Client-CASH: Protecting Master Passwords against Offline Attacks

Jeremiah Blocki^{*}
Microsoft Research
jblocki@microsoft.com

Anirudh Sridhar[†]
Carnegie Mellon University
asridha1@andrew.cmu.edu

ABSTRACT

Offline attacks on passwords are increasingly commonplace and dangerous. An offline adversary is limited only by the amount of computational resources he or she is willing to invest to crack a user's password. The danger is compounded by the existence of authentication servers who fail to adopt proper password storage practices like key-stretching. Password managers can help mitigate these risks by adopting key stretching procedures like hash iteration or memory hard functions to derive site specific passwords from the user's master password on the client-side. While key stretching can reduce the offline adversary's success rate, these procedures also increase computational costs for a legitimate user. Motivated by the observation that most of the password guesses of the offline adversary will be incorrect, we propose a client side cost asymmetric secure hashing scheme (Client-CASH). Client-CASH randomizes the runtime of client-side key stretching procedure in a way that the expected computational cost of our key derivation function is greater when run with an incorrect master password. We make several contributions. First, we show how to introduce randomness into a client-side key stretching algorithms through the use of halting predicates which are selected randomly at the time of account creation. Second, we formalize the problem

of finding the optimal running time distribution subject to certain cost constraints for the client and certain security constraints on the halting predicates. Finally, we demonstrate that Client-CASH can reduce the adversary's success rate by up to 21%. These results demonstrate the promise of the Client-CASH mechanism.

Keywords

Master Passwords; Offline Attacks; Client-side Key-Stretching; Cost Asymmetric Secure Hash

1. INTRODUCTION

Passwords are often the primary entryway to access a user's confidential information on a website, and are thus a focus of attention for attackers. Offline attacks against passwords are particularly powerful. An offline adversary has access to the cryptographic hash of a user's password and can check a vast number of password possibilities without interacting with the authentication server. This adversary is only restricted by the computational resources that he is willing to invest into breaching each account. Offline attacks are increasingly commonplace due to recent high-profile security breaches at organizations like LinkedIn, Sony, eBay, and Ashley Madison¹.

Several factors contribute to the danger of offline attacks. First, users struggle to remember high entropy passwords for multiple accounts. Second, many organizations had failed to implement proper password storage techniques at the time they were breached. Finally, password cracking resources (e.g., hardware and password dictionaries) continue to improve allowing an adversary to mount cheaper and more effective attacks.

Client-side password management tools (e.g., PwdHash [35]) allow the user to generate multiple passwords from one master password and apply secure procedures like key stretching or salting on the client-side. While password managers offer many benefits they also introduce a single point of failure (i.e., the master password) through which the adversary could attack to gain access to all of a user's accounts. A user's master password is not necessarily immune to offline attacks just because a password manager is used. The recent breach of LastPass² highlights this dangerous possibility. Thus, key stretching procedures (e.g., hash iteration,

^{*}This work was completed in part while the Jeremiah was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant #CNS-1523467.

[†]Anirudh Sridhar was supported by grant #PS14011 from the PNC Center for Financial Services Innovation. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the author(s) and do not necessarily reflect the views of Microsoft, the PNC Center for Financial Services Innovation at Carnegie Mellon University or the National Science Foundation.

The full version [7] of this paper is available at <http://arxiv.org/abs/1603.00913>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASIA CCS '16, May 30-June 03, 2016, Xi'an, China

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4233-9/16/05...\$15.00

DOI: <http://dx.doi.org/10.1145/2897845.2897876>

¹For example, see <http://www.privacyrights.org/data-breach/> (Retrieved 9/1/2015).

²For example, see <https://blog.lastpass.com/2015/06/lastpass-security-notice.html> (Retrieved 9/1/2015).

memory hard functions) are recommended to mitigate security risks in the event of an offline attack. However, key-stretching increases costs for the honest party as well as the adversary. It is thus desirable to make authentication costs asymmetric so that (on average) a user authenticating with a correct password will incur lower costs than an offline adversary attempting to check an incorrect password guess. Previous work by Manber [27] and by Blocki and Datta [4] achieves this goal of cost-asymmetric key-stretching, but these solutions are only appropriate for server-side key stretching (see discussion in Section 2).

Contributions.

In this work we present Client-CASH a client-side key-stretching algorithm which achieves the goal of asymmetric costs, and we demonstrate that this system can protect user passwords from offline attacks. The core idea behind Client-CASH is to randomize the runtime of the client-side key stretching so that on average, the cost of verifying a correct password is smaller than the cost of rejecting an incorrect password. This is achieved by the use of halting predicates $P : \{0, 1\}^* \rightarrow \{0, 1\}$, which tells us when to stop hashing the password (e.g., stop after t rounds of hash iteration if $P(\mathbf{H}^t(pwd)) = 1$). These halting predicates are chosen randomly at the time of account creation and are stored by the client (e.g., on a local computer or on the cloud).

Because these halting predicates are stored by the client we must take care to ensure that the predicates themselves do not leak too much information about the user’s password. The key challenge is to select the predicates in a way that satisfies two seemingly conflicting requirements: 1) the halting predicates should induce a cost-asymmetry in the validation of correct/incorrect passwords, and 2) the adversary should not be able learn much about any user’s master password even if he observes these halting predicates. We borrow ideas from differential privacy [19, 29] to satisfy both of these conflicting requirements.

We formalize the properties of Client-CASH and give an example of a system that satisfies these properties based on the Exponential Mechanism [29], a powerful tool in differentially private analysis. We additionally formalize the problem of minimizing the percentage of passwords cracked by the adversary as an optimization over the parameters of the system and show that it can be solved efficiently as a linear program. In this work we analyze the security of Client-CASH systems that use up to two iterations of an underlying hash function $\mathbf{H}_k$ for password key stretching (the two-round case) or up to three iterations of $\mathbf{H}_k$ for key stretching (the three-round case). Here, the parameter k specifies the cost of the underlying hash function (e.g., BCRYPT [33] uses hash iteration to control costs and SCRYPT [31] uses a memory hard function to control costs). If we let $\Pr[E_{i,pwd}]$ denote the probability that Client-CASH terminates after the i ’th round given the correct password pwd as input then the expected cost to verify a correct password guess is $k \sum_i i \cdot \Pr[E_{i,pwd}]$.

We compare Client-CASH to deterministic key-stretching techniques $\mathbf{H}_{k'}$ with equivalent costs $k' = k \sum_i i \cdot \Pr[E_{i,pwd}]$ to Client-CASH by looking at the probability that an offline adversary with a finite computational budget B could crack the user’s password — we use $P_{adv,B}$ (resp. $P_{det,B}$) to denote the adversary’s success rate against Client-CASH (resp. deterministic key-stretching). For the two-round case, we

obtain a reduction $P_{det,B} - P_{adv,B}$ in passwords breached of up to 12%, and up to 21% in the three round case. Although we only show results for 2 and 3 round systems, we formulate the essential groundwork for analysis of systems with more rounds. From the significant decreases in passwords breached from two round to three round systems, we might expect that introducing more rounds will decrease the number of passwords breached in an offline attack even further.

In our analysis we make no assumptions on the password storage practices of the authentication server. If the authentication server adopts techniques like key stretching and salting then this will only make the adversary’s task harder. However, we do assume that we face an optimal offline adversary with a finite computational budget B , and that users choose passwords uniformly at random. Because the later assumption does not hold for general users we only recommend our solution to users who choose passwords (nearly) uniformly at random. However, we argue that this assumption is plausible for many of the security conscious users that would opt to use a client-side key stretching algorithm (finding Client-CASH distributions that are optimal for protecting non-uniform passwords is an important direction for future work).

Overview of Client-CASH.

When Alice creates an account on server W using the master password pwd_A our client-side application first selects a sequence of halting predicates $ow = (P_{u_1}, \dots, P_{u_{n-1}})$ using a randomized function $\mathbf{SelPreds}(pwd_A)$. Afterwards our application sends the message $(Alice, H)$ to the server where $H = \mathbf{H}_k^{S_W(pwd_A)}(pwd_A)$. Here, $\mathbf{H}_k$ denotes a collision resistant hash function which costs k work units to compute one time, $\mathbf{H}_k^i$ denotes the hash function iterated i times and $S_W(pwd_A)$ denotes the stopping time for password pwd_A given by the halting predicates ow — that is the smallest number $i \geq 1$ for which $P_{u_i}(\mathbf{H}_k^i(pwd_A)) = 1$. During authentication, Alice recomputes the derived password $H = \mathbf{H}_k^{S_W(pwd_A)}(pwd_A)$, sends $(Alice, H)$ to the server and gains access to her account. An incorrect password guess $p_g \neq p_A$ would be rejected with high probability since $\mathbf{H}_k$ is collision resistant.

2. RELATED WORK

Halting Puzzles. At a high level our use of halting predicates is similar to Boyen’s [12] halting puzzles. In Boyen’s solution the chosen halting predicate will never return ‘halt’ unless we run the key derivation function with the correct password so the key derivation algorithm never halts (or only halts after the maximum possible number of rounds). The key difference between our work and the work of Boyen [12] is that we carefully bound the amount of information that the chosen halting predicate(s) can leak about the user’s password. Thus, unlike [12], we can ensure that an adversary who only breaches the client will not be able to execute an offline attack against the user’s password.

Cost-Asymmetric Server-Side Key Stretching. Manber [27] proposed the use of hidden salt values (e.g., ‘pepper’) to make it more expensive to reject incorrect passwords. Blocki and Datta [4] refined this idea using game theoretic tools. While our work closely follows the work of Blocki and Datta [4], we stress that neither work [4, 27] addresses the issue of client side key-stretching. In both of these schemes the authentication server selects a secret salt

value $t \in \{1, \dots, m\}$ (e.g., “pepper”) and stores the cryptographic hash $\mathbf{H}(pwd, t)$ — the value t is not stored on the authentication server. An adversary would need to compute the hash function m times in total to reject an incorrect password, while the authentication server will need to compute it *at most* $(m + 1)/2$ times on average to verify a correct password guess because it can halt immediately after it finds the correct value of t . This approach is not suitable for client-side key-stretching because we would produce m different derived keys, but the client program would not know which one is correct — neither of these value t or $\mathbf{H}(pwd, t)$ should be stored on the client. Since we are performing key-stretching on the client-side we need to ensure that the final derived password that is sent to the authentication server is consistent among different authentication sessions.

Password Management Software. Password managers like PwdHash [35] allow the user to generate multiple passwords from one master password. PwdHash uses a public key-derivation function to (re)generate each of the user’s passwords from a single master password. Since the key-derivation function is public an adversary who breaks into any of the third party authentication servers could still execute an offline attack against the user’s master password. By contrast, password managers like KeePass [34] store an encrypted password vault on the client and are not necessarily vulnerable in the previous scenario because the adversary would not have the password vault. The vault, which contains all of the user’s passwords, is encrypted with the user’s master password. However, any adversary who breaks into the client and steals a copy of this vault could execute an offline attack against the user’s master password. Commercial applications like LastPass rely on a trusted server to derive passwords from the user’s master password. Unlike PwdHash and KeePass these commercial applications are typically not open source so it is not always possible for independent researchers to verify their security properties. In theory these password managers could be designed so that an adversary would need to break into multiple servers (e.g., at LinkedIn and LastPass) before he can mount an offline attack on the user’s master password. However, the recent breach at LastPass³ demonstrates that we cannot rule out this dangerous possibility. Similarly, Client-CASH is designed so that the adversary would need to breach both the client computer and a third party authentication server to mount an offline attack.

Deterministic Key Stretching Techniques. Advances in computing hardware (e.g., GPUs [25], ASICs [17]) make offline attacks increasingly dangerous. An offline adversary can often try millions of password guesses per second. Morris and Thompson [30] proposed the idea of key-stretching to make the hash function more expensive to evaluate so that an offline attack is more expensive for the adversary. Other defenses (e.g., distributing the storage/computation of cryptographic hash values so that an adversary who only breaches one server does not learn anything about the user’s password [14] and [16]) require multiple dedicated authentication servers. Finding good key-stretching techniques is an active area of research⁴. Hash iteration (e.g., PBKDF2 [24],

BCRYPT [33]) alone is often viewed as an insufficient key-stretching technique because an offline adversary can often significantly reduce costs by building customized hardware to evaluate the iterated hash function. While computational speeds may vary greatly between different devices, memory latency speeds are relatively consistent [18]. Thus, modern password hash functions like SCRYPT [31] or Argon2 [2] typically use memory hard functions [18] for key-stretching purposes. Our work is largely orthogonal to these lines of research. In particular, we stress that Client-CASH is compatible with both forms of key-stretching (hash iteration and memory hard functions).

Password Alternatives. Although researchers have been working on alternatives to text-passwords (e.g., graphical passwords [23, 15, 1] or biometrics [9]) text-passwords are likely to remain entrenched as the dominant form of authentication for many years [9]. While we focus on text passwords in this paper we stress that the applications of Client-CASH are not necessarily limited to text passwords. Client-side key stretching is a valuable primitive that could be used to protect any lower entropy secret whether that secret is a text password, a graphical password or a biometric signal.

3. DESCRIPTION OF THE MECHANISM

In this section we introduce the Client-CASH mechanism and describe the account creation and authentication protocols. In our presentation we will use $\mathbf{H}$ to denote a cryptographic hash function (e.g., SHA256 or Blake2b) and we will use C_H to denote the cost of computing $\mathbf{H}$ one time. We will also use $\mathbf{H}_k$ to denote a deterministic hash function that is k times more expensive to compute than $\mathbf{H}$ (i.e., $C_{H_k} = k \cdot C_H$). This might be achieved by hash iteration [33, 24] or by the adoption of memory hard functions [31]. We also use $\mathcal{P}$ to denote the set of passwords a user can pick.

Account Creation. When a user u creates an account a with a master password $pwd_u \in \mathcal{P}$ Client-CASH will execute the following steps: First, the client will run a randomized algorithm **SelPreds**(pwd_u) to obtain a sequence of $n - 1$ halting predicates $o_u = (P_{u_1}, \dots, P_{u_{n-1}})$. Here, a halting predicate $P_u : \{0, 1\}^* \rightarrow \{0, 1\}$ is simply a function that will tell us when to halt the key-derivation process. Second, the client will then store the tuple (a, u, s_u, o_u) , where $s_u \leftarrow \mathbf{Unif}(\{0, 1\}^L)$ is a random L -bit salt value. The client will then run the algorithm **Reproduce** (described below) to derive the password for account a .

We intentionally omit the workings of **SelPreds** and treat it as a black box for now. However, we stress that outcome o_u selected by the randomized algorithm **SelPreds**(pwd_u) may depend both on the master password pwd_u and on a security parameter ϵ which bounds the amount of information the outcome o_u might leak about pwd_u . In later sections we will show how to construct a randomized algorithm **SelPreds** which minimizes the adversary’s success rate $P_{adv,B}$ subject to certain security and cost constraints. The account creation protocol is described formally in Algorithm 1.

Authentication. When the user u attempts to access the account a with the password guess $pwd_g \in \mathcal{P}$, the client first locates the record (a, u, s_u, o_u) on the client. Then we execute the algorithm **Reproduce** function to derive the schemes (e.g., [2, 22])

³See <https://blog.lastpass.com/2015/06/lastpass-security-notice.html/> (Retrieved 9/1/2015).

⁴For example, the Password Hashing Competition (<https://password-hashing.net/index.html>) was developed to encourage the development of alternative password hashing

Algorithm 1 Create Account (Client Side)

Input: account name a , username u , password pwd_u ; random bit strings r_1, r_2
System Parameters: rounds n , $\mathcal{O}$, iterations k , ϵ , L
 $o_u \leftarrow \text{SelPreds}(pwd_u; r_1)$
 $key \leftarrow (u, a); s_u \leftarrow \text{Unif}(\{0, 1\}^L; r_2); value \leftarrow (o_u, s_u)$
StoreOnClient($key, value$)
 $H \leftarrow \text{Reproduce}(u, pwd_u, s_u, n)$
SendToServer(u, H)

password $\mathbf{H}_k^{S(pwd_g, o_u)}(pwd_g, s_u)$ for account a .

Algorithm 2 Reproduce (Client Side)

Input: account name a , username u , password pwd_g
 $key \leftarrow (u, a); value \leftarrow \text{FindClientRecord}(key)$
If $value = \emptyset$ **then return** “Account does not exist.”
 $(o_u, s_u, n, \mathcal{O}, \epsilon) \leftarrow value$
 $(P_1, \dots, P_{n-1}) \leftarrow o_u; H \leftarrow \mathbf{H}_k(pwd_g, s_u)$
for $m = 1, \dots, n - 1$ **do**
 If $P_m(H) = 1$ **then Break**
 $H \leftarrow \mathbf{H}_k(H)$
end for
SendToServer(u, H)

Here, we use $\mathcal{S}(pwd, o)$ to denote the implicitly defined stopping time for each password $pwd \in \mathcal{P}$ given the sequence $o = (P_{u_1}, \dots, P_{u_{n-1}})$ of halting predicates. Formally, $\mathcal{S}(pwd, o) = i$ if and only if (1) $P_{u_i}(\mathbf{H}_k^i(pwd, s_u)) = 1$ and $P_{u_j}(\mathbf{H}_k^j(pwd, s_u)) = 0$ for all $j < i$ or (2) $i = n$ and $P_{u_j}(\mathbf{H}_k^j(pwd, s_u)) = 0$ for all $j < n$, where n denotes the maximum possible rounds of hash iteration.

Thus, to compute $\mathbf{H}_k^{S(pwd_g, o_u)}(pwd_g, s_u)$ we initially compute $H_1 \leftarrow \mathbf{H}_k(pwd_g, s_u)$ and check if $P_{u_1}(H_1) = 1$. If it is then we return the derived password H_1 . Otherwise, we compute $H_2 \leftarrow \mathbf{H}_k(H_1)$ and return H_2 if and only if $P_{u_2}(H_1) = 1$. This process is repeated until either $P_{u_i}(H_i) = 1$ or $i = n$. The derived password is sent to the server to be accepted or rejected. Authentication is guaranteed when $pwd_g = pwd_u$ and is very unlikely when $pwd_g \neq pwd_u$ because $\mathbf{H}$ is collision resistant. The client-side algorithm **Reproduce** is presented formally as Algorithm 2. We note that, unlike the account creation process, the algorithm **Reproduce** is entirely deterministic.

Remark. We omit any description of how the authentication server stores the derived password as this is an orthogonal issue. In an ideal world the authentication server would add salt and apply a strong key-stretching algorithm before storing the derived password. Unfortunately, many authentication servers have failed to adopt these standard security practices [10]. Furthermore, users will not necessarily know what security practices have been adopted until the authentication server is actually breached. Thus, in our security analysis we will assume that the authentication server does not do any key-stretching. By applying salting and key-stretching algorithms on the client-side we can help protect users even when organizations fail to adopt these security practices. Of course if the authentication server does perform additional key stretching then the adversary’s task will be even harder.

Notation and Customizable Parameters We use B to

denote the budget of the adversary (i.e., the maximum number of times the adversary is willing to evaluate $\mathbf{H}$ in an attempt to crack the user’s password), and we will use C_{srv} to denote the maximum cost that the client is willing to bear per authentication session (in expectation).

We use $\mathcal{O}$ to denote the range of the **SelPreds** function (e.g., the space of all valid predicate sequences of length $n - 1$, where n denotes the maximum number of rounds of hashing for **Reproduce**). We adopt the following notational conventions: Given a randomized algorithm like **SelPreds** we use $o \leftarrow \text{SelPreds}(pwd)$ to denote a random sample from the distribution induced by an input pwd . If we fix the random input bits r then we will use $o := \text{SelPreds}(x; r)$ to denote the deterministic result. We will use $\mathcal{P}$ to denote the space of all possible passwords that the user might select.

For a user with password pwd_u and predicate sequence o_u , the hash used by that user on the instance of a correct password is $\mathbf{H}_k^{S(pwd_u, o_u)}$. We will use $\mathcal{O}_{j, pwd} \doteq \{o \in \mathcal{O} \mid \mathcal{S}(pwd, o) = j\} \subseteq \mathcal{O}$ to denote the subset of outcomes which yield stopping time j for the password pwd . Finally, we will use the parameter $\epsilon > 0$ to quantify the maximum amount of information leaked about a user’s password by the output $o \in \mathcal{O}$ of the **SelPreds** function. For readers familiar with the notion of ϵ -differential privacy we remark that we use the same notation intentionally.

4. DESIGN OF THE MECHANISM

In the previous section we outlined the Client-CASH mechanism using the randomized algorithm **SelPreds** in a black-box manner. We now examine the exact formulation of this algorithm.

4.1 Security Requirements

The probability that **SelPreds**(pwd) yields a particular outcome $o \in \mathcal{O}$ may depend on the input password pwd . Indeed, our goal is to introduce a cost asymmetry so that, in expectation, the stopping time $\mathcal{S}(pwd, o)$ for the correct password is less than stopping time $\mathcal{S}(pwd', o)$ for any incorrect passwords $pwd' \neq pwd$. One natural way to achieve this asymmetry might be to define a family of predicates

$$P_{pwd, j}(x) = \begin{cases} 1 & \text{if } x \equiv H_k^j(pwd, s_u) \\ 0 & \text{else} \end{cases}$$

for each password $pwd \in \mathcal{P}$ and $j < n$. Then **SelPreds**(pwd) might select a stopping time j at random from some distribution and set $o = P_{u_1}, \dots, P_{u_{n-1}}$ where $P_{u_i} = P_{pwd, j}$ for each $i < n$. This solution could provide an extreme cost asymmetry. In particular, we would have $\mathcal{S}(pwd, o) = j$ for the correct password and with high probability we would have $\mathcal{S}(pwd', o) = n$ for all other passwords. Indeed, this solution is simply a reformulation of Boyen’s [12] halting puzzles.

The downside to this solution is that it would enable an adversary to mount an offline attack using *only* state from the client (e.g., without breaching the authentication server) because the output o of **SelPreds** is stored on the client and this output would include the hash value $H_k^j(pwd, s_u)$.

In contrast to Boyen [12], we will require the **SelPreds** function to satisfy a very stringent information theoretic security property. In particular, for every $o \in \mathcal{O}$ and for all

$pwd, pwd' \in \mathcal{P}$ with $pwd \neq pwd'$, we will require that

$$\frac{\Pr[\text{SelPreds}(pwd) = o]}{\Pr[\text{SelPreds}(pwd') = o]} \leq e^\epsilon. \quad (1)$$

Here, ϵ is the security parameter which upper bounds the amount of information about the user's password that is leaked by the outcome o . This requirement is exactly equivalent to the powerful notion of ϵ -differential privacy [19]. Intuitively, the probability that we produce a particular outcome $o \in \mathcal{O}$ cannot depend too much on particular password pwd that the user chose.

Discussion. Differential privacy is an information theoretic guarantee, which holds even if the adversary has background knowledge about the user. The security guarantees are quite strong. In particular, let Attack_z denote *any* attack that an adversary with background knowledge z might use after observing the outcome o from $\text{SelPreds}(pwd)$ and let $\text{Bad} \subseteq \text{Range}(\text{Attack}_z)$ denote the set of outcomes that our user Alice would consider harmful. It is easy (e.g., see [20]) to show that

$$\frac{\Pr[\text{Attack}_z(\text{SelPreds}(pwd)) \in \text{Bad}]}{\Pr[\text{Attack}_z(\text{SelPreds}(pwd')) \in \text{Bad}]} \leq e^\epsilon.$$

Intuitively, the numerator represents the adversary's actual success rate and the denominator represents the adversary's success rate if the output of SelPreds were not even correlated with Alice's real password. Thus, the constraint implies that the probability the adversary's attack succeeds cannot greatly depend on the output of the SelPreds function.

4.2 Cost Requirements

Performing client-side key stretching requires computational resources on the client side. We will require that the amortized cost of hashing a password does not exceed the maximum amortized cost the user is willing to bear. Let C_{srv} denote the user's maximum amortized cost and let $k \cdot C_H$ denote the cost of executing $\mathbf{H}_k$ one time. Finally, for a fixed password pwd let $\Pr[E_{j,pwd}] \doteq \Pr[\text{SelPreds}(pwd) \in \mathcal{O}_{j,pwd}]$ denote the probability that SelPreds returns an outcome $o \in \mathcal{O}$ which yields stopping time $\mathcal{S}(pwd, o) = j$ for the password pwd . Given a password pwd the expected cost of executing **Reproduce** in Client-CASH is

$$\mathbf{E}[\mathcal{S}(pwd, o) \cdot k \cdot C_H] = k \cdot C_H \sum_{i=1}^n i \cdot \Pr[E_{i,pwd}],$$

where the expectation is taken over the randomness of the algorithm SelPreds . Thus, we require that

$$\sum_{i=1}^n i \cdot \Pr[E_{i,pwd}] \leq \frac{C_{srv}}{k \cdot C_H}. \quad (2)$$

4.3 Symmetric Predicate Sets

In this section we show how to simplify the security and cost constraints from the previous sections by designing $\mathcal{O}$ with certain symmetric properties. In particular, we need to design the SelPreds algorithm in such a way that the cost constraint (eq. 2) is satisfied for *all passwords* $pwd \in \mathcal{P}$, and we need to ensure that our security constraint (eq. 1) holds for *all pairs* of passwords $pwd \neq pwd'$. Working with $|\mathcal{P}|^2$ different security constraints is unwieldy. Thus, we need a way to simplify these requirements.

One natural way to simplify these requirements is to construct a *symmetric* $\mathcal{O}$; that is, for all $pwd, pwd' \in \mathcal{P}$ and for all $j \in \{1, \dots, n\}$, $|\mathcal{O}_{j,pwd}| = |\mathcal{O}_{j,pwd'}|$. We stress that this restriction does not imply that the sets $\mathcal{O}_{j,pwd}$ and $\mathcal{O}_{j,pwd'}$ are the same, simply that their sizes are the same. We will also require that $\forall pwd, pwd' \in \mathcal{P}, \forall j \leq n, \forall o \in \mathcal{O}_{j,pwd}, o' \in \mathcal{O}_{j,pwd'}$

$$\Pr[\text{SelPreds}(pwd) = o] = \Pr[\text{SelPreds}(pwd') = o'] \quad (3)$$

as there is no clear reason to favor one outcome over the other since both outcomes $o, o' \in \mathcal{O}$ yield the same halting time for their respective passwords pwd and pwd' (i.e., $\mathcal{S}(pwd, o) = \mathcal{S}(pwd', o')$).

Notation for Symmetric $\mathcal{O}$. Given a symmetric $\mathcal{O}$, we define $\mathcal{O}_j \doteq |\mathcal{O}_{j,pwd}|$. We define $\tilde{p}_j \doteq \Pr[\text{SelPreds}(pwd) = o]$ for an arbitrary password $pwd \in \mathcal{P}$ and output $o \in \mathcal{O}_{j,pwd}$ since, by symmetry, the choice of pwd and o do not matter. Intuitively, $\tilde{p}_j$ denotes the probability that $\text{SelPreds}(pwd)$ outputs o with stopping time $\mathcal{S}(pwd, o) = j$. We will also write $\Pr[E_j]$ instead of $\Pr[E_{j,pwd}] = \mathcal{O}_j \cdot \tilde{p}_j$ because the value will be the same for all passwords $pwd \in \mathcal{P}$.

Now the SelPreds algorithm can be completely specified by the values $\tilde{p}_j$ for $j \in \{1, \dots, n\}$. It remains to construct a symmetric $\mathcal{O}$. The construction of a symmetric $\mathcal{O}$ used in this work is simple we will focus on stopping predicates with the following form:

$$P_{i,j}(x) = \begin{cases} 1 & \text{if } x \equiv i \pmod{j} \\ 0 & \text{else} \end{cases}$$

Theorem 1 says that we can use these predicates to construct a symmetric set $\mathcal{O}$.

THEOREM 1. *Let integers $\ell_1, \dots, \ell_{n-1} \geq 2$ be given and suppose that $\mathcal{O} = \{P_{0,\ell_1}, \dots, P_{\ell_1-1,\ell_1}\} \times \dots \times \{P_{0,\ell_{n-1}}, \dots, P_{\ell_{n-1}-1,\ell_{n-1}}\}$. Then $\mathcal{O}$ is symmetric, i.e. for all $j \in \{1, \dots, n\}$ and for all $pwd, pwd' \in \mathcal{P}$ we have $|\mathcal{O}_{j,pwd}| = |\mathcal{O}_{j,pwd'}|$.*

One of the primary advantages of using a symmetric $\mathcal{O}$ is that we can greatly reduce the number of security and cost constraints. Table 1 compares the security and cost constraints with and without symmetry. Without symmetry we had to satisfy separate cost constraints for all passwords $pwd \in \mathcal{P}$. With symmetry we only have to satisfy one cost constraint. Similarly, with symmetry we only need to satisfy one security constraint for each $i, j \in \{1, \dots, n\}$ instead of multiple constraints for each pair of passwords $pwd \neq pwd'$ that the user might select. While the space of passwords may be very large, n , the number of rounds of hashing, will typically be quite small (e.g., in this paper $n \in \{2, 3\}$). See the full version of this paper for proofs of the statements in Table 1.

In our analysis of Client-CASH we will focus on two simple symmetric constructions of $\mathcal{O}$:

- “Two Round” Case: set $n = \ell_1 = 2$ in Theorem 1 so $\mathcal{O} = \{P_{0,2}, P_{1,2}\}$.
- “Three Round” Case: set $n = \ell_1 = \ell_2 = 3$ in Theorem 1 so $\mathcal{O} = \{P_{0,3}, P_{1,3}, P_{2,3}\}^2$.

We do not rule out the possibility that other cases would yield even stronger results. However, we are already able to obtain significant reductions in the adversary's success rate with these simple cases.

No Symmetry	Constraint	With Symmetry
$\forall pwd, pwd' \in \mathcal{P}, \forall o \in \mathcal{O},$ $\frac{\Pr[\mathbf{SelPreds}(pwd)=o]}{\Pr[\mathbf{SelPreds}(pwd')=o]} \leq e^\epsilon$	Security [7, Thm 4]	$\forall i, j \in \{1, \dots, n\},$ $\frac{\tilde{p}_i}{\tilde{p}_j} \leq e^\epsilon$
$\forall pwd \in \mathcal{P},$ $\sum_{j=1}^n j \cdot \Pr[E_{j,pwd}] \leq \frac{C_{srv}}{k \cdot C_H}$	Cost [7, Thm 5]	$\sum_{j=1}^n j \cdot O_j \tilde{p}_j \leq \frac{C_{srv}}{k \cdot C_H}$

Table 1: Security and cost constraints for **SelPreds** with symmetric predicate set $\mathcal{O}$.

5. THE ADVERSARY MODEL

In this section we formalize the strategies that an offline adversary might adopt so that we can analyze the effectiveness of the Client-CASH mechanism against an offline adversary. In an *online* attack, the adversary must interact with and query the authentication server to have a password verified or rejected, and there are a variety of effective methods to deal with such an attack (e.g., a k -strikes policy). By contrast, the offline adversary we consider is limited only by the resources that he is willing to invest to crack the user's password. The adversary's goal is to maximize the odds that he cracks the user's password given a finite computational budget B . We will use $P_{adv,B}$ (resp. $P_{det,B}$) to denote the probability that the *optimal* offline adversary successfully cracks a password protected with Client-CASH (resp. deterministic key-stretching) with a budget B denoting the maximum number of times that the adversary is willing to compute $\mathbf{H}$. To simplify notation we will often write P_{adv} (resp. P_{det}) when B is clear from context.

The Adversary's Knowledge. We consider an adversary who has breached the client and one of the third party authentication servers (e.g., LinkedIn) where the user has an account. This adversary has access to all code and data on the client and on a third party authentication server. Suppose for example that user u created account a with a password pwd_u . The client will store the tuple $(u, a, o_u, s_u, n, \mathcal{O}, \epsilon)$ and the third party authentication server will have some record of the derived password $\mathbf{H}_k^{S(pwd_u, o_u)}(pwd_u, s_u)$. In our analysis we do not assume that the third party authentication server adopts proper password storage procedures like key-stretching — this is the unfortunate reality for many authentication servers [10]. If the authentication server does apply key-stretching then the adversary's offline attack will be even more expensive.

Offline attack against the traditional mechanism.

We first analyze the adversary's success rate $P_{det,B}$ against a traditional deterministic key stretching algorithm with comparable cost to Client-CASH. In particular, let $k' = \frac{C_{srv}}{C_H}$. Instead of adopting Client-CASH the adversary might instead use a deterministic key-stretching algorithm $\mathbf{H}_{k'}$ with the same cost. In this case the adversary could check at most $\frac{B}{k'}$ passwords, and each password has probability $\frac{1}{|\mathcal{P}|}$ of being the correct password. Thus,

$$P_{det,B} = \begin{cases} \frac{B}{k'|\mathcal{P}|} & \text{if } \frac{B}{k'} \leq |\mathcal{P}| \\ 1 & \text{otherwise} \end{cases} \quad (4)$$

Offline attack against Client-CASH.

We now consider the adversary's optimal success rate $P_{adv,B}$ against Client-CASH. While each password $pwd \in \mathcal{P}$ is chosen with equal probability, we might have different stopping times $\mathcal{S}(pwd, o_u)$ for each password. The adversary can compute $\mathbf{H}_k$ up to $\frac{B}{k}$ times so he can complete B/k rounds of hashing in total. For each password guess pwd_g the adversary can decide how many rounds of hashing to complete for that particular guess. There are many different strategies that the adversary might adopt, and some may produce a larger probability of success than others. Of course the optimal adversary will never complete more rounds of hashing than necessary (e.g., if $\mathcal{S}(pwd_g, o) = i$ then the adversary will never complete $i+1$ rounds of guessing for guess pwd_g).

It turns out that we can represent the optimal adversary's strategies as a vector of n numbers $\mathbf{b} = (b_1, \dots, b_n) \in \mathbb{R}^n$ where $b_i \in [0, \min\{1, k|\mathcal{P}|/B\}]$ denotes the fraction of the adversary's budget spent hashing passwords i 'th round (see the full version of this paper for a more formal justification). Formally, given a symmetric set $\mathcal{O}$ from Theorem 1 we use $F_B \subseteq \mathbb{R}^n$ to denote the set of feasible strategies for the adversary

$$F_B = \left\{ \mathbf{b} \in \mathbb{R}^n \mid \sum_{i=1}^n b_i = 1 \wedge \forall i \leq n \right. \\ \left. \left(b_i \in \left[0, \min \left\{ 1, \frac{k|\mathcal{P}|}{B} \right\} \right] \wedge b_{i+1} \leq b_i - \frac{b_i}{\ell_i} \right) \right\}.$$

We require that $b_{i+1} \leq b_i \cdot \frac{\ell_i - 1}{\ell_i}$ because the adversary can only hash a password pwd_g in round $i+1$ if he previously completed round i for the same password and the optimal adversary will not hash pwd_g in round $i+1$ if $\mathcal{S}(pwd_g, o) = i$ — for a random password the halting predicate will output 1 with probability $1/\ell_i$ in the i 'th round. Given a particular strategy $\mathbf{b}$ the adversary succeeds with probability

$$\frac{B}{k|\mathcal{P}|} \left(\Pr[E_1]b_1 + \sum_{i=2}^n \Pr[E_i]b_i \prod_{j=1}^{i-1} \left(\frac{\ell_j}{\ell_j - 1} \right) \right).$$

Intuitively, $b_i \frac{B}{k|\mathcal{P}|}$ denotes the fraction of password selected by the adversary to be hashed at least i time, and the $\prod_{j=1}^{i-1} \left(\frac{\ell_j - 1}{\ell_j} \right)$ term represents the fraction of passwords with

stopping time $\mathcal{S}(pwd_g, o) > i - 1$. Thus, $\Pr[E_i] \prod_{j=1}^{i-1} \left(\frac{\ell_j}{\ell_j - 1} \right)$ denotes the conditional probability that $\mathcal{S}(pwd_g, o) = i$ given that $\mathcal{S}(pwd_g, o) > i - 1$.

Because the user selects passwords uniformly at random from the set $\mathcal{P}$ we might conjecture that the optimal adversary will follow the same strategy for every password guess (e.g., for some $i \leq n$ the optimal adversary will hash each guess pwd_g for $\min\{i, \mathcal{S}(pwd_g, o)\}$ rounds before moving on to the next guess). This intuition turns out to be correct. Formally, $\mathcal{A}_i \doteq \{\mathbf{b} \in F_B : \forall \mathbf{b}' \in F_B, b_i \geq b'_i\}$ denotes the set of strategies in which the adversary tries each password guess for $\min\{i, \mathcal{S}(pwd_g, o)\}$ rounds before giving up. Our next results states that for some $i \leq n$ the adversary will follow a strategy in $\mathcal{A}_i$.

THEOREM 2. *Assume that $\mathcal{O}$ is constructed as in Theorem 1. The dominant adversary strategies are given by the collection of sets $\mathcal{A}_i = \{\mathbf{b} \in F_B : \forall \mathbf{b}' \in F_B, b_i \geq b'_i\}$ for*

$1 \leq i \leq n$. Let $F_B^* := \bigcup_{i=1}^n \mathcal{A}_i$. Then

$$P_{adv,B} = \frac{B}{k|\mathcal{P}|} \max_{\mathbf{b} \in F_B^*} \left\{ \Pr[E_1]b_1 + \sum_{i=2}^n \Pr[E_i]b_i \prod_{j=1}^{i-1} \left(\frac{\ell_j}{\ell_j - 1} \right) \right\}$$

Discussion. In our analysis of the adversary we assume that the user picks passwords uniformly at random from $\mathcal{P}$. Thus, we can only recommend Client-CASH to users for whom this assumption holds. A large body of research has explored the security of user selected passwords (e.g., [28, 13, 32, 8]), how users cope with multiple passwords [21] and how users respond to password restrictions [26, 6]. These results indicate that many users do not select their password uniformly at random — contrary to the assumption we made in our security analysis. However, several research results indicate that users are capable of remembering truly random system-assigned passwords with practice [3, 5, 11, 36]. Our assumption would hold for users who adopt the four random words strategy popularized by the web comic XKCD⁵ or the password management strategies proposed by Blocki et al. [3]. Thus, we believe that this assumption could be reasonable for many of the security conscious users who would adopt Client-CASH.

5.1 Two Round Strategies

In the two-round case we have $F^* = \mathcal{A}_1 \cup \mathcal{A}_2$ in Theorem 2, so it suffices to calculate these two sets. We have $\mathcal{A}_2 = \{(\frac{2}{3}, \frac{1}{3})\}$ and

$$\mathcal{A}_1 = \begin{cases} \{(1, 0)\} & \text{if } \frac{B}{k|\mathcal{P}|} \leq 1 \\ \left\{ \left(\frac{k|\mathcal{P}|}{B}, 1 - \frac{k|\mathcal{P}|}{B} \right) \right\} & \text{otherwise.} \end{cases}$$

If $\mathbf{b} = (1, 0)$ then, using the fact that $O_1 = 1$, we have

$$\frac{B}{k|\mathcal{P}|} \left(O_1 \tilde{p}_1 b_1 + \sum_{i=2}^n O_i \tilde{p}_i b_i \prod_{j=1}^{i-1} \left(\frac{\ell_j}{\ell_j - 1} \right) \right) = \tilde{p}_1 \frac{B}{k|\mathcal{P}|}.$$

Thus, working out each case for $\mathbf{b} \in F^*$ Theorem 2 simplifies to:

$$P_{adv,B} \geq \begin{cases} \max \left\{ \tilde{p}_1 \frac{B}{k|\mathcal{P}|}, \frac{2}{3} \frac{B}{k|\mathcal{P}|} \right\} & \text{if } \frac{B}{k|\mathcal{P}|} \leq 1 \\ \max \left\{ \tilde{p}_1 + 2 \left(\frac{B}{k|\mathcal{P}|} - 1 \right) \tilde{p}_2, \frac{2}{3} \frac{B}{k|\mathcal{P}|} \right\} & \text{otherwise.} \end{cases}$$

5.2 Three Round Strategies

In the three-round case we have $F^* = \mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{A}_3$ in Theorem 2, so it suffices to calculate these three sets. To bound $P_{adv,B}$ we note that it is sufficient to find the extremal points of each region since the adversary's objective function is linear in the b_i 's. Letting $\overline{\mathcal{A}}_i$ denote the extremal points of $\mathcal{A}_i$, we thus have

$$\overline{\mathcal{A}}_1 = \begin{cases} \{(1, 0, 0)\} & \text{if } \frac{B}{k|\mathcal{P}|} \leq 1 \\ \left\{ \left(\frac{k|\mathcal{P}|}{B}, 1 - \frac{k|\mathcal{P}|}{B}, 0 \right), \right. \\ \left. \left(\frac{k|\mathcal{P}|}{B}, \frac{3}{5} \left(1 - \frac{k|\mathcal{P}|}{B} \right), \frac{2}{5} \left(1 - \frac{k|\mathcal{P}|}{B} \right) \right) \right\} & \text{if } \frac{B}{k|\mathcal{P}|} \in [1, \frac{5}{3}] \\ \left\{ \left(\frac{k|\mathcal{P}|}{B}, \frac{3}{5} \left(1 - \frac{k|\mathcal{P}|}{B} \right), \frac{2}{5} \left(1 - \frac{k|\mathcal{P}|}{B} \right) \right) \right\} & \text{otherwise} \end{cases}$$

⁵See <http://xkcd.com/936/> (retrieved 9/8/2015).

$$\overline{\mathcal{A}}_2 = \begin{cases} \left\{ \left(\frac{3}{5}, \frac{2}{5}, 0 \right) \right\} & \text{if } \frac{B}{k|\mathcal{P}|} \leq \frac{5}{3} \\ \left\{ \left(\frac{k|\mathcal{P}|}{B}, \frac{2}{3} \frac{k|\mathcal{P}|}{B}, 1 - \frac{5}{3} \frac{k|\mathcal{P}|}{B} \right) \right\} & \text{otherwise} \end{cases}$$

$$\overline{\mathcal{A}}_3 = \left\{ \left(\frac{9}{19}, \frac{6}{19}, \frac{4}{19} \right) \right\}$$

Thus, we can express $P_{adv,B}$ in terms of the following bounds.

$$\begin{aligned} (1) & 3\tilde{p}_1 \frac{B}{k|\mathcal{P}|} & (2) & \left(\frac{9}{5}\tilde{p}_1 + \frac{6}{5}\tilde{p}_2 \right) \frac{B}{k|\mathcal{P}|} \\ (3) & \frac{9}{19} \frac{B}{k|\mathcal{P}|} & (4) & 3\tilde{p}_1 + 3\tilde{p}_2 \left(\frac{B}{k|\mathcal{P}|} - 1 \right) \\ (5) & 3\tilde{p}_1 + \left(\frac{9}{5}\tilde{p}_2 + \frac{18}{5}\tilde{p}_3 \right) \left(\frac{B}{k|\mathcal{P}|} - 1 \right) \\ (6) & 3\tilde{p}_1 + 2\tilde{p}_2 + 9\tilde{p}_3 \left(\frac{B}{k|\mathcal{P}|} - \frac{5}{3} \right) \end{aligned}$$

For example, if $\mathbf{b} = (\frac{3}{5}, \frac{2}{5}, 0)$ then, using the fact that $O_1 = 3, O_2 = 2$ and $\ell_j = 3$, we have

$$\begin{aligned} & \frac{B}{k|\mathcal{P}|} \left(O_1 \tilde{p}_1 b_1 + \sum_{i=2}^n O_i \tilde{p}_i b_i \prod_{j=1}^{i-1} \left(\frac{\ell_j}{\ell_j - 1} \right) \right) \\ &= \frac{B}{k|\mathcal{P}|} \left(3\tilde{p}_1 \frac{3}{5} + 2\tilde{p}_2 \frac{2}{5} \left(\frac{3}{2} \right) \right) = \frac{B}{k|\mathcal{P}|} \left(\tilde{p}_1 \frac{9}{5} + \frac{6}{5}\tilde{p}_2 \right) \end{aligned}$$

as in (2). Thus, in the three-round case Theorem 2 simplifies to:

$$P_{adv,B} \geq \begin{cases} \max\{(1), (2), (3)\} & \text{if } \frac{B}{k|\mathcal{P}|} \leq 1 \\ \max\{(2), (3), (4), (5)\} & \text{if } \frac{B}{k|\mathcal{P}|} \in [1, \frac{5}{3}] \\ \max\{(3), (5), (6)\} & \text{otherwise.} \end{cases}$$

6. THE EXPONENTIAL MECHANISM

In this section we show how to construct the **SelPreds** function so that it satisfies our security and cost constraints (equations 1 and 2). Our construction is based on the Exponential Mechanism of McSherry and Talwar [29], a powerful tool from the differential privacy literature. We demonstrate that this mechanism is feasible and that it leads to a significant reduction in the adversary's success rate. In particular the exponential mechanism can reduce the adversary's success rate by up to 12% when n , the maximum number of hashing rounds, is two and by up to 18% when $n = 3$.

6.1 Constructing SelPreds via the Exponential Mechanism

To define the **SelPreds** function it suffices to specify the probability of each outcome $o \in \mathcal{O}$ given an input password $pwd \in \mathcal{P}$. Consider the utility function $U(pwd, o) \doteq \frac{1 - S(pwd, o)}{n-1}$ which assigns a utility score to each outcome $o \in \mathcal{O}$ given a password pwd . We can now use $U(pwd, o)$ to specify the probability of an outcome o given password pwd . We define $\text{SelPreds}_{exp, \epsilon}$ such that

$$\Pr[\text{SelPreds}_{exp, \epsilon}(pwd) = o] \doteq \frac{e^{\epsilon U(pwd, o)}}{\sum_{o' \in \mathcal{O}} e^{\epsilon U(pwd, o')}}.$$

Intuitively, given a fixed password pwd the exponential mechanism assigns a higher probability to outcomes o with shorter stopping times $\mathcal{S}(pwd, o)$.

We can ensure that our cost constraints are satisfied by tuning the parameter k (the number of hash iterations per round). After ϵ and $\mathcal{O}$ have been fixed we will set k to be the maximum integer such that:

$$\sum_{i=1}^n iO_i \tilde{p}_i = \frac{1}{W} \sum_{i=1}^n iO_i e^{\epsilon U(i)} \leq \frac{C_{srv}}{k \cdot C_H}.$$

where the left hand side is a constant.

Observe that whenever $\mathcal{O}$ is symmetric it is possible to precompute the stopping time probabilities $\tilde{p}_j \doteq \frac{O_j \cdot e^{\epsilon \cdot j}}{\sum_{i=1}^n O_i \cdot e^{\epsilon \cdot i}}$. Thus, to sample from $\mathbf{SelPreds}_{exp, \epsilon}(pwd)$ it suffices to sample the stopping time j , compute $\mathbf{H}^{jk}(pwd, s_u)$ and sample uniformly at random from the set $\mathcal{O}_{j, pwd}$.

Theorem 3 states that the exponential mechanism above satisfies our desired security constraint. Theorem 3 is similar to the general result of McSherry and Talwar [29], which would imply that our mechanism $\mathbf{SelPreds}_{exp, \epsilon}$ satisfies the security constraint with security parameter 2ϵ . In our particular setting we can exploit symmetry (e.g., $|\mathcal{O}_{j, pwd}| = |\mathcal{O}_{j, pwd'}|$) to obtain a tighter bound with security parameter ϵ .

THEOREM 3. *For any symmetric $\mathcal{O}$ we have*

$$\forall pwd, pwd' \in \mathcal{P}, \forall o \in \mathcal{O}, \frac{\Pr[\mathbf{SelPreds}_{exp, \epsilon}(pwd) = o]}{\Pr[\mathbf{SelPreds}_{exp, \epsilon}(pwd') = o]} \leq e^\epsilon.$$

Proof of Theorem 3. Let $w_{o, pwd} \doteq e^{\epsilon U(pwd, o)}$ denote the weight of outcome $o \in \mathcal{O}$ given password pwd and let $W_{pwd} \doteq \sum_{o \in \mathcal{O}} w_{o, pwd}$ denote the cumulative weight of all $o \in \mathcal{O}$. Because $\mathcal{O}$ is symmetric we have $W_{pwd} = W_{pwd'}$ for any pair of passwords $pwd, pwd' \in \mathcal{P}$. Thus,

$$\begin{aligned} \frac{\Pr[\mathbf{SelectPredicate}(pwd) = o]}{\Pr[\mathbf{SelectPredicate}(pwd') = o]} &= \frac{w_{o, pwd} / W_{pwd}}{w_{o, pwd'} / W_{pwd'}} \\ &= \frac{w_{o, pwd}}{w_{o, pwd'}} \leq e^{|U(pwd', o) - U(pwd, o)|} \leq e^\epsilon. \end{aligned}$$

□

6.2 Analysis of the Exponential Mechanism

We will use $G_B \doteq P_{det, B} - P_{adv, B}$ to denote the gains from adopting Client-CASH (i.e., the reduction in the probability that the adversary cracks the password). Recall that $P_{adv, B}$ (resp. $P_{det, B}$) denotes the probability that the optimal offline adversary successfully cracks a password protected with Client-CASH (resp. deterministic key-stretching) with a budget B . Figure 1a and 1b plots G_B as a function of B for various values of the security parameter ϵ . Theorem 2 allows us to compute $P_{adv, B}$ and G_B efficiently.

6.2.1 Two-Round Exponential Mechanism

Note that when $n = 2$, $\mathcal{O} = \{P_{0,2}, P_{1,2}\}$ so $(O_1, O_2) = (1, 1)$. Thus for a fixed value of ϵ , we can define the exponential mechanism as follows:

$$W = 1 + e^{-\epsilon}, \quad \tilde{p}_1 = \frac{1}{W}, \quad \tilde{p}_2 = \frac{e^{-\epsilon}}{W}$$

$$\Pr[E_1] + \Pr[E_2] = \tilde{p}_1 + \tilde{p}_2 = 1$$

$$\Pr[E_1] + 2\Pr[E_2] = \tilde{p}_1 + 2\tilde{p}_2 \leq \frac{C_{srv}}{k \cdot C_H}$$

Although we have an inequality in the last equation, we emphasize that k is determined to be the maximum integer such that the inequality holds, as described in §6.1. Since k is large, and thus marginal changes in $\frac{C_{srv}}{k \cdot C_H}$ are small, we assume for simplicity that $\tilde{p}_1 + 2\tilde{p}_2 = \frac{C_{srv}}{k \cdot C_H}$ and round k downwards if it is not an integer.

6.2.2 Three-Round Exponential Mechanism

When $n = 3$, $\mathcal{O} = \{P_{0,3}, P_{1,3}, P_{2,3}\} \times \{P_{0,3}, P_{1,3}, P_{2,3}\}$ so $(O_1, O_2, O_3) = (3, 2, 4)$. Thus for a fixed ϵ , we can define the exponential mechanism as follows:

$$W = 3 + 2e^{-\epsilon/2} + 4e^{-\epsilon}$$

$$(\tilde{p}_1, \tilde{p}_2, \tilde{p}_3) = \left(\frac{1}{W}, \frac{e^{-\epsilon/2}}{W}, \frac{e^{-\epsilon}}{W} \right)$$

$$\Pr[E_1] + \Pr[E_2] + \Pr[E_3] = 3\tilde{p}_1 + 2\tilde{p}_2 + 4\tilde{p}_3 = 1$$

$$\Pr[E_1] + 2\Pr[E_2] + 3\Pr[E_3] = 3\tilde{p}_1 + 4\tilde{p}_2 + 12\tilde{p}_3 \leq \frac{C_{srv}}{k \cdot C_H}$$

Discussion.

Our analysis shows that exponential mechanism can be used to reduce the percentage of passwords breached by an optimal offline adversary. When $B = k'|\mathcal{P}|$ we will always have $P_{det, B} = 1$. Thus, we expect that the gain G_B will decrease monotonically after this point which is indeed what we observe in the plots. In the two round case we always had positive gain (e.g., for all $B \geq 0$ we had $G_B > 0$ for every value of ϵ that we tried). In the three round case we always had $G_B > 0$ whenever $\epsilon < 2.3$. Observe that G_B does not always increase monotonically as we increase ϵ (e.g., as we relax the security constraint). In both the two round and the three round case G_B increases monotonically with ϵ as long as $\epsilon \leq 0.5$, but after this the graphs intersect indicating that different values of ϵ are better than others for different B when the exponential mechanism is used.

We note that the exponential mechanism does not incorporate any information about the adversary's budget B . The exponential mechanism $\mathbf{SelPreds}_{exp, \epsilon \approx 2.3}$ with the largest value of ϵ performs poorly against smaller budgets B , but it provided the optimal defense in two rounds against an adversary with larger budgets. If we know B in advance then we can often improve our construction of the $\mathbf{SelPreds}$ function (see §7).

7. Client-CASH OPTIMIZATION

In this section we show how to optimize the $\mathbf{SelPreds}$ function under the assumption that the size of the password space $|\mathcal{P}|$ and the adversary's budget B are known in advance. This may not be an unreasonable assumption for many users. For example, Bonneau and Schechter estimated that for the SHA256 hash function $C_H = \$7 \times 10^{-15}$ [11], and Symantec reports that passwords generally sell for $\$4 - \30 on the black market. Thus, we might reasonably expect that a rational adversary will have a budget $B \leq \frac{30}{7 \times 10^{-15}} = 4.29 \times 10^{15}$. If the user selects passwords uniformly at random then it is reasonable to assume that $|\mathcal{P}|$ is known. For example, we might have $|\mathcal{P}| = 10^{12}$ for users who memorize three random person-action-object stories and use the associated action-object pairs for their password [3, 5].

Given a symmetric set $\mathcal{O}$ of outcomes our goal is to find probability values $\tilde{p}_1, \dots, \tilde{p}_n$ and a parameter k which min-

Figure 1: Percent reduction in passwords breached $G_B = P_{det,B} - P_{adv,B}$ given as a function of the adversary's total budget B with a normalizing factor of $\frac{1}{k|\mathcal{P}|} = \frac{C_H}{C_{srv}|\mathcal{P}|}$ for various values of ϵ . (a) and (b) show the two and three round exponential mechanisms, respectively, while (c) and (d) show the optimal gains for the two and three round mechanisms when the adversary's budget is known in advance.

imize the adversary's success rate $P_{adv,B}$ in the event of an offline attack. We must choose these values in a way that satisfy our security and cost constraints (Table 1). This goal is stated formally as Optimization Goal 1.

Constraints (1) and (2) specify that the $\tilde{p}_i$ values must define a valid probability distribution and constraints (3) and (6) are the security and cost constraints from Table 1. Constraint (4) states that k , the parameter controlling the cost of the underlying hash function $\mathbf{H}_k$ in each round of Client-CASH, can be at most $\frac{C_{srv}}{C_H}$. Otherwise it would be impossible to satisfy our cost constraint because just a single round of hash iteration would cost more than C_{srv} . Finally, constraint (5) implies that we must minimize $P_{adv,B}$ subject to the assumption that the adversary responds with his optimal strategy (see Theorem 2).

We note that, for a fixed value of k , Optimization Goal 1 becomes a linear program. Thus, we can solve Optimization Goal 1 efficiently by solving the resulting linear program for each value of $k \in \{1, \dots, \frac{C_{srv}}{C_H}\}$ and taking the best solution.

<p>Input Parameters: $C_{srv}, C_H, \epsilon, \mathcal{O}, B, \mathcal{P} , n$</p> <p>Variables: $\tilde{p}_1, \dots, \tilde{p}_n, k, P_{adv,B}$</p> <p>minimize $P_{adv,B}$ subject to</p> <p>(1) $\forall i \in \{1, \dots, n\}, \tilde{p}_i \in [0, 1]$ (2) $\sum_{i=1}^n O_i \cdot \tilde{p}_i = 1$</p> <p>(3) $\forall i, j \in \{1, \dots, n\}, \frac{\tilde{p}_i}{\tilde{p}_j} \leq e^\epsilon$ (4) $k \in \left\{1, \dots, \frac{C_{srv}}{C_H}\right\}$</p> <p>(5) $P_{adv,B} \geq \frac{B}{k \mathcal{P} } \max_{\mathbf{b} \in F^*} \left\{ O_1 \tilde{p}_1 b_1 + \sum_{i=2}^n O_i \tilde{p}_i b_i \prod_{j=1}^{i-1} \left(\frac{\ell_j}{\ell_j - 1} \right) \right\}$</p> <p>(6) $\sum_{j=1}^n j O_j \tilde{p}_j \leq \frac{C_{srv}}{k C_H}$</p>

Optimization Goal 1: Minimizing the adversary's probability of success.

Figure 2: Comparison of the percent reduction in cracked passwords G_B between the optimal distribution (blue) and the exponential distribution (orange) for as a function of B with a normalizing factor of $\frac{1}{k'|\mathcal{P}|} = \frac{C_H}{C_{srv}|\mathcal{P}|}$ for various values of ϵ . The results for the two round optimal mechanism is shown in (a), while the three round optimal mechanism is shown in (b).

7.1 Analysis

Figures 1c and 1d compare the optimal Client-CASH defense to the cost-equivalent deterministic key-stretching defense. We use the same symmetric predicate sets $\mathcal{O}$ as in previous sections (two-round and three-round). The full version of this paper [7] includes additional figures comparing the exponential mechanism and the optimal mechanism directly.

Discussion.

Our results demonstrate that it is often possible to obtain even greater reductions in the adversary’s success rate if we know the budget B in advance. For example, in the three round case $n = 3$ we can reduce the adversary’s rate of success by as much as $G_B = 21\%$. However, we note that the solution to Optimization Goal 1 may not be optimal if the estimate of the adversary’s budget B is wrong. Thus, if the user is uncertain about the adversary’s budget he may better off adopting the exponential mechanism.

Unlike the exponential mechanism we expect for the gain G_B to increase monotonically with larger ϵ because we only loosen constraint (3). This is what we observe in Figures 1c and 1d. We also note that when the adversary’s budget B exceeds $k'|\mathcal{P}|$ we have $P_{det,B} = 1$. Thus, G_B must decrease monotonically after this point because $P_{adv,B}$ will increase monotonically with B .

8. PRACTICAL GUIDELINES

All the Client-CASH mechanisms examined give reductions in the percentage of passwords breached by an optimal offline adversary. We also make the following recommendations of when a user should use a certain mechanism.

Two Rounds vs. Three Rounds.

Figures 3a and 3b compare the two-round Client-CASH mechanism to the three-round Client-CASH mechanism. Based on the results in Figure 3, we recommend using

three rounds of hashing in general. In particular, the three round mechanisms tend to achieve greater gains in security. The one possible exception is when the adversary’s budget is smaller, but we are uncertain about the adversary’s exact budget. In this one case the two round exponential mechanism might be the right choice⁶. Against adversaries with a larger budget the three round mechanism is the clear winner — note that the gains remain positive for much larger values of the adversary budget B .

Information Leakage.

In general we would recommend that any implementation use $\epsilon \leq 2.08$ so that the outcome o leaks at most 3 bits of information about the user’s password. We note that if the adversary breaches only the client then he would still need to complete a brute-force search over all passwords $pwd \in \mathcal{P}$ to extract this information from o . In our demo implementation we selected security parameter $\epsilon = 1.609$ so that the outcome o of **SelPreds** will leak at most 2.32 bits of information about the user’s password. Future versions might help the client to tune ϵ himself. We selected $\epsilon = 1.609$ because it yielded the highest reductions against an adversary with a larger budget B .

Adversary Budget.

If we have an accurate estimation of the adversary’s budget B then we should compute the **SelPreds** distribution using our algorithm from §7 before implementing Client-CASH. However, if this estimation is not precise then it may be better to use the exponential mechanism from §6.

Password Distribution.

We only show how to optimize Client-CASH under the assumption that users choose passwords uniformly at random. We hope that future work will extend this analysis

⁶In Figure 3a, the two round exponential mechanism will outperform the three round mechanism when $\frac{B}{k'|\mathcal{P}|} \leq 0.5$ for security parameters $\epsilon \geq 1.6$.

Figure 3: Comparison of the percent reduction in passwords between two-round mechanisms (blue) and three-round mechanisms (orange) as a function of B with a normalizing factor of $\frac{1}{k'|D|} = \frac{C_H}{C_{srv} |D|}$ for various values of ϵ . The results for the exponential mechanism are shown in (a), while the optimal mechanism is shown in (b).

to more general password distributions. Until then Client-CASH should only be adopted by users who select uniformly random passwords from some space.

Amortization of Costs.

Our guarantee that the client’s computational costs during authentication at most C_{srv} only holds in expectation. Of course the user may get unlucky when we run $\text{SelPreds}(pwd)$ and end up with the maximum stopping time $\mathcal{S}(pwd, o) = n$. We note that if key-stretching was performed by a trusted third party (e.g., LastPass) then these costs could be amortized over different users. If key-stretching is only performed on the user’s local computer then we could give an unlucky user the option to speed up the key-stretching process by memorizing an additional random character and appending it to his password (e.g., see [11]).

8.1 Demo Implementation

A demo implementation of the three round exponential Client-CASH mechanism is available on client-cash.github.io⁷. We implemented our algorithms in Javascript so the key-stretching is done locally by the user’s web browser. In our implementation we used $k = 100,000$ iterations of the SHA256 hash function H . Thus, Client-CASH perform between 100,000 and 300,000 iterations of the SHA256 hash function to derive our final password. We selected $n = 3$ because the exponential mechanism performs better with $n = 3$ rounds of hashing (see Figure 3a).

9. CONCLUSIONS AND FUTURE WORK

In this work, we have introduced Client-CASH as a novel client-side key stretching mechanism which gives better security guarantees than traditional key-stretching procedures. We

⁷The current demo implementation does not permanently save state for future authentication sessions though we anticipate adding this functionality soon.

make several contributions. First, we show how to introduce randomness into a client-side key stretching algorithms through the use of halting predicates which are selected randomly at the time of account creation. Second, we formalize the problem of finding the optimal running time distribution subject to certain cost constraints for the client and certain security constraints on the halting predicates. Finally, we demonstrate that Client-CASH can reduce the adversary’s success rate by up to 21% (in the $n = 3$ round case). These results demonstrate the promise of the Client-CASH mechanism. In the future we hope that researchers will extend our analysis of Client-CASH to handle non-uniform password distributions. We also speculate that it may be possible to obtain even larger reductions in the adversary’s success rate by using predicate sequences $\mathcal{O}$ with more rounds of hashing $n > 3$. Future work might extend our analysis of Client-CASH to explore this possibility.

10. REFERENCES

- [1] R. Biddle, S. Chiasson, and P. van Oorschot. Graphical passwords: Learning from the first twelve years. *ACM Computing Surveys (CSUR)*, 44(4):19, 2012.
- [2] A. Biryukov, D. Dinu, and D. Khovratovich. Fast and tradeoff-resilient memory-hard functions for cryptocurrencies and password hashing. Cryptology ePrint Archive, Report 2015/430, 2015. <http://eprint.iacr.org/>.
- [3] J. Blocki, M. Blum, and A. Datta. Naturally rehearsing passwords. In K. Sako and P. Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013*, volume 8270 of *Lecture Notes in Computer Science*, pages 361–380. Springer Berlin Heidelberg, 2013.
- [4] J. Blocki and A. Datta. Cash: A cost asymmetric secure hash algorithm for optimal password protection. *CoRR*, abs/1509.00239, 2015.

- [5] J. Blocki, S. Komanduri, L. F. Cranor, and A. Datta. Spaced repetition and mnemonics enable recall of multiple strong passwords. In *22nd Annual Network and Distributed System Security Symposium, NDSS 2015, San Diego, California, USA, February 8-11, 2014*, 2015.
- [6] J. Blocki, S. Komanduri, A. Procaccia, and O. Sheffet. Optimizing password composition policies. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 105–122. ACM, 2013.
- [7] J. Blocki and A. Sridhar. Client-cash: Protecting master passwords against offline attacks. *CoRR*, abs/, 2016.
- [8] J. Bonneau. The science of guessing: analyzing an anonymized corpus of 70 million passwords. In *Security and Privacy (SP), 2012 IEEE Symposium on*, pages 538–552. IEEE, 2012.
- [9] J. Bonneau, C. Herley, P. C. van Oorschot, and F. Stajano. The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In *IEEE Symposium on Security and Privacy*, pages 553–567. IEEE, 2012.
- [10] J. Bonneau and S. Preibusch. The password thicket: technical and market failures in human authentication on the web. In *Proc. of WEIS*, volume 2010, 2010.
- [11] J. Bonneau and S. Schechter. "toward reliable storage of 56-bit keys in human memory". In *Proceedings of the 23rd USENIX Security Symposium*, August 2014.
- [12] X. Boyen. Halting password puzzles. In *Proc. Usenix Security*, 2007.
- [13] S. Boztas. Entropies, guessing, and cryptography. *Department of Mathematics, Royal Melbourne Institute of Technology, Tech. Rep*, 6, 1999.
- [14] J. G. Brainard, A. Juels, B. Kaliski, and M. Szydło. A new two-server approach for authentication with short secrets. In *USENIX Security*, volume 3, pages 201–214, 2003.
- [15] S. Brostoff and M. Sasse. Are Passfaces more usable than passwords: A field trial investigation. In *People and Computers XIV-Usability or Else: Proceedings of HCI*, pages 405–424, 2000.
- [16] J. Camenisch, A. Lysyanskaya, and G. Neven. Practical yet universally composable two-server password-authenticated secret sharing. In *Proceedings of the 2012 ACM conference on Computer and Communications Security*, pages 525–536. ACM, 2012.
- [17] M. Dürmuth and T. Kranz. On password guessing with gpus and fpgas.
- [18] C. Dwork, A. Goldberg, and M. Naor. On memory-bound functions for fighting spam. In *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 426–444. Springer, 2003.
- [19] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography*, pages 265–284. Springer, 2006.
- [20] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407, 2013.
- [21] D. Florencio and C. Herley. A large-scale study of web password habits. In *Proceedings of the 16th international conference on World Wide Web*, pages 657–666. ACM, 2007.
- [22] C. Forler, S. Lucks, and J. Wenzel. Catena: A memory-consuming password scrambler. *IACR Cryptology ePrint Archive*, 2013:525, 2013.
- [23] I. Jermyn, A. Mayer, F. Monrose, M. K. Reiter, and A. D. Rubin. The design and analysis of graphical passwords. In *Proceedings of the 8th conference on USENIX Security Symposium - Volume 8*, pages 1–1, Berkeley, CA, USA, 1999. USENIX Association.
- [24] B. Kaliski. Pkcs# 5: Password-based cryptography specification version 2.0. 2000.
- [25] K. Kim, S. Lee, D. Hong, and J.-C. Ryou. Gpu-accelerated password cracking of pdf files. *KSII Transactions on Internet and Information Systems (TIIS)*, 5(11):2235–2253, 2011.
- [26] S. Komanduri, R. Shay, P. Kelley, M. Mazurek, L. Bauer, N. Christin, L. Cranor, and S. Egelman. Of passwords and people: measuring the effect of password-composition policies. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pages 2595–2604. ACM, 2011.
- [27] U. Manber. A simple scheme to make passwords based on one-way functions much harder to crack. *Computers & Security*, 15(2):171–176, 1996.
- [28] J. Massey. Guessing and entropy. In *Information Theory, 1994. Proceedings., 1994 IEEE International Symposium on*, page 204. IEEE, 1994.
- [29] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 94–103. IEEE, 2007.
- [30] R. Morris and K. Thompson. Password security: A case history. *Communications of the ACM*, 22(11):594–597, 1979.
- [31] C. Percival and S. Josefsson. The scrypt password-based key derivation function. 2012.
- [32] J. Pliam. On the incomparability of entropy and marginal guesswork in brute-force attacks. *Progress in Cryptology-INDOCRYPT 2000*, pages 113–123, 2000.
- [33] N. Provos and D. Mazières. Bcrypt algorithm.
- [34] D. Reichl. Keepass password safe, 2013. *Retrieved July*, 10, 2013.
- [35] B. Ross, C. Jackson, N. Miyake, D. Boneh, and J. C. Mitchell. Stronger password authentication using browser extensions. In *Usenix security*, pages 17–32. Baltimore, MD, USA, 2005.
- [36] R. Shay, P. Kelley, S. Komanduri, M. Mazurek, B. Ur, T. Vidas, L. Bauer, N. Christin, and L. Cranor. Correct horse battery staple: Exploring the usability of system-assigned passphrases. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, page 7. ACM, 2012.