To be strictly logical, one must also consider the possibility that allowing the multiple tables to be of different sizes may lead to better tradeoff coefficients. We have analyzed the case of three tables with the most general table sizes and came to the conclusion optimal tradeoff performance is achieved at equal sized tables. Our method of analyzing this possibility can probably be extended to larger number of tables, but the computations will be much more complicated than what was presented here. Since the examination of the 3-table case showed that we are not likely to gain anything from the more general analysis, we chose to work with equal sized tables. In comparison, for the case of perfect rainbow

tables, we have reasons to believe that this extra flexibility will bring about better tradeoff performance.

Finally, we want to provide an argument that is analogous to what was discussed at the end of Section 8.2. One can check that

$$R_{pc} = R_{msr} l = 2l \left\{ \left(\frac{1}{1 - R_{ps}} \right)^{\frac{1}{2l}} - 1 \right\} \quad (27)$$

is a decreasing function of l , for each fixed R_{ps} . Hence, use of an l count that is larger than what is suggested by Table 1 will decrease the pre-computation requirement at the cost of reduced tradeoff efficiency. This may be preferable in some situations. On the other hand, use of an l count that is smaller than the optimal count will have bad effects on both the pre-computation cost and tradeoff efficiency, and should be avoided.

9 Comparison of tradeoff performances

All the tools required for a fair comparison of performances between the tradeoff algorithms are now ready.

9.1 Conversion of the tradeoff coefficients to a common unit

Discussion of the previous section has made it clear that for any comparison of tradeoff algorithms to be fair, the algorithms must be made to present the same probability of success. One must also consider the pre-computation cost required by each algorithm, but this aspect will be considered later. We are also aware that the tradeoff coefficient is a measure of tradeoff performance. Hence let us assume that the DP, Hellman, and rainbow tradeoff algorithms display the respective tradeoff curves

$$M_D^2 T_D = D_{tc} N^2, \quad M_H^2 T_H = H_{tc} N^2, \quad \text{and} \quad M_R^2 T_R = R_{tc} N^2, \quad (28)$$

at the same success rate. We will discuss how to interpret the ratios $D_{tc} : H_{tc}$, $D_{tc} : R_{tc}$, and $H_{tc} : R_{tc}$ of the tradeoff coefficients as ratios of tradeoff performances.

Unit for T . Let us first consider the time variable T . The appearance of the time value T_D in a DP tradeoff curve signifies that there are parameters t_D , m_D , l_D , and \hat{t}_D with which the DP algorithm will display running time corresponding to T_D . To be more exact, the expected online execution time will be that consumed by $T_D = \Theta(t_D^2)$ invocations of the one-way function and at most $l_D = \Theta(t_D)$ table lookups. In comparison, the value T_H for the Hellman tradeoff corresponds to $T_H = \Theta(t_H^2)$ one-way function computations and, as testified through Lemma 27, table lookups of the same $\Theta(t_H^2)$ order.

Hence, even if we are working with two parameter sets for the DP and Hellman tradeoffs which leads to identical time $T_D = T_H$, the real-world execution time of the two algorithms will be different. For a fair interpretation of the tradeoff coefficient ratio $D_{tc} : H_{tc}$ as a ratio of tradeoff performances, the difference in the time units used by the two algorithms must be taken into account.

To continue the discussion, we recall the online time complexity of the rainbow tradeoff. In this case, we can expect the appearance of the time value T_R to call for $T_R = \Theta(t_R^2 l_R)$ one-way function computations and, according to Lemma 33, table lookups of order $\Theta(t_R l_R)$.

In both the DP and rainbow tradeoffs, the number of table lookups is of strictly smaller order than the number of one-way function computations. Hence, in these cases, we may ignore the time taken by table lookups and treat T_D and T_R as the count of just one-way function invocations. In the Hellman case, we can combine (16) and Lemma 27 to conclude that a T_H must be treated as T_H one-way function computations and $\frac{6}{6+H_{msr}}T_H$ table lookups.

Ignoring any issues concerning the storage count M for the moment, we can state that to compare the DP or rainbow tradeoff algorithm against the Hellman tradeoff, one should look at the ratios

$$D_{tc} : \left(1 + \frac{6}{6+H_{msr}} \frac{\text{single table lookup time}}{\text{single one-way function computation time}}\right) H_{tc}$$

and

$$\left(1 + \frac{6}{6+H_{msr}} \frac{\text{single table lookup time}}{\text{single one-way function computation time}}\right) H_{tc} : R_{tc}$$

rather than the raw ratio $D_{tc} : H_{tc}$ and $H_{tc} : R_{tc}$.

If the one-way function is computationally very heavy and all the pre-computed tables are to reside on the fast online memory during the online phase, then the table lookup time could be insignificant in comparison to the one-way function computation time. In such a case, the above ratios would essentially reduce to $D_{tc} : H_{tc}$ and $H_{tc} : R_{tc}$. On the other hand, if huge pre-computed tables are to be accessed through slow network storage and the one-way function is computationally very light, the above conversion of units will be necessary.

It is clear that when comparing the DP tradeoff against the rainbow tradeoff, no conversion of the time units is necessary. At least when the issue of the storage unit is ignored, the ratio $D_{tc} : R_{tc}$ is equal to the tradeoff performance ratio.

Unit for M . Let us now discuss the storage unit. In all of the three tradeoff algorithms, M represents the number of starting and ending point pairs that need to be stored, but the actual number of bits required to store each table entry will be different for different tradeoff algorithms. We saw through Section 5.3 that, for the DP tradeoff, slightly more than $\log m_D$ bits are required to store a single starting and ending point pair. On the other hand, slightly more than $\log m_H + \log t_H$ bits are required for the Hellman tradeoff, and the rainbow tradeoff requires slightly more than $\log m_R$ bits to store each table entry.

The implementation environment and tradeoff requirements will place the choice of suitable parameters into a certain range, and it is reasonable to assume that the parameters chosen to be used with each algorithm will be related by

$$\log m_D \approx \log m_H, \quad \log t_D \approx \log t_H, \quad \text{and} \quad \log m_R \approx \log m_D + \log t_D \approx \log m_H + \log t_H.$$

Some readers may object that our discussion on the number of bits required for each table entry makes $m_D = 2m_H$ a more reasonable choice, but this difference by a factor two is lost in the approximation when, as in the above assumption, their logarithm values are compared.

Given the same amount of physical storage, the number of table entries that can be stored by the DP tradeoff will be greater than the number of entries that can be stored by the Hellman tradeoff by a factor of

$$\frac{\log m_H + \log t_H}{\log m_D} \approx 1 + \frac{\log t_D}{\log m_D} \approx 1 + \frac{\log t_H}{\log m_H}.$$

Noting that the change in storage affects the tradeoff performance through a square factor, and ignoring effects of time unit differences for the moment, we can state that, to compensate for the storage unit differences, the ratio

$$D_{tc} : \left(1 + \frac{\log t_H}{\log m_H}\right)^2 H_{tc}$$

should be used instead of $D_{tc} : H_{tc}$ for comparison of tradeoff performances. Similarly, ignoring the time unit, since we have $\frac{\log m_R}{\log m_D} \approx 1 + \frac{\log t_D}{\log m_D}$, the ratio $D_{tc} : R_{tc}$ should be converted into

$$D_{tc} : \left(1 + \frac{\log t_D}{\log m_D}\right)^2 R_{tc}.$$

The remaining ratio $H_{tc} : R_{tc}$ requires no conversion to deal with storage units, since we have $\log m_H + \log t_H \approx \log m_R$.

Combined unit conversion. It now suffices to combine the two arguments concerning units of time and storage to give fair comparisons of different tradeoff algorithms.

Proposition 39 *Consider different tradeoff algorithms that are set to run with specific corresponding parameters. Suppose that the tradeoff coefficients for the algorithms are given as D_{tc} , H_{tc} , and R_{tc} . Then the tradeoff performance ratios between tradeoff algorithms are given by the ratios*

$$D_{tc} : \left(1 + \frac{6}{6 + H_{msr}} \frac{\text{single table lookup time}}{\text{single one-way function computation time}}\right) \left(1 + \frac{\log t_H}{\log m_H}\right)^2 H_{tc},$$

$$D_{tc} : \left(1 + \frac{\log t_D}{\log m_D}\right)^2 R_{tc},$$

and

$$\left(1 + \frac{6}{6 + H_{msr}} \frac{\text{single table lookup time}}{\text{single one-way function computation time}}\right) H_{tc} : R_{tc}.$$

In our further discussions below, we shall mainly work with parameter sets that roughly satisfy

$$\log m_D \approx \log m_H \approx \log t_D \approx \log t_H \approx \log t_R \approx \frac{1}{3} \log N \quad \text{and} \quad \log m_R \approx \frac{2}{3} \log N.$$

and also mostly assume that the table lookup time is negligible in comparison to the one-way function computation time. Under these assumptions, the ratios that need to be studied are

$$D_{tc} : 4H_{tc}, \quad D_{tc} : 4R_{tc}, \quad \text{and} \quad H_{tc} : R_{tc}.$$

Hence, it suffices to compare the values $\frac{1}{4}D_{tc}$, H_{tc} , and R_{tc} against each other.

We shall refer to the above situation as the *typical situation*, as it often appears during theoretic developments of the tradeoff technique, but we do not claim this to be typical in practical applications of the tradeoff technique.

We emphasize that our further discussions given below concerning tradeoff performance comparisons will only be valid under the typical environment assumption that was just explained. If the environment and tradeoff performance requirements make parameter choices such that $m_H \gg t_H$ or $m_H \ll t_H$ more appropriate, or if the table lookup time is not negligible in comparison to the one-way function computation time, the conclusions will be different. Still, one will be able to start from Proposition 39 and similarly discuss these other situations.

9.2 DP tradeoff versus Hellman tradeoff

The focus of this work is with practical uses of the tradeoff algorithms, and we shall restrict discussion of the DP tradeoff to the case when $\hat{t} \gg t$. As discussed in the previous subsection, it suffices to compare $\frac{1}{4}D_{tc}$ against H_{tc} for a fair comparison between the DP and Hellman tradeoffs. Note that we are assuming the typical situation explained at the end of the previous subsection and any conclusion we make could be different under different circumstances.

The contents of Proposition 36 and 37 show that the optimal performances of the two algorithms are given by

$$\frac{1}{4}D_{tc} = 1.75264 D_{ps} \{ \ln(1 - D_{ps}) \}^2 \quad \text{and} \quad H_{tc} = 1.50217 H_{ps} \{ \ln(1 - H_{ps}) \}^2.$$

One may be lead to believe that the Hellman tradeoff, with the smaller tradeoff coefficient, will be more efficient, but this is only true when the pre-computation is totally ignored. In practice, pre-computation cost is the largest barrier to any large scale deployment of tradeoff algorithms and hard to ignore.

The pre-computation costs required to achieve the above tradeoff performances are

$$D_{pc} = 1.22871 \ln\left(\frac{1}{1 - D_{ps}}\right) \quad \text{and} \quad H_{pc} = 1.35021 \ln\left(\frac{1}{1 - H_{ps}}\right).$$

The pre-computation cost of the DP tradeoff is seen to be lower and we are faced with the problem of how to compare low tradeoff performance at lower pre-computation against better tradeoff performance at higher pre-computation cost.

A moment of thought shows that such a comparison can not be objective. It is closely related to the relative value of the tradeoff performance against the pre-computation effort, and there is no unit with which to express either of these values. As an extension of this thought, one must question whether it is reasonable to compare the two tradeoffs at parameters giving their respective optimal tradeoff performances. Non-optimal parameters may be preferable under many situations in view of lower pre-computation cost.

We can conclude that all we can do is present the range of choices that can be made and allow the users to make their conclusions based on their situation. The crucial information that must be displayed to allows easy judgement of which tradeoff is more suitable is the relation between tradeoff performance and pre-computation cost. This must be done under each fixed requirement for success rate.

As was previously noted through (24) and (23), when under a fixed probability of success, both D_{pc} and H_{pc} are functions of their respective D_{msr} and H_{msr} values. The tradeoff coefficients D_{tc} and H_{tc} , under a fixed success rate requirement, were similarly expressed as functions of the corresponding D_{msr} and H_{msr} values in (20) and (21).

For a comparison of the DP tradeoff against the Hellman tradeoff, it now suffices to present the graphs

$$\{(D_{pc}[D_{msr}], \frac{1}{4}D_{tc}[D_{msr}]) \mid D_{msr} \leq 0.562047\} \quad (29)$$

and

$$\{(H_{pc}[H_{msr}], H_{tc}[H_{msr}]) \mid H_{msr} \leq 2.25433\}, \quad (30)$$

where the bounds on D_{msr} and H_{msr} were placed in accordance to the discussion at the end of Section 8.2. These graphs are given in Figure 5.

Since the two graphs are to be compared at identical success rate requirements $D_{ps} = H_{ps}$, we have removed the common parts that depend on the success probability from both of the




Fig. 5 The tradeoff coefficient $\frac{1}{4}D_{tc}$ (dotted) and H_{tc} (dashed) in relation to their respective pre-computation cost

cases before plotting the graphs. Hence, the graphs do not depend on the success rate and are valid for all success rate requirements. Both graphs extend further upwards, but the right ends, corresponding to the optimal tradeoff performances, are clearly marked with dots.

The two graphs are very close to each other. Even though slightly better tradeoff performance can be obtained with the Hellman tradeoff at higher pre-computation cost, in practice, unless parameters far from the typical $m \approx t \approx N^{\frac{1}{3}}$ region are to be used, the DP tradeoff will be favored in view of less number of table lookups. For example, if the table lookup time makes $\frac{1}{5}D_{tc} : H_{tc}$ a more appropriate measure of tradeoff performance ratio than the current $\frac{1}{4}D_{tc} : H_{tc}$, the dotted curve for the DP tradeoff would move down and present itself as a more advantageous algorithm.

If table lookup time is absolutely negligible in comparison to the one-way function computation time, there is a small range of parameters with which the Hellman tradeoff can slightly outperform the DP tradeoff using the same amount of pre-computation. If table lookup time is negligible and pre-computation is not to be considered, the Hellman tradeoff can be somewhat better.

9.3 Rainbow tradeoff versus DP and Hellman tradeoffs

As was discussed in Section 9.1, we assume the typical situation concerning the approximate range of parameters and table lookup time, and consider comparisons between $\frac{1}{4}D_{tc}$, H_{tc} , and R_{tc} to be fair.

In addition to the graphs (29) and (30), we need to plot all possible (R_{pc}, R_{tc}) points. We can first check through (27) that R_{pc} can be seen as a function of the table count l , when success rate requirement R_{ps} is fixed. As for the tradeoff coefficient, equation (26) presents it as a function of just l , when R_{ps} is fixed. Given any requirement on the success rate R_{ps} , it is now possible to draw the graph

$$\{(R_{pc}[l], R_{tc}[l]) \mid l \geq \text{optimal table count for } R_{ps}\}, \quad (31)$$

where the optimal table count can be obtained from Table 1. Note that this is no longer a continuous graph, but a discrete set of points. In the strict sense, previous graphs were also discrete set of points, but when N is large, the points will be extremely close to each other.

Unlike our comparison between DP and Hellman tradeoffs, the parts that depend on R_{ps} appearing in the expressions (27) and (26) are not identical to those appearing in the corre-




Fig. 6 Tradeoff coefficient $\frac{1}{4}D_{tc}$ (dotted), H_{tc} (dashed), and R_{tc} (large dots) in relation to their respective pre-computation cost at success rates 25%, 50%, 75%, 90%, 95%, and 99% (X-axis: D_{pc} , H_{pc} , and R_{pc} ; Y-axis: $\frac{1}{4}D_{tc}$, H_{tc} , and R_{tc})

sponding expressions (24), (23), (20), and (21). Hence separate graphs need to be drawn for each success rate. This is given in Figure 6 for some success rates.

In all of the graphs, one can see that the curve for the rainbow tradeoff sits closer to the origin than the curves for DP and Hellman tradeoffs. Note that a graph sitting lower shows better tradeoff performance and being positioned more to the left implies lower pre-computation cost. In all the cases except for the ones corresponding to 25% and 50% success rates, given any position on the curve for either the DP or Hellman tradeoff there is a rainbow tradeoff position that presents better tradeoff performance at a smaller pre-computation cost. Hence use of the rainbow tradeoff is definitely advisable in these cases.

The existence of better rainbow position is also mostly true in the 50% case. The exception is marked with an \otimes on the curve for the Hellman tradeoff. This position is very slightly to the left of the optimal rainbow position and hence corresponds to less pre-computation than the optimal rainbow position. At the same time, it is positioned lower than the second best rainbow position and hence shows better tradeoff performance than this second best position. Hence, there can be no rainbow tradeoff parameter set that can replace the Hellman position marked with an \otimes without at least very slightly sacrificing either the pre-

computation cost or the tradeoff efficiency. Still, anybody can agree that this exception is quite unreasonable and one would normally choose to sacrifice the extremely small amount of either the pre-computation cost or the tradeoff performance for a somewhat better value of the other factor.

The 25% case also displays the rainbow tradeoff requiring less pre-computation than the other two tradeoffs in achieving the equal tradeoff performance, but the awkward exceptional position discussed for the 50% can be found here as rather large segments. In addition, the best performance achievable by the rainbow tradeoff falls short of what is reachable by the other two algorithms. Hence there may be situations where the DP or Hellman tradeoffs may be preferable over the rainbow tradeoff, when required to achieve 25% success rate.

The relative advantage of using rainbow tradeoff is clearly seen to grow with the increase in the success rate requirement. For the 99% success rate case, it seems almost safe to say that the rainbow tradeoff is approximately two times more efficient than the other two tradeoffs in any of their reasonable usages.

In conclusion, the use of rainbow tradeoff is advisable for high success rate requirements and there may occasionally be low success rate applications with special situations where the other two tradeoffs are preferable. We emphasize once more that this conclusion is only valid under the typical situation assumption explained in Section 9.1. For example, if we must work with parameters such that $2 \log m_D \approx \log t_D$ and $2 \log m_H \approx \log t_H$, then comparison of the coefficients $\frac{1}{9}D_{TC}$, H_{TC} , and R_{TC} would be appropriate, which would bring the curve for the DP tradeoff lower, leading to different conclusions.

10 Conclusion

In the first part of this work, we solidified the basis on which analysis of tradeoff algorithms may be discussed. Logical gaps in common arguments were identified and plausible explanations for ignoring them were given. We next studied the performance of DP, Hellman, and rainbow tradeoffs, and summarized each as a tradeoff curve that is correct even up to the small multiplicative factor. These results were used in the last part of this work to compare the performance of tradeoff algorithms against each other.

Even though we did provide explicit statements comparing the three tradeoff algorithms, our conclusions are only true under a certain assumption on the tradeoff situation. We emphasize once more that one should not extend our conclusions to other situations. Rather, one should see this work as providing the tools that allow for a fair comparison of tradeoff algorithms, and use these to arrive at their own final judgements.

One conclusion we can provide about the relative performance of different tradeoff algorithms is that any difference in performance will be rather small. The practical inconvenience of having to align each entry of the pre-computed tradeoff table at a byte boundary has not been considered in this work, and the performance differences between algorithms can be so small that such obscure issues may be of more importance in practice. In the extreme case, issues as small as which algorithm is easier to implement may affect the choice of algorithms.

The fact that algorithm performances are not very different is disappointing to us as authors of this work, but this fact should be relieving to practitioners of the tradeoff algorithm that are not concerned with small performance differences. Still, even if one decides to ignore small performance differences, graphs of the previous section show that meaningful reduction in pre-computation cost can be achieved with only a small sacrifice to tradeoff

performance, and being able to take advantage of this knowledge will be of practical importance.

This work did not consider the use of checkpoints [1], which can be used to reduce the cost of false alarms. This decision was mostly based on the work [9], where the effect of checkpoints in reducing the online time of non-perfect Hellman and rainbow tradeoffs was shown to be quite smaller than 10% at typical parameters. Since its introduction will add much complication to the analysis, while having only a small effect on the final difference of tradeoff performances, we chose not to consider its use. Still, the effect of checkpoints on the DP tradeoff has not yet been accurately analyzed, and there is a small possibility that its behavior on the DP tradeoff will be different from that on the other algorithms.

Analysis of perfect table versions of the tradeoff algorithms analogous to what is given here also remains to be done, with some partial results available from [1, 9, 13]. Due to the larger pre-computation cost, the perfect table cases may be of less practical interest, but they are certainly interesting theoretically.

References

1. G. Avoine, P. Junod, P. Oechslin, Characterization and improvement of time-memory trade-off based on perfect tables. *ACM Trans. Inform. Syst. Secur.*, **11**(4), 17:1–17:22 (2008). Preliminary version in INDOCRYPT 2005
2. E. Barkan, E. Biham, A. Shamir, Rigorous bounds on cryptanalytic time/memory tradeoffs, in *Advances in Cryptology—CRYPTO 2006*, LNCS, vol. 4117, (Springer, 2006), pp. 1–21
3. A. Biryukov, A. Shamir, Cryptanalytic time/memory/data tradeoffs for stream ciphers, in *Advances in Cryptology—ASIACRYPT 2000*, LNCS, vol. 1976, (Springer, 2000), pp. 1–13
4. D. E. Denning, *Cryptography and Data Security* (Addison-Wesley, 1982)
5. P. Flajolet, A. M. Odlyzko, Random mapping statistics, in *Advances in Cryptology—EUROCRYPT '89*, LNCS, vol. 434, (Springer, 1990), pp. 329–354
6. S. Goldwasser, M. Bellare, Lecture Notes on Cryptography. Unpublished manuscript, July 2008. Available at: <http://cseweb.ucsd.edu/~mihir/papers/gb.html>
7. J. Dj. Golić, Cryptanalysis of alleged A5 stream cipher, in *Advances in Cryptology—EUROCRYPT '97*, LNCS, vol. 1233, (Springer, 1997), pp. 239–255
8. M. E. Hellman, A cryptanalytic time-memory trade-off. *IEEE Trans. on Infor. Theory*, **26**, 401–406 (1980)
9. J. Hong, The cost of false alarms in Hellman and rainbow tradeoffs. *Des. Codes Cryptogr.*, to appear
10. D. Ma, J. Hong, Success probability of the Hellman trade-off. *Inf. Process. Lett.*, **109**(7), 347–351 (2009)
11. A. J. Menezes, P. C. van Oorschot, S. A. Vanstone, *Handbook of Applied Cryptography* (CRC Press, 1997)
12. P. Oechslin, Making a faster cryptanalytic time-memory trade-off. in *Advances in Cryptology—CRYPTO 2003*, LNCS, vol. 2729, (Springer, 2003) pp. 617–630
13. F.-X. Standaert, G. Rouvroy, J.-J. Quisquater, J.-D. Legat, A time-memory tradeoff using distinguished points: New analysis & FPGA results, in *Cryptographic Hardware and Embedded Systems—CHES 2002*, LNCS, vol. 2523, (Springer, 2003), pp. 593–609