# DALock: <u>D</u>istribution <u>A</u>ware Password Throttling

Jeremiah Blocki
Purdue University
jblocki@purdue.edu

Wuwei Zhang
Purdue University
zhan1015@purdue.edu

## ABSTRACT

Large-scale online password guessing attacks are wide-spread and continuously qualified as one of the top cyber-security risks. The common method for mitigating the risk of online cracking is to lock out the user after a fixed number ($K$) of consecutive incorrect login attempts. Selecting the value of $K$ induces a classic security-usability trade-off. When $K$ is too large a hacker can (quickly) break into a significant fraction of user accounts, but when $K$ is too low we will start to annoy honest users by locking them out after a few mistakes. Motivated by the observation that honest user mistakes typically look quite different than the password guesses of an on-line attacker, we introduce DALock a *distribution aware* password lockout mechanism to reduce user annoyance while minimizing user risk. As the name suggests, DALock is designed to be aware of the frequency and popularity of the password used for login attacks while standard throttling mechanisms (e.g., $K$-strikes) are oblivious to the password distribution. In particular, DALock maintains an extra "hit count" in addition to "strike count" for each user which is based on (estimates of) the cumulative probability of *all* login attempts for that particular account. We empirically evaluate DALock with an extensive battery of simulations using real world password datasets. In comparison with the traditional $K$-strikes mechanism we find that DALock offers a superior security/usability trade-off. For example, in one of our simulations we are able to reduce the success rate of an attacker to 0.05% (compared to 1% for the 10-strikes mechanism) whilst simultaneously reducing the unwanted lockout rate for accounts that are not under attack to just 0.08% (compared to 4% for the 3-strikes mechanism).

## KEYWORDS

Authentication Throttling; Password; Dictionary Attack

## 1 INTRODUCTION

An online password attacker repeatedly attempts to login to an authentication server submitting a different guess for the target user's password on each attempt. The human tendency to pick weak ("low-entropy") passwords has been well documented e.g., [5]. An untargetted online attacker will typically submit the most popular password choices consistent with the password requirements (e.g., "Password1"), while a targeted attacker [53] might additionally incorporate background knowledge about the specific target user (e.g., birthdate, phone number, anniversary etc...). To protect user's against online attackers most authentication servers incorporate some form of throttling mechanism. In particular, the $K$-strikes mechanism temporarily locks a user's account if $K$-consecutive incorrect passwords are attempted within a predefined time period (e.g., 24 hours). Setting the lockout parameter $K$ induces a classic security-usability trade-off. Selecting small values of $K$ (e.g., $K = 3$) provides better protection against online attackers, but may result

in many unwanted lockouts when an honest user miss-types (or miss-remembers) their password. Selecting a larger value of $K$ (e.g., $K = 10$) will reduce the unwanted lockout rate, but may increase vulnerability to online attacks.

Bonneau et al. [6] considered many proposed replacements for password authentication finding that all proposals have some drawbacks compared when compared with passwords. For example, passwords are easy to revoke unlike biometrics. Similarly, hardware tokens are expensive to deploy and require users to carry them around. By contrast, passwords are easy to deploy and do not require users to carry anything around. Put simply we have not found a "silver bullet" replacement for passwords. Thus, despite all of their short-comings (and many attempts to replace them) passwords will likely remain entrenched as the dominant form of authentication on the internet [26]. Thus, protecting passwords against online attacks without locking out legitimate users remains a crucial challenge for the foreseeable future [36, 44, 46].

One approach to protect user's against online guessing attacks is to adopt strict password composition policies to prevent user's from selecting weak passwords. However, it has been well documented that users dislike restrictive policies and often respond in predictable ways [29]. Another defense is to store cookies on the user's device to prove that the next login attempt came from a known device. Similarly, one can also utilize features such as IP address, geographical location, device and time of day [18, 23, 40] to help distinguish between malicious and benign login attempts. While these features can be helpful indicators they are not failproof. Honest users oftentimes travel and login from different devices at unusual times. Similarly, an attacker may attempt to mimic login patterns of legitimate users e.g., using a botnet the online attacker can submit guesses from a vide variety of IP addresses and geographical locations.

### 1.1 Contributions

We introduce DALock, a novel <u>D</u>istribution <u>A</u>ware throttling mechanism which can achieve a better balance between usability and security. The key intuition behind DALock is to base lockout decisions on the *popularity* of the passwords that are being guessed. An online attacker will typically want to attempt to login with the most popular passwords to maximize his chances of success. By contrast, when an honest user miss types (or miss remembers) his password he will typically not be globally popular password. In addition to keeping track of $K_u$ (the number of consecutive incorrect login attempts), DALock keeps track of a "hit count" $\Psi_u$ for each user $u$, where $\Psi_u$ intuitively represents the cumulative probability mass of all incorrect login attempts for user $u$'s account. When $\Psi_u$ exceeds the threshold $\Psi$ we decide to lock the account.

*Example 1: Usability.* **Figure** 1 compares the usability of DALock with the standard $K = 3$ strikes mechanism. In this example scenario our user John Smith registers an account with the somewhat complicated password "J.S.UsesStr0ngpwd!" based on the story "John Smith uses a strong password." Later when he tries to login he remembers the basic story, but not the exact password. Did he use his first name and his last name? With or without abbreviation? Did he add a punctuation mark at the end? Which letters are capitalized? If we use the 3-Strike mechanism John Smith will be locked out quickly e.g., after trying the incorrect password guesses "JohnUseStrongPassword," "JohnUsesStrongPassword" and "JohnUsesStrongpwd." However, since none of these passwords are overly popular DALock would allow our user to continue attempting to login until he recovers the correct password.

*Example 2: Security.* **Figure** 2 compares DALock with the $K = 10$ strike mechanism. In this scenario our user registers an account with a weak password "letmein." Because the password is globally popular it is likely that an online attacker will attempt this password within the first 10 guesses and break into the user's account. By contrast, DALock will quickly lockdown the account after the attacker submits two globally popular passwords.

We evaluate DALock empirically by simulating an authentication server in the presence of an online password attacker. We compare DALock with the the traditional $K$ strikes mechanism for $K \in \{3, 10\}$. Our experiments show that when the hit count threshold $\Psi$ is tuned appropriately that DALock significantly outperforms our $K$-strike mechanisms. In particular, when user accounts are under attack we find that the fraction of accounts that are compromised is significantly lower for DALock than either $K$ strikes mechanism — even for the strict $K = 3$ strikes policy. We also evaluate the unwanted lockout rate of user accounts which are not under attack. We find that the unwanted lockout rate for DALock is much lower compare to $K = 3$-strikes mechanism. The unwanted lockout rate for DALock and the more lenient $K = 10$-strikes mechanism were comparable. A more detailed description of our experiments can be found in **section** 6.

To deploy DALock we need a way to estimate the frequency of each incorrect login attempt to update $\Psi_u$. We propose two methods for doing this: password strength meters (e.g. ZXCVBN[54]) and a differentially private count sketch data structure. Our empirical experiments show DALock outperforms traditional lockout mechisms with either approach. However, DALock performs best when we instantiate with a differentially private count sketch. On a positive note we show that even if the differentially private count sketch is trained on a small subset of user passwords that the estimates will still be high enough for DALock to be effective.

## 2 RELATED WORKS AND BACKBROUNDS

### 2.1 Authentication Throttling

**K-strike Mechanism** K-strike mechanism is a straight-forrward implementation for authentication throttling. As its name suggests, throttling occurs when $K$ consecutive incorrect login attempts are detected. To reduce the cost of expensive overhead caused by unwanted throttling, Brostoff [7] et.al suggest setting threshold $K$ to be 10 instead of 3. They argue the increment risk is limited when strong
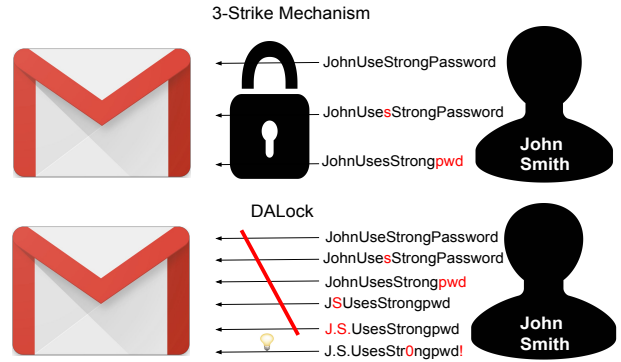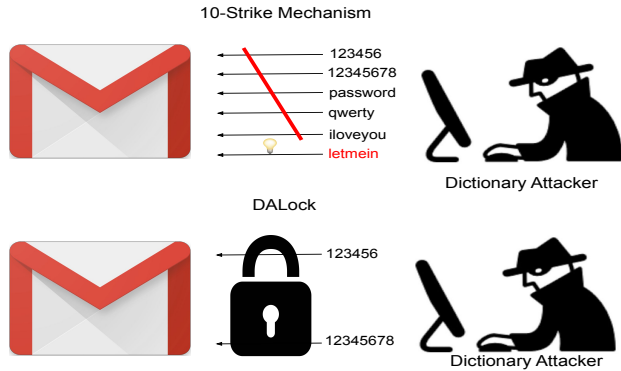


Figure 1: Usability Comparison



Figure 2: Security Comparison

password policy is enforced. However, this argument is challenged by empirical analyses of password composition policies [29][4]. Many password composition policies do not rule out all low entropy password choices. For instance, it turns out that banning dictionary words does not increase entropy as expected. [29]

**Feature-Based Mechanism** To improve performance, modern throttling mechanisms[40][23] often times use features such geographical location, IP-address, device information, and etc in addition to the correctness of attempting password. These features can be used to train sophisticated machine learning models to help distinguish between malicious and benign login attempts [18]. DALock takes an orthogonal approach and relies instead on the popularity of the password guesses. One can combine those models with a rigious throttling system for a better performance.

**Password-Distribution Aware Throttling** In an independent line of work Tian et al. [56] developed an IP-based throttling mechanism which exploits differences between the distribution of honest login attempts and attacker guesses. In particular, they propose to "silently block" login attempts from a particular IP address $a$ if the system detects too many popular passwords being submitted from that IP address. In more detail StopGuessing uses a data-structure called the binomial ladder filter [42] to (approximately) track the frequency $F(p)$ of each incorrect password guess. For each IP address the StopGuessing protocol maintains an associated counter $I_a = \sum_{p \in \mathcal{P}} F(p)$ where $\mathcal{P}$ is a list of incorrect password guesses that

have been (recently) submitted from that IP address — $I_a$ can be updated without storing $\mathcal{P}$ explicitly. Intuitively (and oversimplifying a bit) if $I_a > T$ then login attempts from address $a$ are silently blocked i.e., even if the attacker (or honest user) submits a correct password the system will respond that authentication fails. The authors also suggest protecting accounts with weak passwords by setting a user specific threshold $T(F(u_p))$ based on the strength $F(u_p)$ of the password $pw_u$ of user $u$. Now if $I_a > T(F(u_p))$ then the system will silently reject any password from address $a$. Both StopGuessing and DALock exploit differences between the distribution of user passwords and attacker guesses. One of the key difference is that StopGuessing focuses on identifying malicious IP addresses (by maintaining a score $I_a$) while DALock focuses on protecting individual accounts by maintaining a "hit-count" parameter $\Psi u$ for each user u. There are several other key differences between the two approaches. First, in DALock the goal of our frequency oracle (e.g., count-sketch, password strength meter) is to estimate the total fraction of users who have actually selected that particular password — as opposed to estimating the frequency with which that password has been *recently* submitted as a incorrect guess. Second, DALock does not require silent blocking of login attempts which could create usability concerns if an honest user is silently blocked when they enter the correct password.

## 2.2 Passwords

**Password Distribution** Password distribution naturally represents the chance of success in the setting of statcial guessing attacks. Password distribution has been extensively studied since last decades[17][32]. Using leaked password corpora[39][31][5] is a straight forward way to describe the distribution of passwords. In recent works of Wang et al. [50–52] argue that password distributions follows Zipf's law i.e., leaked password corpora nicely fit Zipf's law distributions. Blocki et al. [3] later found that Zipf's law nicely fits the Yahoo! password frequency corpus [2, 5].

**Password Typos** To test the usability of DALock, it's crucial to reasonably simulate users' mistakes. Recent studies[11, 12] from Chatterjee et. al have summarized probabilities of making (various of) typos when one enters his or her password based on users' studies. Based on the empirically measured data, they purposed two typo-tolerant authentication without sacrificing security. If fact, such mechanism has already been deployed in industry[57][20].

## 2.3 Eliminating Dictionary Attacks

**Increasing Cost of Authentcation** Pinkas and Sanders [37] proposed the use of puzzles (e.g., proofs of work or CAPTCHAs) as a way to throttle online password crackers. CAPTCHAs are hard AI challenges meant to distinguish people from bots [48]. For example, reCAPTCHA [49] has been widely deployed in online web services such as Google, Facebook, Twitter, CNN, and etc. Assuming that CAPTCHAs are only solvable by people, one can mitigate automated online dictionary attacks without freezing users' accounts [8, 9]. However, an attacker can always pay humans to solve these CAPTCHA challenges[38]. Increasingly sophisticated CAPTCHA solvers [19, 55] powered by neural networks make it increasingly difficult to design CAPTCHA puzzles that are also

easy for a human to solve. Golla et al. [21] proposed a fee-based password verification system where a small deposit is necessary to authenticate, which is refunded after successful authentication. A password cracker risks loosing its deposit if it is not able to guess the real password.

**Eliminating Popular Passwords** One mediation for dictionary attacks is eliminating the existence of weak or popular passwords. Schechter et. al [41] show that it is possible to forbid the existence of over popular password by maintaining the password distribution securely. Industrial solutions such as "Have I been pwned?" [25] and "Password CheckUp" [45] " prevent users to choose weak passwords based on data breaches. Password strength meters such as ZXCVBN [54] are also been widely deployed to help users choosing stronger passwords.

## 2.4 Privacy Perserving Aggregate Statistics Releasing.

DALock relies on the distribution of passwords to perform throttling. Storing/Releasing aggregate statistics naively often causes privacy leakage [34] [35]. To answer the challenge, Cynthia Dwork purposed Differential Privacy [15] for aggregated data releasing. Informally speaking, differential private algorithm makes powerful adversaries unable of telling the existence of a record in the dataset. We defer the formal definition of differential privacy to section 3.2. Blocki et.al released the statics of Yahoo! dataset consist of 70 millions of passwords. [2]. Recent work by Naor et at [33] also demonstrates that releasing the distribution of password privately is feasible. In industry, Differential Privacy has been considered as the golden tool for various of tasks [16, 27, 43].

## 3 PRELIMINARIES

### 3.1 Count Sketch

Count (Median) Sketch [10] and it's variants are widely in the tasks for finding frequent items such as popular passwords [33], homepage settings [16], and frequently used chat emojis [27]. In this work, we uses CS as a tool for password popularity estimation. Formally speaking, we define Count Sketch as follows

*Definition 3.1 (Count Sketch [13] [10]).* A Count Sketch of state $\sigma : \mathsf{R}^{d \times w} \times \mathsf{R}$ is represented by a two-dimensional $d \times w$ array counts CS, a total frequency counter CS.T, and d + 1 hash functions $(h_1, \ldots, h_d, h_\pm)$ chosen uniformly at random from a pairwise-independent family.

$$h_1 \cdots h_d : \{pw\} \to \{1 \cdots w\}$$
$$h_\pm : \{pw\} \to \{1, -1\}$$

In this work we consider the following four classic APIs for Count Sketch: initialize, Add, Estimate, and TotalFreq. Additionally, we consider an extra operation DP which is used to construct differentially private Count Sketch from a standard one.

$\sigma_0 \leftarrow$ **Initialize**$(d, w)$ : This API initialize and return a Count Sketch of state $0^{d \times w} \times 0$, i.e. an all zero table.

$\sigma_{new} \leftarrow$ **Add**$(pw, \sigma)$**:** Add operation updates the stored frequency count password $pw$ based on a CS state $\sigma$, and outputs the updated state $\sigma_{new}$.

In addiition, given a multiset $\mathcal{D}_{\mathcal{U}} = \{pw_1, ..., pwd_N\}$, we use the following notation $\sigma_{\mathcal{D}_{\mathcal{U}}} = Add(\mathcal{D}_{\mathcal{U}}, \sigma)$ $= Add(pw_1, Add(\{pw_2, ..., pw_N\}, \sigma))$ to ease presentation. Further more, we omit subscript $\mathcal{U}$ and simply use $\sigma$ to denote $\sigma_{\mathcal{D}_{\mathcal{U}}}$.

**Estimate$(pw, \sigma)$** : This interface returns the estimated frequency of password $pw$ based on the given Count Sketch State $\sigma$.

To implement DALock with high accuracy, we want the estimator has the following correctness Property: Estimate$(pw, \sigma) \approx$ F$(pw, \mathcal{D}_{\mathcal{U}})$.

**TotalFreq$(\sigma)$** : This opeartion returns the total number of passwords based on state $\sigma$.

Based on the above definition, we denote the *estimated popularity of password pw* by $\sigma$ with $\mathsf{p}(pw, \sigma) = \frac{\mathrm{Estimate}(pw, \sigma)}{\mathrm{TotalFreq}(\sigma)}$. For the rest of discussions, we sometimes omit $\sigma$ when there is no ambiguity to simplify presentation. e.g. $\mathsf{p}(pw) = \mathsf{p}(pw, \sigma)$. In addition, we allow the above APIs to take a set of passwords as argument and return the summed results. i.e.. $\mathsf{p}(S) = \sum_{pw \in S} \mathsf{p}(S)$.

$\sigma_{dp} \leftarrow$ **DP$(\epsilon, \sigma)$** :This function outputs an differentially private state $\sigma_{dp}$ of $\sigma$ with privacy budget $\epsilon$.

### 3.2 Differential Privacy

Differential Privacy [15] is one of the industrial golden standard tools for private aggregated statical releasing. Intuitively speaking, individual record has limited impact on computing the final results to be published if one applies differential. For instance, consider the following two password datasets: A password dataset$\mathcal{D}_{\mathcal{U}}$ consists of users' passwords , and it's neighboring dataset $\mathcal{D}_{\mathcal{U} - pw_u}$ obtained by removing user u's password $pw_u$ from $\mathcal{D}_{\mathcal{U}}$. Differential privacy guarantees that *with high probability* the published results based on $\mathcal{D}_{\mathcal{U}}$ and $\mathcal{D}_{\mathcal{U} - pw_u}$ are the same. Therefore, adversary is not able to infer the existence of $u$ in dataset $\mathcal{D}_{\mathcal{U}}$.

In this work, we adopt differential privacy Count Sketch to reduce the risk of privacy leakage. Based on the our notion of Count Sketch, one can define differential privacy as follows

*Definition 3.2 ($\epsilon$-Differential Privacy [15]).* A randomized mechanism $\mathcal{M}$ gives $\epsilon$-differential privacy if for any pair of neighboring datasets $\mathcal{D}_{\mathcal{U}}$ and $\mathcal{D}'_{\mathcal{U}}$, and any $\sigma \in Range(\mathcal{M})$,

$$\Pr\left[\mathcal{M}(\mathcal{D}_{\mathcal{U}}) = \sigma\right] \le e^{\epsilon} \cdot \Pr\left[\mathcal{M}(\mathcal{D}'_{\mathcal{U}}) = \sigma\right].$$

We consider two datasets $\mathcal{D}_{\mathcal{U}}$ and $D'$ to be neighbors i.f.f. either $\mathcal{D}_{\mathcal{U}} = \mathcal{D}'_{\mathcal{U}} + pw_u$ or $\mathcal{D}'_{\mathcal{U}} = D + pw_u$, where $\mathcal{D}_{\mathcal{U}} + pw_u$ denotes the dataset resulted from adding the tuple $pw_u$(a new password) to the dataset $\mathcal{D}_{\mathcal{U}}$. We use $\mathcal{D}_{\mathcal{U}} \simeq \mathcal{D}'_{\mathcal{U}}$ to denote two neighboring datasets. This protects the privacy of any single tuple, because adding or removing any single tuple results in $e^{\epsilon}$-multiplicative-bounded changes in the probability distribution of the output. If any adversary can make certain inference about a tuple based on the output, then the same inference is also likely to occur even if the tuple does not appear in the dataset.

**Laplace Mechanism.** The Laplace mechanism is a classic tool achieve differential privacy. It computes a differential private state $\sigma$ on the dataset $\mathcal{D}_{\mathcal{U}}$ by adding a random noise. The magnitude of the noise depends on $\mathrm{GS}_{\sigma}$, the *global sensitivity* or the $L_1$ sensitivity of

$\sigma$. $\mathrm{GS}_{\sigma}$ quantify the maximum impact on $\sigma$ by adding or removing any record.

**Differentially Private Count Sketch** Given a CS of state $\sigma$, adding (removing) any password $pw_u$ to(from) it can result in at most d + 1 changes for $l_1$ norm. Because each $pw_u$ contributes to d entries in the $d \times w$ table and total count. Therefore, To release $\sigma$ with privacy budget $\epsilon$, it suffices to add $\mathsf{Lap}(\frac{d+1}{\epsilon})$ to all entries in $\sigma$.

**Privacy-Preserving Password Corpus Relasing Mechanism** Naor et.al[33] purposed an algorithm to release password distribution using local differential privacy. In our work, we focused on a centralized version of differential privacy which is expected to have less noisy compare to local setting. StopGuessing[56] uses a binomial ladder to identify "heavy hitters" (popular passwords), though the data-structure does not provide any formal privacy guarantees such as differential privacy. The data-structure is not suitable for DALock as it provides a binary classification i.e., either the password is a "heavy hitter" or it is not. For DALock requires a more fine grained estimate of a passwords popularity.

### 3.3 Notation Summary

In this section, we summarize frequently used notations in this paper across all sections in **Table** 1. For a password $pw \in \mathcal{P}$ we use $\mathsf{P}(pw)$ to denote the probability each user selects the password $pw$. We assume that there is some underyling distribution over user passwords and use $\mathsf{P}(pw)$ to denote the probability of the password $pw \in \mathcal{P}$. It will be convenient to assume that all passwords $\mathcal{P} = \{pw_1, pw_2, \ldots\}$ are sorted in descending order of probability i.e., so that $\mathsf{P}(pw_1) \ge \mathsf{P}(pw_2) \ldots$. We will use $pw_r = \mathsf{P}(pw_r)$ to denote the probability of the $r$th most likely password in the distribution.

We use $\mathcal{U} = \{u_1, \ldots, u_N\}$ to denote a set of $N$ users and $D_{\mathcal{U}} \subseteq \mathcal{P}$ is a multiset of user passwords i.e., $D_{\mathcal{U}} = \{pw_{u_1}, \ldots, pw_{u_N}\}$. We typically view $D_{\mathcal{U}}$ as $N$ independent samples from an underlying distribution over $\mathcal{P}$ and write $\mathsf{F}(pw, \mathcal{D}_U) = \left|\{i : pw_{u_i} = pw\}\right|$ to denote the number of times the password $pw$ was observed in our sample. We often omit $D_{\mathcal{U}}$ when the dataset is clear from context and simply write $\mathsf{F}(pw)$.

We remark that $\mathsf{P}(pw) = \frac{\mathbb{E}[\mathsf{F}(pw, \mathcal{D}_U)]}{N}$ and thus for popular passwords we expect that the estimate $\mathsf{P}(pw) \approx \frac{\mathsf{F}(pw, \mathcal{D}_U)}{N}$ will be accurate for sufficiently large $N$. However, because the underlying password distribution is unknown and an authentication server cannot store a plaintext encoding of $\mathcal{D}_{\mathcal{U}}$ we will often use other techniques to estimate $\mathsf{P}(pw)$ and/or $\mathsf{F}(pw, \mathcal{D}_U)$. In particular, we consider a Count (Median) Sketch data structure CS trained on $\mathcal{D}_U$ (or a small subsample of $\mathcal{D}_U$) which allows us to generate an estimate $\mathsf{p}(pw)$ for the popularity of each password. Similarly, we can also use password strength meters to estimate $\mathsf{P}(pw)$.

## 4 THE DALOCK MECHANISM

In this section, we present the DALock mechanism, discuss how DALock might be implemented and the strategies that an attacker might use when DALock is deployed. Intuitively, DALock punishes incorrect password guesses more harshly if the guessed password $pw$ is overly popular since an attacker will want to submit popular

| Notation | Description |
|---|---|
| $\mathcal{A}$ | $\mathcal{A}$dversary |
| $\mathcal{U}$ | The set of $\mathcal{U}$sers |
| $u$ | A user $u \in \mathcal{U}$ |
| $\mathcal{P}$ | The set of all potential user $\mathcal{P}$asswords |
| $\mathcal{D}_{\mathcal{U}} \subseteq \mathcal{P}$ | a multiset of $N$ sampled passwords for users $u_1, \ldots, u_N \in \mathcal{U}$ |
| $pw_u$ | User $u$'s password |
| $pw_r$ | The $r$'th most likely password in $\mathcal{D}_{\mathcal{U}} \subseteq \mathcal{P}$ |
| CS | Count (Median) Sketch data structure |
| $F(pw, \mathcal{D}_{\mathcal{U}})$ | Frequency of password $pw$ in dataset $\mathcal{D}_{\mathcal{U}}$ |
| $P(pw)$ | Empirical probability of password $pw$ |
| Estimate($pw$) | Estimated frequency of password $pw$ |
| p($pw$) | Estimated probability of password $pw$ |
| $\Psi$ | Hit count threshold |
| $\Psi_u$ | Remaining hit count threshold of user u. The account gets locked out if $\Psi_u$ reaches $\Psi$ |
| $K$ | traditional $K$-strike threshold. |
| $K_u$ | Remaining strike threshold on $u$'s account. The account gets locked if $K_{\mathcal{U}}$ exceeds $K$. |

**Table 1: Notation Summary**

password guesses to maximize his chances of cracking the user password.

### 4.1 DALock

In the classic $K$-strike throttling mechanism we keep track of a parameter $K_u$ which tracks the number of consecutive incorrect login attempts for each user $u$. After each consecutive login attempt the parameter is $K_u$ and the parameter $K_u$ is reset to 0 whenever $u$ authenticates successfully. If we ever have $K_u \geq K$ then throttling mechanism kicks in and the authentication server will lock down the account until the user takes some action[1].

The key-idea behind DALock is to additionally maintain an extra "hit count" variable $\Psi_u$ for each user $u$. Intuitively, $\Psi_u$ measures the total probability mass of all incorrect password guesses submitted for user $u$. Initially, when a new user registers we will have $\Psi_u = 0$ (and $K_u = 0$). After each attempted login with an incorrect password $pw \neq pw_u$ the hit count is incremented so that $\Psi_u += p(()pw)$. Here, p($()pw$) denotes an estimate for the probability of the password $pw$ so that incorrect passwords are punished more severely when $pw$ is an overly popular password. Unlike the consecutive strikes parameter $K_u$ which is reset to 0 after each successful login, the hit count parameter can only be incremented. DALock throttles $u$'s account if the "hit count" exceeds $\Psi$ (i.e., $\Psi_u \geq \Psi$) or if there are too many consecutive mistakes (i.e., $K_u \geq K$) For example, suppose that the (estimated) probability of the passwords "aaa," "bbb" and "ccc" were 3%, 1.7% and 0.8%. If a user registers with a password "ddd" and then attempts to login with the previous three passwords then $\Psi_u$ will be set to $0.055 = 0.03 + 0.017 + 0.008$.

Each time the user (or attacker) attempts to login with a password $pw$ the response will either be (1) "locked" if $\Psi_u \geq \Psi$ or if

$K_u \geq K$, (2) "correct" if the the guessed password matches the user password i.e., $pw = pw_u$[2] or (3) "incorrect password" otherwise. We demonstrate the login flow in **Algorithm** 1, Appendix. We remark that the authentication server could intentionally blur this distinction between cases (1) and (3), but that this comes at a usability cost e.g., an honest user would be annoyed if they were repeatedly informed that their password is incorrect whenever the account is actually locked.

**Remark:** One could optionally consider initializing the hit count parameter $\Psi_u$ based on the strength of the user's password. For example, if $u$ registers with a weak password then we might initialize $\Psi_u = \Psi/2$ for stronger protection i.e., so that the account is locked down faster. Similarly, a user with a strong password might be awarded by setting $\Psi_u = \Psi$. However, because $\Psi_u$ and $K_u$ are stored on the authentication server this would leak information about the strength of $pw_u$ to an offline attacker e.g., if an offline attacker sees that $\Psi_u = \Psi/2$ he might reasonably infer that the user picked a weak password. [3]

### 4.2 DALock **Authentication Server**

To implement DALock we need an efficient way to estimate the probability p($()pw$) of each incorrect password $pw$. We consider several instantiations of this frequency oracle. One option is to use password strength meters such as ZXCVBN or more sophisticated password cracking models e.g., Markov Models, Probabilistic Context Free Grammars, or Neural Networks. Another naive approach would be to simply maintain a plaintext list of all user passwords along with their frequencies. However, this approach is inadvisable due to the risk of leaking this plaintext list. Herley and Schechter [41] proposed the use of the Count-Sketch data-structure which would allow us to estimate the frequency of each password without explicitly storing a plaintext list although there are no formal privacy guarantees to this approach. We chose to adopt a Differentially Private Count-Median-Sketch. The authentication server initializes the Count-Sketch $\sigma_{dp} \leftarrow DP(\epsilon, \sigma)$ by adding Laplace Noise to preserve $\epsilon$-differential privacy and each time a new user $u$ registers a new password $pw_u$ would be added to the Count Sketch.

We remark that maintaining a Differentially Private Count-Sketch has many other potentially beneficial applications e.g., one could use the Count-Sketch to ban weak passwords [41] and/or to help identify IP addresses associated with malicious online attacks [56]. One disadvantage is that the attacker will also be able to view the Count-Sketch data-structure if the data-structure is leaked. The usage of differential privacy helps to minimize these risks. Intuitively, differential privacy hides the influence of any individual password ensuring that an attacker will not be able to use the Count-Sketch data-structure to help identify any unique passwords. However, an attacker may still be able to use the data-structure to learn that a

---

[1]For example, the user might be asked to resetting his password via e-mail or wait for some fixed amount of time. In some settings the user might simply be asked to solve a CAPTCHA challenge. The latter approach has some usability advantages and security drawbacks e.g., a malicious password might pay to solve the CAPTCHA challenges so that he can continue attempting to guess the user's password

[2]To ease presentation we omit the description of the password hashing algorithm when we describe the authentication server. In practice, we recommend that the authentication server only stores salted password hashes using a moderately expensive key derivation function to increase guessing costs for an offline attacker.

[3]One could potentially avoid storing $\Psi_u$ unencrypted if one is willing to implement a silent lockout policy where the user cannot distinguish between an incorrect guess and a locked account, but we wish to avoid solutions that blur this distinction.

particular password is globally popular (without linking that password to a particular user). We argue that this is not a major risk as most attackers will already know about globally popular passwords e.g., from prior breaches.

## 5 EXPERIMENTAL DESIGN

We evaluate the performance of DALock through an extensive battery of empirical simulations. In this section we describe the modeling choices we made when designing our experiments. To simulate the authentication ecoystem we need to simulate the behavior of honest users, the authentication server running DALock and an online attacker.

Briefly, when simulating users we need to model the distribution over user passwords, the distribution over honest login mistakes (e.g., typos or recall errors) as well as the user's login schedule. When simulating the distribution over user passwords we use three empirical datasets (RockYou, Yahoo!, and LinkedIn) to define the underlying password distribution. We use a Poisson arrival process to model the frequency of user login attempts [1]. Our model for user mistakes is informed by recent empirical studies of password typos [11, 12] and is augmented to simulate other mistakes i.e., recall errors. The key question to answer for simulating an authentication server running DALock is how the (password) frequency oracle $p(\cdot)$ is implemented. We consider two concrete implementations: password strength models (e.g., ZXCVBN, Markov Models, Neural Networks) and (differentially private) count sketch. When simulating the attacker we consider an untargetted one who knows the distribution over user passwords as well as the DALock mechanism — including the frequency oracle $p(\cdot)$ . We leave the question of tuning DALock to protect against targetted online attackers [53] as an important direction for future research. We elaborate on each of these key model components below. We begin by with an overview of the empirical datasets $\mathcal{D}_{\mathcal{U}}$ that we used in our experiments.

### 5.1 Experimental Datasets

In this work we use three publicly available password datasets: LinkedIn[31], RockYou[39], and Yahoo[2, 5]. In addition, we construct datasets via subsampling and password banning mentioned in **section** 5.3. We summarize the characteristics of datasets in Table 2 to help reader catch the essence.

**RockYou**[39] is a plaintext password corpus contains 14,341,564 unique passwords from a data leakage in 2009. RockYou stored user data in unencrypted database and leaked the passwords due to SQL-injection attacks.

**LinkedIn**[31] is a plaintext password corpus (partially) recovered constructed from a leakage in 2012. While the passwords were originally hashed LinkedIn was using a weak (unsalted) password hashing algorithm and almost all of the passwords in the dataset have been cracked. The dataset we used has approximately 68 million passwords. We remark that the actual size of the leak is larger and that there is a larger (differentially private) frequency corpus based on 174+ million LinkedIn passwords [24] that is publicly available. However, this dataset does not include any plaintext passwords. We chose to use the smaller dataset in our experiments so that we could evaluate with frequency oracles based on password models (e.g., ZXCVBN, PCFGs, Neural Networks).

**Yahoo!** The Yahoo! frequency corpus is a sanitized password frequency dataset collected [5] with permission from Yahoo! It consists of anonymized password histograms representing almost 70 million Yahoo! users who logged into their account during a 48 hour window in May, 2011. Yahoo! later authorized the public release of a differentially private version of this dataset [2]. We remark that this frequency corpus does not contain any plaintext passwords so we did not use password strength models in our experiments involving the Yahoo! dataset.

| Dataset | Unique Passwords | Accounts | $P(pw_1)$ |
|---------|------------------|----------|-----------|
| LinkedIn | 6,840,885 | 68,361,064 | 1.53% |
| Yahoo | 33,895,873 | 69,301,337 | 1.1% |
| RockYou | 14,341,564 | 32,603,388 | 0.89% |

**Table 2: Summary of dataset**

**Remark:** In this work, we assume each dataset (approximately) reflects the actual distribution due to their tremendous size. To ease presentation, we refer "actual distribution" to the datasets described in this section. We do acknowledge there could be misalignment between the above ones and their undelying true distributions. But we focus on analyzing popular passwords where empirical distribution and the real distribution will be very similar.

**Ethics** Some of the datasets we used (LinkedIn[31] and RockYou[39]) contain passwords that were previously stolen and subsequently leaked online. The use of this data raises important ethical considerations. We remark that the password lists are already publicly available online so our use of the data does not exacerbate the prior harm to users. We did not crack any new user passwords. Furthermore, the data we use has been cleaned of all identifying information beyond the passwords themselves. In summary, we believe that our use of the leaked data will not exacerbate prior harm to users, and that the lockout mechanism we develop and evaluate may help to protect user passwords in the future.

### 5.2 Modeling Users

Our model to simulate the behavior of honest users consists of three key components: user password selection, login frequency, and mistake model.

*5.2.1 Simulating Users' Choice of Password.* In each simulation we fix a dataset which is used to simulate user password selection. We use three large empirical password datasets: RockYou, LinkedIn and Yahoo! In particular, a dataset consists of a multiset $\mathcal{D}_{\mathcal{U}} = pw_1, ..., pw_N$ of $N$ user passwords which can be compressed into pairs $(pw, \mathsf{F}(pw, \mathcal{D}_{\mathcal{U}}))$ — recall that $\mathsf{F}(pw, \mathcal{D}_{\mathcal{U}})$ denotes the number of time the password $pw$ occurs in the dataset $\mathcal{D}_{\mathcal{U}}$. See **Section** 5.2.1 for a more details about each dataset. Each of these datasets $\mathcal{D}_{\mathcal{U}}$ induces an empirical distribution over user passwords where the probability of sampling each password $pw$ is simply $\mathsf{F}(pw, \mathcal{D}_{\mathcal{U}})/N^4$. Each simulated user $u$ in our experiment samples 6 passwords from this empirical distribution and registers with the first password. Intuitively, the five extra sampled passwords will

---

[4]In our analysis we will assume that the empirical distribution is the *real distribution* over user passwords and which is also known to the attacker. Given a password datÆŠaset $\mathcal{D}_{\mathcal{U}}$ sampled from the real distribution over user passwords we remark that when analyzing online attacks we focus on popular passwords where empirical distribution and the real distribution will be very similar.

be used to help simulate recall errors e.g., they represent the user's passwords for other websites.

We remark that the Yahoo! dataset [2, 5] only contains frequencies without actual passwords i.e., instead of recording the pair $(pw, \mathsf{F}(pw, \mathcal{D_U}))$ the dataset simply records $\mathsf{F}(pw, \mathcal{D_U})$. We generate a complete password dataset by designating a unique string for each password. As we avoid using password models like ZXCVBN to analyze DALock with the Yahoo! dataset since frequency estimation requires accesss to the original passwords. However, we are still able to evaluate DALock with the Yahoo! dataset using the Count-Sketch frequency oracle.

**Banlists** We additionally consider the setting where the authentication server chooses to ban users from selecting the top $B$ passwords e.g., $B = 10^4$ passwords. We use the normalized probabilities model [4] to simulate user password selection under such restrictions. In this model we simply use rejection sampling to avoid sampling one of the top $B$ passwords. Equivalently, we can let $\mathcal{D}_{\mathcal{U}, B}$ denote the dataset $\mathcal{D_U}$ with the $B$ most common passwords removed and sample from the empirical distribution corresponding to the updated dataset $\mathcal{D}_{\mathcal{U}, B}$.

*5.2.2 Simulating user's login patterns.* To simulate each user we need to model the frequency with which our honest user attempts to login to the authentication server. In particular, we aim to simulate the login behavior over a 180 day time span. For each user $u$ we want to generate a sequence $0 < t_1^u < t_2^u < \cdots < 4320 = 180 \times 24$ where each $t_i^u \in \mathbb{N}$ represents the time (hour) of the $i$th user visit. Following prior work (e.g., see [1, 28]) we use a Poisson arrival process to generate these login times. The Poissuon arrival process is parameterized by an arrival rate $T_u$ (hours) which encodes the expected time between consecutive login attempts $T_u = \mathbb{E}[t_{i+1} - t_i]$, and the arrival process is memoryless so the actual gap $t_{i+1} - t_i$ is independent of $t_i$. Since some users are more active than others we pick a different arrival rate $T_u$ for each user $u$ where each $T_u$ is sampled uniformly random from $\{12, 24, 24 \times 3, 24 \times 7, 24 \times 14, 24 \times 30\}$. The parameter $T_u = 12$ (hours) corresponds to users who visit multiple times per day on average, while the parameter $T_u = 24 \times 30$ corresponds to a user who visits the site once per month. For each user $u$ we use the Poisson arrival process with parameter $T_u$ to generate the sequence of visits $0 < t_1^u < t_2^u < \cdots < 4320 = 180 \times 24$ over a time span of 180 days (4320 hours). Each time a user visits we assume that they will continue attempting to login until they succeed or get locked out.

We remark that we do not simulate a client device which automatically attempts to login on the user's behalf. It may be desirable to have the authentication server store the (salted) hash of the user's previous password(s) so that we can avoid locking the user's account in settings where a client device might repeatedly attempt to login with an outdated password. Alternatively, the authentication server could store an encrypted cache of incorrect login attempts using public key cryptography where each incorrect login attempt $pw'_u \neq pw_u$ would be encrypted with a public key $pk_u$ and stored on the authentication server. The encrypted cache could only be decrypted when the user authenticates with the correct password[5]. The encrypted cache could be used as part of a personalized typo

corrector [12] and could also be used to avoid penalizing repeat mistakes [12, 56]. One potential downside to this approach is that the cache might inadvertently contain credentials from other user accounts making cached data valuable to the attacker. More empirical study would be needed to determine the risks and benefits of maintaining such a cache.

*5.2.3 Simulating User Mistakes.* The last component of our user model is a mechanism to simulate user mistakes during the authentication process. Our model relies upon recent empirical studies of password typos [11, 12] and additionally incorporates other common user mistakes e.g., recall errors. The aforementioned studies show that roughly 7.5% of login attempts are mistakes and at least 68% of these mistakes are (most likely) typos i.e., within edit distance 2 of the original passwords. Accordingly, we set the mistake rate to be 7.5% for simulation. When simulating each login attempt the user will enter the correct password with probability 92.5%. Otherwise, if the user makes a mistake we simulate a typo with probability 68% and we simulate a recall error with probably 32%. To simulate a recall error we randomly select one of the user's five alternate passwords to model a user who forgot which of his passwords was associated with this particular account — if the user recalls the wrong password they might additionally miss-type it (with probability $0.075 \cdot 0.68$). We refer an interested reader to the appendix for a more detailed discussion of our mistake model including a flow chart (see Figure 9) and more finegrained typo statistics (see Table 3).

We remark that we do not attempt to simulate a user who completely forgets his password. Of course we expect that this will occasionally happen in reality. However, we observe that a user who forgets his password will *always* need to reset it regardless of the throttling mechanism adopted by the authentication server.

## 5.3 Modeling the Authentication Server

We model an authentication server running DALock with various parameters $K$ and $\Psi$ for the strike count and hit count. Each time a user $u$ (or attacker pretending to be $u$) attempts to login the authentication server updates the parameters $\Psi_u$ and $K_u$ accordingly following the DALock mechanism. We remark that when $\Psi = \infty$ that the authentication server is running the classical $K$-strikes lockout policy. To deploy DALock with a finite hitcount parameter $\Psi$ an authentication server needs to use a frequency oracle to update the hit count after each incorrect login attempt. In this work we consider two concrete approaches the authentication server might adopt: (differentially private) Count Sketch estimator and Password Strength Models. We use $\mathsf{p}(pw, \mathsf{Estimator})$ to denote the estimated popularity (probability) of a password $pw$ using the estimator Estimator e.g., given a Count-Sketch $\sigma$ we would use $\mathsf{p}(pw, \sigma) = \frac{\mathsf{Estimate}(pw, \sigma)}{\mathsf{TotalFreq}(\sigma)}$. We remark that the authentication server might (optionally) chose to ban overly popular passwords to flatten the password distribution to protect user accounts against online attackers [41]. If the authentication server adopts such a policy then the frequency oracle would need to be adjusted accordingly to model the new password distribution.

---

[5]Unlike the public encryption key $pk_u$, which would be stored on the authentication server, the secret key $sk_u$ would only be stored in encrypted form i.e., the server would store $c_u = \mathsf{Enc}_{K_u}(sk_u)$ where $K_u = \mathsf{KDF}(pw_u)$ is a symmetric encryption key derived from the user's password.

*5.3.1 Differentially Private Count Sketch Estimator.* The first instantiation of $p(\cdot, \cdot)$ we consider is to build a Count Sketch Estimator $\sigma_{\mathcal{D}_{\mathcal{U}}} = \text{Add}(\mathcal{D}_{\mathcal{U}}, \sigma)$ from our dataset $\mathcal{D}_{\mathcal{U}}$ of user passwords. To build a Count Sketch in practice the authentication server would update the Count Sketch with the new password each time a user registers [6]. There are several issues to consider when deploying the Count Sketch estimator: memory efficiency, privacy, sample size and accuracy.

**Memory Efficiency** We instantiate the Count Sketch with parameters $d = 5$ and $w = 10^6$ so that the entire datastructure requires just 20 MB of space which easily fits in RAM.

**Privacy** As we discussed earlier one concern about storing a Count Sketch $\sigma_{\mathcal{D}_{\mathcal{U}}}$ on the authentication server is that an offline attacker might steal this file and use the data-structure to help identify user passwords. For example, if our user John Smith selects (resp. does not select) the password "J.S.UsesStr0ngpwd!" then we would expect that the true frequency of this password is $F(pw, \mathcal{D}_{\mathcal{U}}) = 1$ (resp. $F(pw, \mathcal{D}_{\mathcal{U}}) = 0$). If the Count Sketch estimator is overly accurate then the attacker would be able to learn that one user (most likely John Smith) picked this password. Without a way to address these privacy concerns an organization might be understandibly wary to deploy a Count Sketch estimator.

To address these privacy concerns we consider an $\epsilon$- differentially private estimator $\sigma_{dp} = \text{DP}(\epsilon, \sigma)$ in our experiments. During initialization we add laplace noise to each of the cells in the Count Sketch where the noise parameter scales with $d/\epsilon$. In our above example, differential privacy ensures that — up to a multiplicative advantage $e^{\epsilon}$ — an attacker cannot use the count sketch to distinguish between a dataset in which John Smith did (resp. did not) pick the password "J.S.UsesStr0ngpwd!' We remark that lower values of $\epsilon$ correspond to stronger privacy guarantees e.g., we use $\epsilon = \infty$ to denote the case with no differential privacy guarantees. In most of our experiments we use a small privacy parameter $\epsilon = 0.1$ which is much smaller than the privacy parameters used in most prior deployments of differential privacy e.g., [2, 16, 43].

**Sample Size and Accuracy** In general the accuracy of a Count Sketch increases with the size of the password dataset. Suppose that the organization does not have millions of users or the that the sample size is decreased because the organization allows users to "opt-in" to the (differentially private) count sketch. One natural question is whether a smaller organization would be able to deploy a Count Sketch to obtain reliable frequency estimates. We investigate this question by subsampling smaller datasets to train the Count Sketch. Given a set $\mathcal{U}$ of $N$ users we use $\mathcal{U}_{r\%}$ to denote a randomly subsampled set of $r\%$ of users. We use $\mathcal{D}_{\mathcal{U}_{r\%}}$ to denote the corresponding subsampled password dataset $\sigma_{r\%} = \text{Add}(\mathcal{D}_{\mathcal{U}}, \sigma)$ to denote the Count Sketch trained on the subsampled data. The question is whether $\sigma_{r\%}$ can be as effective as $\sigma$ for deploying DALock.

In our experiments we consider the following sampling rates: 1%, 5%, and 10%. We find that even when $r = 1\%$ the Count Sketch CS trained on $\mathcal{D}_{\mathcal{U}_{1\%}}$ is sufficiently accurate — even if we additionally add laplace noise to preserve $\epsilon = 0.1$-differential privacy.

**Count Sketch with Banlists** In our simulations we also consider an authentication server that bans the most popular $B = 10^4$ passwords in a dataset to help flatten the password distribution and protect users against online attacks. Theoretical anaylsis indicates that directly banning the most popular passwords is the most effective way to increase the minimum entropy of the password distribution [4]. We remark that one additional benefit of using a Count Sketch datastructure is that it can be used to help implement this type of policy i.e., if a user attempts to register with password $pw$ and $p(pw, \sigma)$ is already too high then the user will be required to pick a different password [41].

We evaluate the performance of DALock in the presence of banlists. Recall that we let $\mathcal{D}_{\mathcal{U}, B}$ denote the dataset $\mathcal{D}_{\mathcal{U}}$ with the $B$ most common passwords removed following the normalized probabilities model of [4] to model how affected users will update their passwords in response to the banlist. In particular, we assume users who are affected by the policy will pick a new passwords following the empirical distribution induced by $\mathcal{D}_{\mathcal{U}, B}$. We then train the Count Sketch on the updated dataset i.e., $\sigma_{-B} = \text{Add}(\mathcal{D}_{\mathcal{U}, B})$ as follows.

*5.3.2 Frequency Oracle from Password Models.* As we previously discussed there are several reasons why an organization might prefer not to use a Count Sketch for frequency estimation e.g., privacy concerns or limited sample size. An alternative is to instantiate the frequency oracle with a password model. This could be a heuristic password strength meter, a more sophisticated model based on Neural Networks, Probabilistic Context Free Grammars or Markov Models or an empirical estimate based on Hashcat. The primary advantage to this approach is that the model can be deployed immediately even before an organization has any users and there are no privacy concerns.

We adopted the ZXCVBN password strength meter [54] as prior empirical studies demonstrate that it is one of the most accurate password strength meters [22]. We used the Password Guessing Service [47] to obtain guessing numbers for Neural Network, PCFG, Hashcat, and Markov Models — we also considered the minimum guessing number across all four models as suggested in [47]. For example, if a password $pw$ had guessing number $g$ we might estimate that $p(pw_i) = 1/g$. One challenge that we needed to address was that the estimates we obtain do not always yield a probability distribution e.g., for ZXCVBN we have $\sum_{i=1}^{10000} p(pw_i) \gg 1$ where $i$ ranges over the top $10^4$ passwords in the dataset. Thus, before deploying the frequency estimator in DALock we renormalized our estimates so that $\sum_{i=1}^{10000} p(pw_i) = 1$.

## 5.4 Modeling the Attacker

The final component of our simulation is a model of the attacker. We take a conservative approach and model an untargetted attacker with complete knowledge of the password distribution. Following Kerckhoff's principle we also assume that the attacker has access to the complete description of the DALock mechanism. In particular, for any password $pw$ we assume that the attacker knows both the true probability $P(pw)$ and the estimated probability $p(pw)$. Finally, we also assume that the attacker is given the complete sequence of login times $t_1^u \leq t_2^u \leq \ldots \leq 24 \times 180$ for each user $u$ over a 180

---

[6] The Count Sketch instantiations we consider would also support a Remove operation which would allow the authentication server to handle password updates efficiently

day time span as well as the outcome of each login attempt e.g., at time $t_i^u$ user $u$ will succeed after 2 incorrect guesses.

**Remark:** We conservatively aim to overestimate the capabilities of an untargetted online attacker. In practice, the online attacker will be able to able to approximate $P(pw)$ and $p(pw)$ overtime by interacting with the DALock server e.g., by setting up dummy accounts to test many times he can submit a particular incorrect guess without exceeding the hit count. Similarly, the attacker would not necessarily know the exact login times for a user, but this conservative assumption makes it feasible to precisely characterize the optimal behavior of an attacker. In practice, an online attacker might wait several days in between guesses to avoid accidently locking the user's account based on the number of consecutive incorrect login attempts.

*5.4.1 Optimizing Attack Strategies.* The goal of the attacker is to maximize the probability of cracking each password within the fixed 180 day time span. For example, the attacker might try to find popular passwords $pw$ where the ratio $p(pw)/P(pw)$ is small so that the increased hit count is smaller than intended. We formalize the attacker's optimal strategy in terms of the Password Knapsack problem (PK). Unsurprisingly, the password knapsack probelm turns out to be NP-hard (as we prove in the appendix), but there are several heuristic algorithms the $\mathcal{A}$ can use which yield nearly optimal strategies in practice.

Recall that we assume that the $\mathcal{A}$ has perfect knowledge of the distribution and probability estimates for each password $pw$. We also assume $\mathcal{A}$ knows the DALock security parameters $K$ and $\Psi$. Furthermore, for each user $u$ we assume that the attacker is given the complete sequence of login times $t_1^u \leq t_2^u \leq \ldots \leq 24 \times 180$ for each user $u$ over a 180 day time span as well as the outcome of each login attempt e.g., at time $t_i^u$ user $u$ will succeed after 2 incorrect guesses. In particular, at any point in time $t < 24 \times 180$ the attacker can infer the current strike threshold and hit count threshold for any user $u$. We denote by $K_{u,t}$ (resp. $\Psi_{u,t}$) the strike (resp. hit count) threshold for user $u$ at time $t$ assuming that the attacker does not submit any of his own guesses.

Supposing that the attacker wishes to avoid locking down the user's account before time $t$ the cummulative (estimated) probability of all guesses submitted before that time should be at most $\Psi'_{u,t} := \Psi - \Psi_{u,t}$. Similarly, we let $M(t)$ denote the maximum number of gueses that the attacker can sneak in over the first $t$ hours without locking down the account i.e., because $K_{u,t'} \geq K$ at some time $t' \leq t$.

Fixing the time parameter $t$ the attacker's goal is to find a subset $S_t \subseteq \mathcal{P}$ of $M(t)$ passwords to check such that

$$\sum_{pw \in S_t} p(pw) \leq \Psi'_{u,t} . \tag{1}$$

After checking the passwords in $S_t$ the attacker can still check one more password $pw_{hold} \notin S_t$ before the account is locked down. Given a set $S_t$ and a holdout password $pw_{hold} \notin S_t$ the probability that the attacker succeeds is

$$P(pw_{hold}) + \sum_{pw \in S_t} P(pw) . \tag{2}$$

Thus, the goal of the attacker is to find a subset $S_t$ of size $|S_t| \leq M(t)$ maximizing his success rate (eq 2) subject to the constaint in equation 1.

**Password Knapsack Problem** Given a password dictionary $\{pw_1, \ldots, pw_n\}$ we formally define the Password Knapsack(PK) problem as the following integer program with indicator variables $s_i \in \{0, 1\}$ and $l_i = \{0, 1\}$ for each password $pw_i$. The attackers goal is to select a holdout password and a separate subset of $M$ ($= M(t)$) passwords with total 'weight' (estimated probability) at most $\Psi'$ ($= \Psi'_{u,t}$)

$$\max \sum_i (s_i + l_i) \cdot P(pw_i)$$

*subject to,*

$$\sum_i s_i \cdot p(pw_i, \sigma)) \leq \Psi'$$
$$\sum_i s_i \leq M$$
$$\sum_i l_i \leq 1$$
$$\forall i \; l_i + s_i \leq 1$$

*where,*

$$\forall i, s_i, l_i \in \{0, 1\}$$

Intuitively, setting $s_i = 1$ means $pw_i$ is selected to be placed in the "password knapsack" $S \subseteq \mathcal{P}$, i.e. to be used for dictionary attack. Setting $l_i = 1$ indicates that password $pw_i$ is used as holdout password. This is equivalent to the following optimization problem. The constraints ensure that $|S| \leq M$ and we pick exactly one holdout password that is not already in $S$.

**Solving the Password Knapsack** To maximize the number of cracked passwords an online attacker can compute $M(t)$ and $\Psi'_{u,t} := \Psi - \Psi_{u,t}$ for each time $t \leq 24 \times 180$ and solve the corresponding Password Knapsack problem. Given optimal solutions $(pw^*_{hold,t}, S^*_t)$ for each time $t$ the attacker will pick the solution that maximizes the number of cracked passwords as in equation 2. We remark that the calculations above need to be repeated for each different user $u$ since the values $M(t)$ and $\Psi'_{u,t}$ may vary due to different visitation schedules.

**Solving Password Knapsack** Unfortunately, the Password Knapsack problem is NP-hard as we prove in **Theorem** A.1 in the Appendix via a straightforward reduction from Subset Sum. In all of instances we considered we found that the optimal choice for the holdout password was simply $pw_1$ the most likely password in the distribution. Once we fix our holdout password our problem reduces to the two dimensional knapsack problem. We remark that PK can be viewed as a two dimensional knapsack problem.

Assuming $P \neq NP$ the two dimensional knapsack probelm does not even admit a polynomial time approximation scheme (PTAS) [30] in contrast to the regular knapsack problem which has fully polynomial time approximation scheme (FPTAS)). Thus, we consider two heuristic approaches to solve the password knapsack problem: Dantizig's Algorithm Based[14] approach (DAB) and Feasible Most Promising Password First approach(FMPPF).

DAB (**Algorithm** 2, Appendix) sorts the remaining passwords $\mathcal{P}_{\tilde{\Pi}} = \{pw_2, \ldots pw_n\}$ based on the ratios $\frac{P(pw_i)}{p(pw_i)}$ and select candidates based on the sorted order until we either select $M$ passwords or until selecting another password would exceed our capacity $\Psi'$. FMPPF (**Algorithm** 3, Appendix) sorts the remaining passwords

based on the true probability $P(pw_i)$ and simply selects password $pw$ in sorted order until we either select $M$ passwords or until selecting another password would exceed our capacity $\Psi'$. We discuss the advantages and disadvantages to both heuristics in the appendix. Intuitively, FMPPF (resp. DAB) will perform better when $M$ (resp. $\Psi'$) is the limiting constraint.

We found that FMPPF generally performs better than DAB despite of its simplicity. In addition, our simuation shows that FMPPF's performance is close to optimal. Practically speaking, one generally expect $p(pw_i) \approx P(pw_i)$ especially when $pw_i$ is a popular password. In such case, DAB can hardly gain advantages from underestimation. Further more, imagine one bucket passwords by probability ranges, there are plenty of passwords in each bucket. Intuitively, picking passwords ordered by $P(pw_i)$ should produce an (almost) optimal solution (quickly). Thus, we choose to present the results based on FMPPF approach.

## 6 EXPERIMENTAL RESULTS

We empirically evaluated the performance of DALock under a variety of scenarios. During each simulation we had $10^6$ honest users register with an authentication server running DALock and login over a period of 180 days. To analyze usability we ran the simulations without an online password attacker and measured the unwanted lockout rate i.e., the fraction of user accounts that were locked due to honest mistakes. To analyze security we added an untargetted online attacker to the simulation and measured the fraction of user passwords that the attacker cracked.

The results of our simulations are summarized in **Figures** 3, 4, 5, 6, 7, and 8[7]. The first four figures evaluate the security (**Figures** 3 and 5) and usability (**Figures** 4 and 6) in the absence of a banlist. The last two figures evaluate the usability (**Figure** 8 ) and security (**Figure** 7 ) of DALock when the authentication server bans the top $B = 10^4$ passwords in our dataset. The X-axis of each plot represents the time span over 180 days. And the Y-axis represents percentage of compromised users (unwanted locked out rate) for security (usability) experiments.

**Implementation Details** In each of our implementations DALock we instantiated $K = 10$ using hit count parameters $\Psi \in \{2^{-7.0}, 2^{-9.375}\}$ (no banlist) and $\Psi = 2^{-11}$ (with banlist). In each batch of experiments we used one of our three password datsets $\mathcal{D}_{\mathcal{U}}$ (RockYou, Yahoo!, or LinkedIn) to define our password distribution and we instantiated the frequency oracle using (1) a variety of password models including ZXCVBN [54], Hashcat, Markov Models, PCFGs and Neural Networks [47][8] and (2) an $\epsilon$-differentially private count sketch trained on a subsample $\mathcal{D}_{\mathcal{U}_{r\%}}$ of containing $r\%$ of the original dataset $\mathcal{D}_{\mathcal{U}}$ with $r \in \{1\%, 5\%, 10\%, \texttt{All}\}$. When instantiated

with banlist we used the dataset $\mathcal{D}_{\mathcal{U},B}$ instead of $\mathcal{D}_{\mathcal{U}}$ i.e., we removing the $B = 10^4$ most common passwords.

**Baseline** We used the classical 3-strike mechanism and the 10-strike mechanisms (recommend by Brostoff et. al [7] to improve usability) as a baseline for comparison. We remark that this is equivalent to $(3, \Psi = \infty)$-DALock and $(10, \Psi = \infty)$-DALock respectively. Our results suggest that one can improve *both* security and usability by replacing the classic 3-strike throttling mechanism with $(10, \Psi) - $DALock. Our results demonstrate that $(10, \Psi)$-DALock greatly outperforms the classic 10-strikes throttling mechanism without significant usability loss – from a usability standpoint decreasing $\Psi$ can only increase the unwanted lockout rate. We discuss our findings in more detail below.

### 6.1 Usability

**Figures** 4 and 6 highlight the usability of DALock in the absence of a banlist with hit count parameters $\Psi = 2^{-9.375}$ and $\Psi = 2^{-7.0}$. When $\Psi = 2^{-7.0}$ we find that DALock *always* outperforms the classical 3-strikes mechanism regardless of how the frequency oracle is instantiated. When $\Psi = 2^{-9.375}$ DALock still outperforms the classical 3-strikes mechanism when instantiated with a countsketch — even when we train on just $r = 1\%$ of the dataset and add laplace noise to achieve $\epsilon = 0.1$-differential privacy. When we instantiate DALock password models the results were mixed e.g., ZXCVBN and Hashcat had superior usability while Markov Models and Probabilistic Context Free Grammars performed poorly.

**Figure** 8 highlights the usability benefit of banning the top $10^4$ passwords. While a few users might be inconvenienced during the password registration our simulations indicate that the unwanted lockout rates for DALock are greatly reduced even when we adopt a stricter $\Psi = 2^{-11}$. Intuitively, the banlist allows us to avoid locking out users who select and overly popular password as one of their five alternate passwords for other accounts. We remark that when we instantiate DALock with a count sketch that the usability results are virtually identical to the 10-strikes policy — even if we use a $\epsilon = 0.1$ -differentially private count sketch trained on just 1% of the data.

**10-strike Mechanism is user friendly** Brostoff et. al [7] proposed that one should replace 3-strike mechanism with 10-strike mechanism to achieve higher usability. Our simulation results clearly align with their recommendations. Based on our plots, 10-strike mechanism results in unwanted locked out rate close to zero. However, deploying this mechanism also has a high security cost as indicated by our simulations.

### 6.2 Security Results

**Figures** 3 and 5 evaluate the security of DALock in the absence of a banlist with hit count parameters $\Psi = 2^{-9.375}$ and $\Psi = 2^{-7}$ respectively. For reference we also plotted the line $\Psi + P(() pw_1)$ which constitutes a theoretical upper bound on the attacker success rate when DALock is instantiated with a perfect frequency oracle. We found that DALock always outperforms even the stricter $K = 3$-strikes policy under *all* insantiations of the frequency oracle (excluding Hashcat with $\Psi = 2^{-7}$). For example, the attacker

[7]We included additional experimental results in **Appendix** C for interested readers.
[8]We relied on the Password Guessing Service to obtain cracking numbers for Hashcat, Markov Models, PCFGs and Neural Networks [47]. We found that the Neural Network model failed to crack many weak passwords in our datasets as it was configured not to guess short passwords. As such we did not directly use Neural Networks to implement our frequency oracle. However, the Neural Network guessing numbers are included in our Min-all frequency oracle which uses the minimum guessing number over all models.
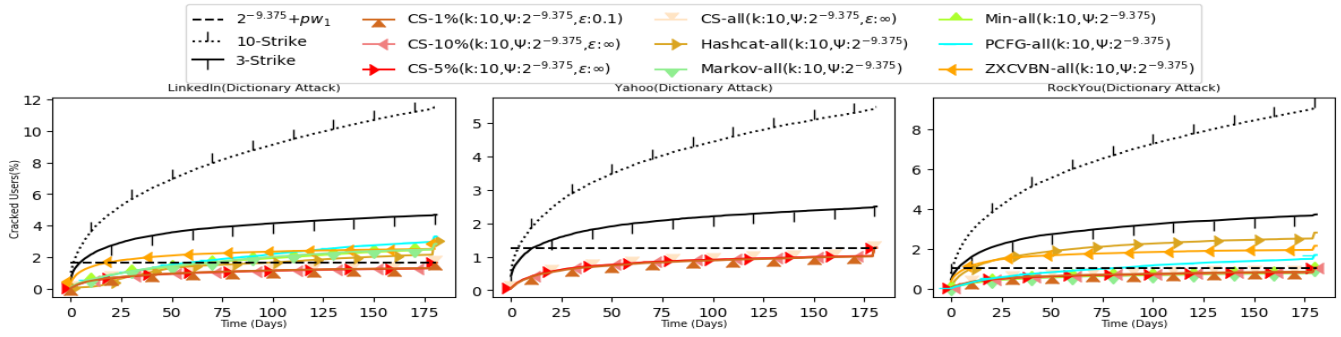
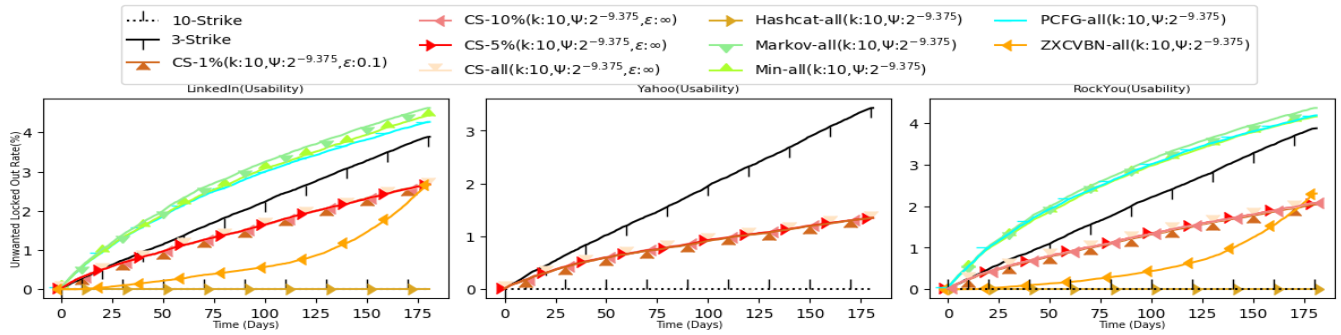**Figure 3: Security Measurement of** DALock - $\Psi = 2^{-9.375}$



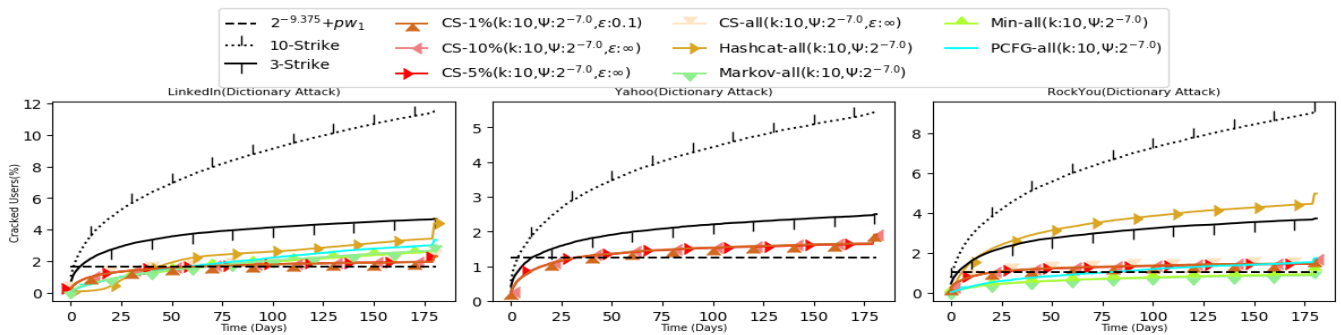**Figure 4: Usability Measurement of** DALock - $\Psi = 2^{-9.375}$



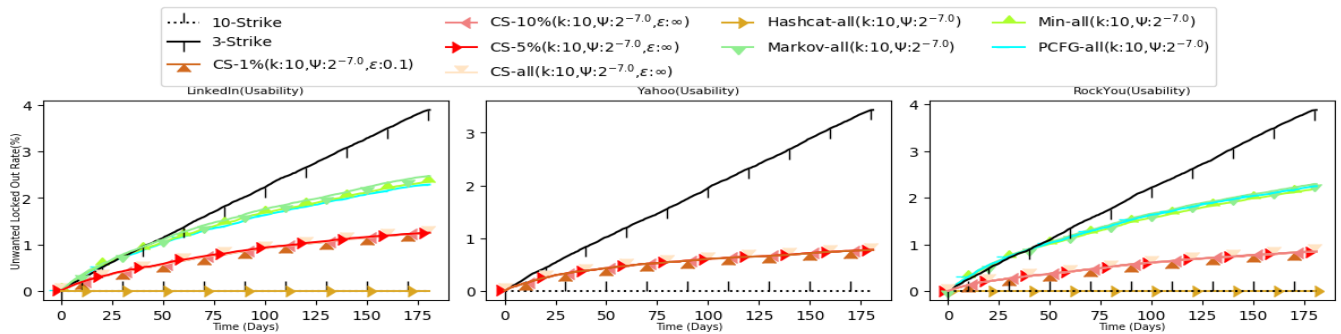**Figure 5: Security Measurement of** DALock - $\Psi = 2^{-7.0}$



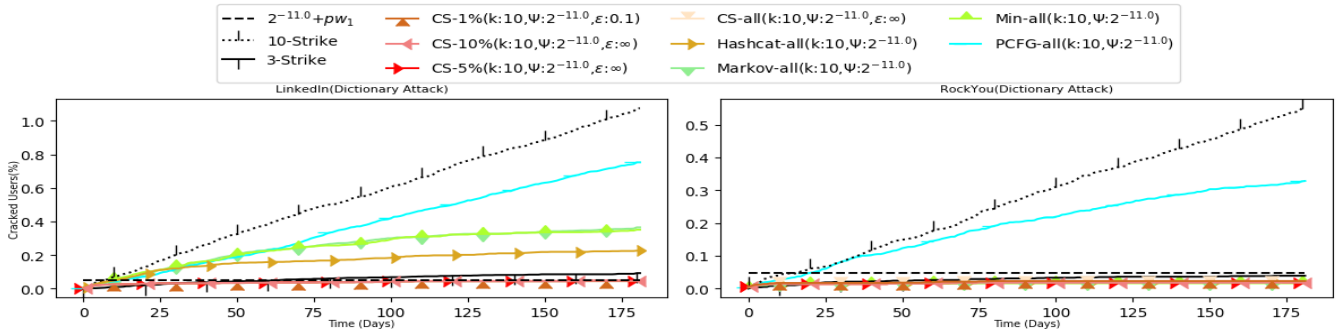**Figure 6: Usability Measurement of** DALock - $\Psi = 2^{-7.0}$

**Figure 7: Security Measurement of** DALock - $\Psi = 2^{-11}$ **(Banning top $B = 10^4$ passwords)**
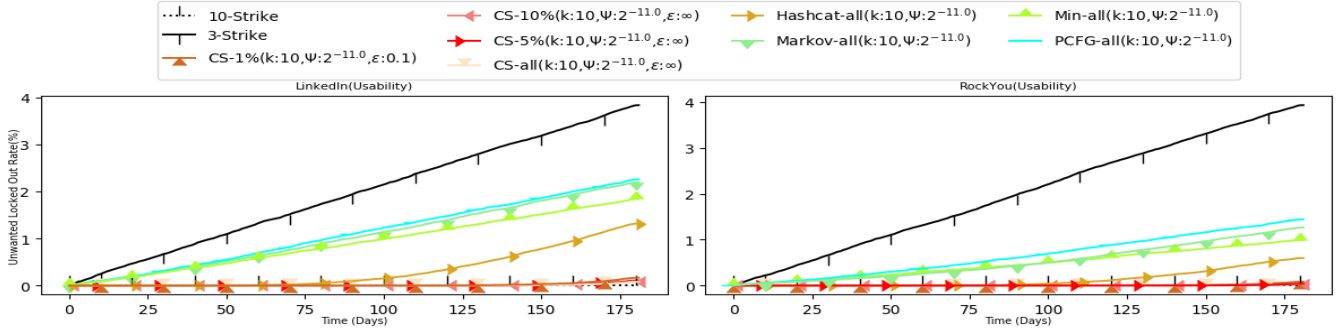


**Figure 8: Usability Measurement of** DALock - $\Psi = 2^{-11}$ **(Banning top $B = 10^4$ passwords)**

cracked roughly 4% of users accounts when facing 3-strike mechanism (LinkedIn + RockYou) compared with 2% when deploying DALock with a differentially private count sketch.

Figure 7 highlights the advantage of banning the most popular $B = 10^4$ passwords. We remark that DALock always outperformed the $K = 10$-strikes policy. We also found that the $K = 3$-strikes policy and DALock (with a differentially private Count Sketch) were both *highly effective* at protecting user accounts with a compromised rate is about close to zero. ( $\approx 0.092\%$ for 3-strike and $\approx 0.048\%$ for differential private Count Sketch).

## 6.3 Summary and Discussion

We find that $(10, \Psi)$-DALock offers a superior security/usability tradeoff to the classical $K$-strikes mechanism. Our experiments also highlight the security *and* usability benefits of banning overly popular passwords. We found that DALock can be reasonably instantiated with password strength models such as ZXCVBN, Markov Models, Probabilistic Context Free Grammars and Neural Networks. However, we obtain the best security/usability tradeoffs when we ban the most popular passwords and when we instantiate the DALock frequency oracle with a differentially private count sketch. We found that the count sketch can be reliable trained from a smaller subsample containing just 1% of the dataset, and even if we add enough noise to preserve $\epsilon - 0.1$-differential privacy (strong privacy). This is promising news for a smaller organization that is considering deploying DALock.

**Limitations** Our empirical security results are all based on simulations. While we aim to model the authentication server, users and an attacker there will inevitably be some differences between the simulated/real-world behavior of the attacker/users. We also remark that our simulations do not model the behavior of targetted

attackers. Extending DALock to protect against targeted attackers is an important research question that is beyond the scope of the current paper. Finally, we remark that larger organizations might distribute the workload across multiple authentication servers. In this case maintaining a synchronized state $(K_u, \Psi_u)$ for each user $u$ could be challenging. To address this challenge it may be necessary to define a relaxation of our DALock mechanism where the states $(K_u, \Psi_u)$ on each authentication server are not always assumed to be perfectly synchronized.

## 7 CONCLUSION

We present a novel *distribution aware* password throttling mechanism DALock that penalizes incorrect passwords proportionally to their popularity. We show that DALock can be reliably instantiated with either a password strength model such as ZXCVBN or with a differentially private count sketch. Our empirical analysis demonstrates that DALock offers a superior balance between security and usability and is particularly effective when used in combination with a short banlist of overly popular passwords. For example, we are able to reduce the success rate of an attacker to 0.05% (compared to 1% for the 10-strikes mechanism) whilst simultaneously reducing the unwanted lockout rate to just 0.08% (compared to 4% for the 3-strikes mechanism).

# A DALOCK

## A.1 DALock **Authentication Algorithm**

We supplement the pseudo code of DALock in this section to help readers understand how to implement DALock for authencation. The authencation process takes four arguments: username $u$, input password $pw$, salt $s_u$, and password popularity estimator $\sigma$. Before verifying the correctness of entered password, DALock first check if $u$'s account has already been locked or not based on $\Psi_u$ and $Ku$. If the account is not locked, DALock proceeds and verify the correctness of the passwords. If the password is valid, DALock resets strike threshold $K_u$ and grant user the access to the service. If the entered password is wrong then in addition to denying the access to the service, the server also increases $\Psi_u$ and $K$ by $p(pw)$ and 1 respectively.

---

**Algorithm 1** DALock: Novel Password Distribution Aware Throttling Mechanism

---
1: **function** LOGIN($u$, $pw_u$, $\sigma$, $s_u$)
2:    **if** $\Psi_u \geq \Psi$ or $K_u \geq K$ **then**
3:      Reject Login
4:    **end if**
5:    **if** $hash(pw, s_u) == hash(pw_u, s_u)$ **then**
6:      Reset $K_u$ to $K$
7:      Grant Access
8:    **else**
9:      $\Psi_u \leftarrow \Psi_u + p(pw, \sigma)$
10:     $K_u \leftarrow K_u + 1$
11:     Deny Access
12:    **end if**
13: **end function**

---

## A.2 **Password Knapsack is** NP **hard**

THEOREM A.1 (HARDNESS OF PASSWORD KNAPSACK). *Find optimal solution for password knapsack is* NP-*hard.*

**Proof:** We first formally define subset sum problem, and then prove password knapsack is NP hard by showing the reduction from subset sum to it.

*Definition A.2 (Subset Sum).* Given Partition instance $x_1, \ldots, x_n \in (0, 2^m]$ and target sum value $T$. The goal is to find $S \subseteq [n]$ s.t. $\sum_{i \in S} x_i = T$?

**Reduction**: One can create the following password knapsack instance

- Set $\gamma = \sum_{i=1}^n x_i$,
- Set $\psi = T/(2\gamma) < \frac{1}{2}$,
- Set $CS(p_i) = f(p_i) = x_i/(2\gamma)$ for $i = 1, \ldots, n$
- Set $f(p_{last}) = 1 - \sum_{i=1}^n p_i = 1/2 > \psi$.

If $S$ exists for partition instance then attacker can use $S$ for password knapsack to crack $p_{last} + T/(2\gamma)$ passwords. On the other hand let $S$ be the optimal password knapsack solution such that $\sum_{i \in S} CS(p_i) \leq \psi$ then the attacker cracks at most $p_{last} + \sum_{i \in S} f(p_i) \leq 1/2 + \psi$ passwords. If equality holds then $\sum_{i \in S} f(p_i) = \psi$ which implies $\sum_{i \in S} x_i = T$ by definition of $\psi$.

## A.3 **Solving** PK **with Heruistics**

In this section we supplement the missing details of algorithm DAB and FMPPF mentioned in **Section** 5.4.

The DAB approach takes three inputs: a sorted password dictionary based on the ratio of actual popularity and estimated popularity $\frac{\mathrm{p}(pw)}{\mathrm{P}(pw)}$: $\mathcal{P}_{\tilde{\Pi}} = \{pw_{\tilde{\Pi}(1)}, \ldots, pw_{\tilde{\Pi}(n)}\}$, attack budget $\Psi$ and K. The algorithm keeps placing passwords into the knapsack S based on the sorted order until it cannot further add some password $pw$. At this points, DAB compares $\mathrm{P}(pw)$ with $\mathrm{P}(S)$ and sets S to be the one with higher values. After that, the algorithm repeat the above process by scanning throught the whole dictionary. At the end, since only $K$ passwords is allowed to be used for guessing, the algorithm returns $K$ passwords based on their actual probability

Primary incentives of using this algorithm are 1) to take advantage of underestimated passwords and 2) to avoid (severely) overestimated ones. There are several drawbacks of DAB. Firstly, the progress can be slow because priority are given to significantly underestimated passwords. Intuitively, for popular passwords ($\mathrm{P}(pw)$ large) the ratio $\frac{\mathrm{P}(pw)}{\mathrm{p}(pw)}$ is likely to be close to 1, therefore, attempts with popular ones are likely to be delayed. Secondly, unlike vanilla version of Knaspack, DAB may not yield a 2-approximation due to the additional constraint on the number of passwords one can place in the Knapsack. Third, computation cost is slightly higher for running DAB though both algorithms terminate reasonably quickly.

---

**Algorithm 2** DAB **Attack**

---
**Return:** An array of password sorted in the order of guessing
1: **function** DAB ATTACK($\mathcal{P}_{\tilde{\Pi}}$, $\Psi$, $M(T)$)
2:    $S$= []
3:    **while** $S$ changes **do**
4:      **for** $pw \in \mathcal{P}_{\tilde{\Pi}}$ **do**
5:        **if** $\mathrm{p}(pw) > \Psi$ **then**
6:          continue
7:        **end if**
8:        **if** $\mathrm{p}(S \cup pw) < \Psi$ and $|S| \leq M(T)$ **then**
9:          $S \leftarrow S \cup pw$
10:       **else if** $\mathrm{P}(pw) > \mathrm{P}(S)$ **then**
11:         $S \leftarrow \{p\}$
12:       **end if**
13:      **end for**
14:    **end while**
15:    **return** Top $M(T)$ passwords from $S$ based on actual popularity.
16: **end function**

---

An alternative to DAB is FMPPF. It takes three input parameters $\mathcal{P} = pw_1, \ldots, pw_n\}$, attack budget $\Psi$ and $M(T)$ and selects passwords greedily. FMPPF differs from DAB in the following two aspects. Firstly, FMPPF uses a password dictionary sorted based on the actual popularity only, which can be easily obtained in reaf life. Secondly, to save computational cost FMPPF terminates once it finds $K$ passwords that are suitable for attacks and stop further explore the dictionary. The pseudo can be found in **Algorithm** 3

In short time attack scenarios, FMPPF offers better chance of success than DAB by attempting popular ones first. For long term

case, FMPPF should still be able to achieve almost optimal results given an abundant choice of passwords. In fact, based on the empirical results (in **section** 6.2), the performance of FMPPF is very close to theoretical upper bounds ($\Psi + P(pw_1)$ ).

---

**Algorithm 3** FMPPF **Approach**

---

**Return:** An array of password sorted in the order of guessing

1: **function** FMPPF($\mathcal{P}, \Psi, M(T)$)
2:     S= []
3:     **for** $pw \in \mathcal{P}$ **do**
4:         **if** $p(S \cup p) < \Psi$ and $|S| < M(T)$ **then**
5:             $S \leftarrow S \cup pw$
6:         **end if**
7:     **end for**
8:     **return** S
9: **end function**

---

## B    SIMULATING USER'S MISTAKES

In this section we elaborate the details for simulating users' mistakes missing in **Section** 5.2. To help reader visualize the process of simulating mistakes, we plot the flowchart in **figure** 9. The starting point is to simulate recall errors. Based on the empirical results of existing literatures[11, 12], we set the probability of makeing an recall error to 2.4%. Recall that when we generate user's profile, each user has five "passwords from other services". Therefore, we simulate recall error by choose one of them to be the *password user intends to enter*. On top of this process, we further simulate typos (on the password intended to enter) with probability $\approx$ 5%. Condition on making typos, we simulate this step by choosing a type of typos with their conditional probability summarized in **Table** 3.
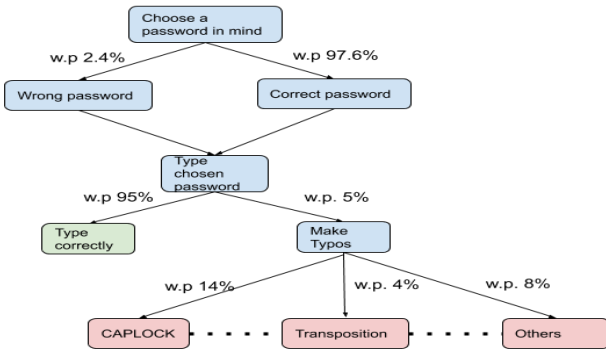


**Figure 9: Flow Chart for Simulating Users' mistake**

## C    MORE EXPERIMENTAL RESULTS

We provide more experimental results for curious readers to better understand the performance of DALock in multiple scenarios. To clearly demonstrates the impact of $\Psi$ on security and usability we choose to visualize the following set of parameters $\{2^{-6}, 2^{-7}, 2^{-8}, 2^{-9}, 2^{-10}\}$(**Figure** 10 and 11). When one set $\Psi$ to an

| Typo Types | Chance of Mistake(Rounded %) |
|---|---|
| CapLock On | 14 |
| Shift First Char | 4 |
| One Extra Insertion | 12 |
| One Extra Deletion | 12 |
| One Char Replacement | 31 |
| Transposition | 4 |
| Two Deletions | 3 |
| Two Insertions | 3 |
| Two Replacements | 10 |
| Others | 8 |

**Table 3: Typo Distributions[12]**

over conservative, e.g. $2^{-10}$ or smaller, dictionary attackers are not able to crack a significant potion of users accounst; however, usability can be a concern as infrequent passwords burns $\Psi_u$ quickly.(see curve CS-all($k : 10, \Psi : 2^{-10}, \epsilon$)) across all figures). Another extreme approach is adopting aggresively large $\Psi$, e.g. $2^{-6}$. Based on the plots(curve CS-all($k : 10, \Psi : 2^{-6}, \epsilon$)), the usability performance is satisfactory while security risk is enlarged (but still better than 3-strike). Beyond standard Count Sketch, we supplement more results on applying differential privacy to Count-Sketch(**Figure** 12 and 13). We discovered that differential privacy hardly have any negative impact on the performance of DALock especially when the dataset is large. Further more, we supplement more results on subsampling to show $\sigma_{r\%}$ is effective for a wide range of $\Psi$ even when the sampling rate is 1%.(**Figure** 14 and 15) Finally, we present the results of differentially private DALock trained on subsampling dataset $\mathcal{D}_{U_{1\%}}$. Despite the datasets become 100 times smaller, DP($\sigma_{1\%}$,) still performs reasonly well.(**Figure** 16 and 17).

**Almost Optimal Heruistic**  If one assumes that $P(pw) \approx p(pw)$ for popular passwords, then theoretically $\mathcal{A}$ can compromise at most $\Psi + P(pw_1)$ users accounts by using $pw_1$ as holdout passwords. Based on our simulation results, we found that $\mathcal{A}$ is very close to achieve such threshold by adopting FMPPF algorithm metioned in **Section** 5.4. To help readers identify this property, we highlight the upperbouds a series of upperbounds(e.g. $2^{-6} + P(pw_1)$ etc)
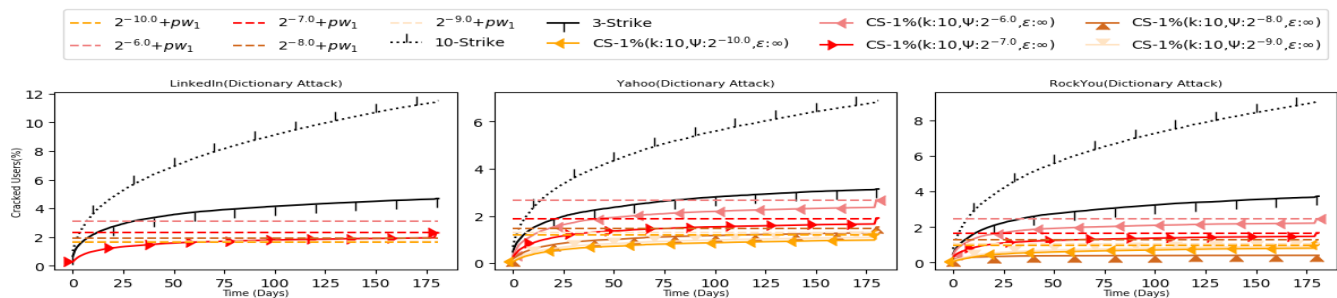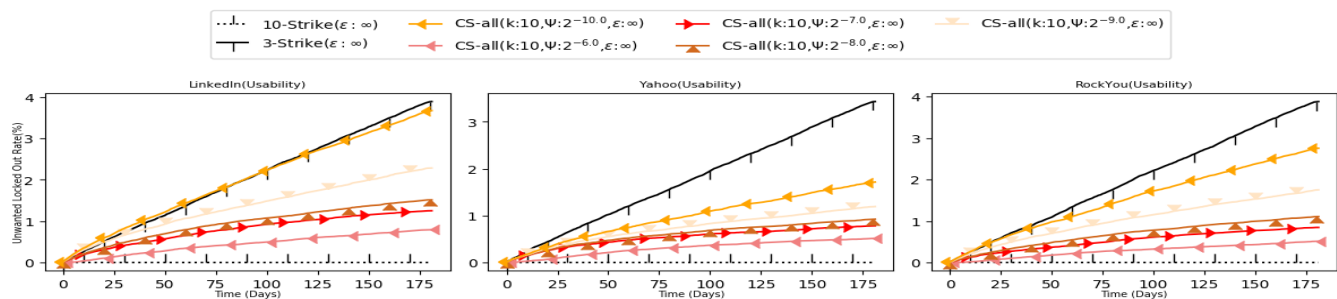
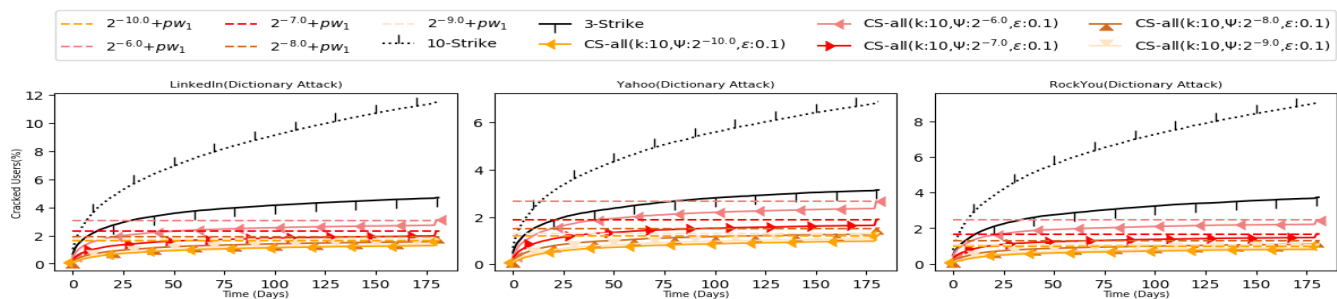**Figure 10: Security:** $\sigma$



**Figure 11: Usability:** $\sigma$



**Figure 12: Security: DP**$(0.1, \sigma)$
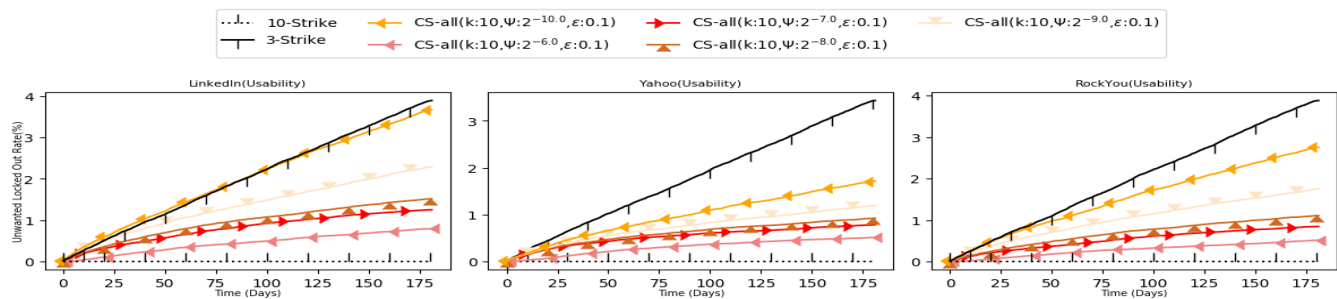


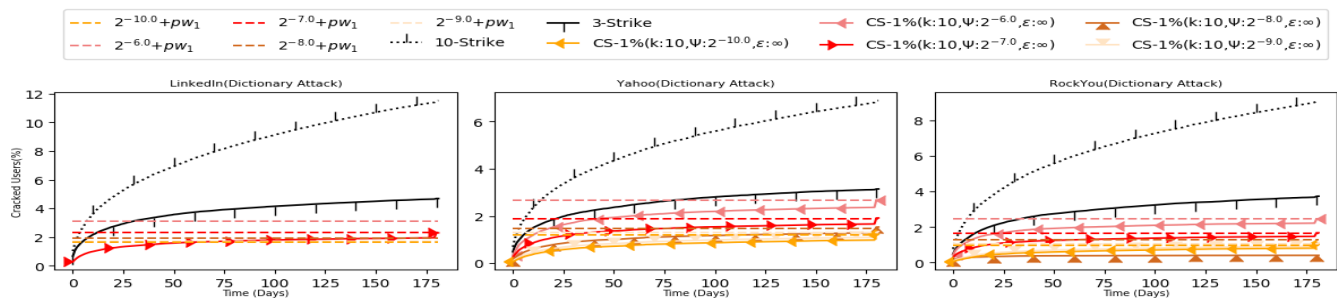**Figure 13: Usability: DP**$(0.1, \sigma)$

**Figure 14: Security:** $\sigma_{1\%}$
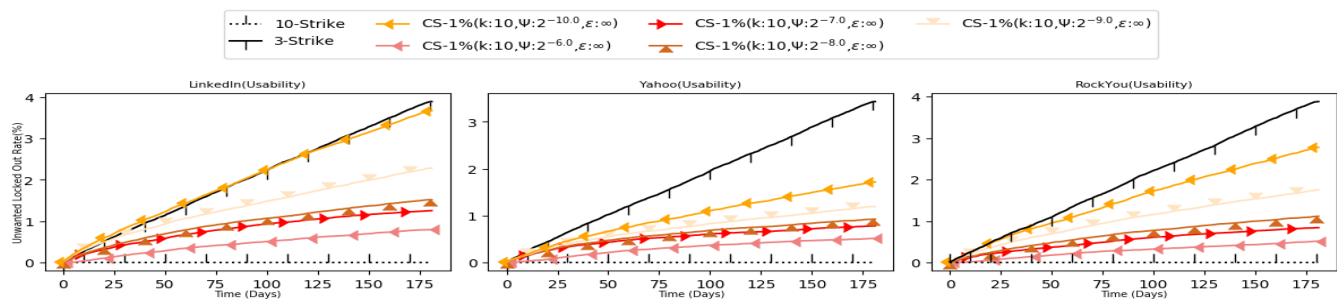


**Figure 15: Usability:** $\sigma_{1\%}$
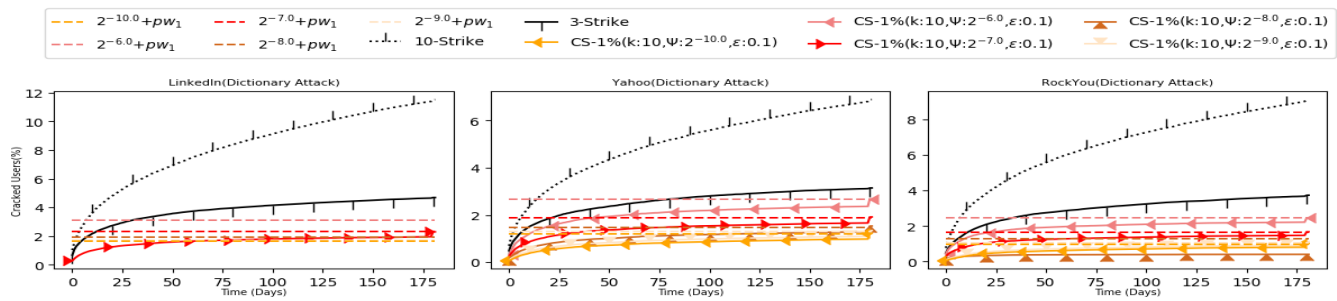


**Figure 16: Security: DP**$(0.1, \sigma_{1\%})$
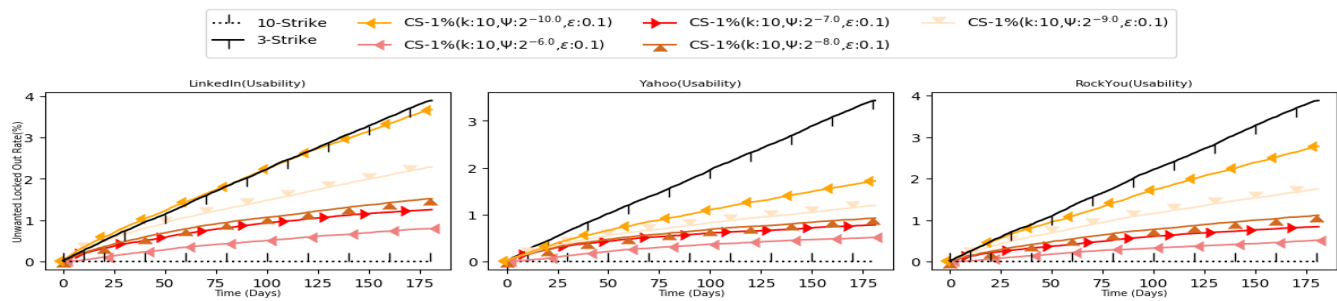


**Figure 17: Usability:DP**$(0.1, \sigma_{1\%})$

## ACKNOWLEDGMENTS

## REFERENCES

[1] Jeremiah Blocki, Manuel Blum, and Anupam Datta. 2013. Naturally Rehearsing Passwords. 361–380. https://doi.org/10.1007/978-3-642-42045-0_19
[2] Jeremiah Blocki, Anupam Datta, and Joseph Bonneau. 2016. Differentially Private Password Frequency Lists.
[3] Jeremiah Blocki, Benjamin Harsha, and Samson Zhou. 2018. On the Economics of Offline Password Cracking. 853–871. https://doi.org/10.1109/SP.2018.00009
[4] Jeremiah Blocki, Saranga Komanduri, Ariel Procaccia, and Or Sheffet. 2013. Optimizing password composition policies. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 105–122.
[5] Joseph Bonneau. 2012. The Science of Guessing: Analyzing an Anonymized Corpus of 70 Million Passwords. 538–552. https://doi.org/10.1109/SP.2012.49
[6] Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. 2012. The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes. 553–567. https://doi.org/10.1109/SP.2012.44
[7] Sacha Brostoff and Angela Sasse. 2003. Ten strikes and you're out: Increasing the number of login attempts can improve password usability. (07 2003).
[8] Elie Bursztein, Steven Bethard, Celine Fabry, John C. Mitchell, and Daniel Jurafsky. 2010. How Good Are Humans at Solving CAPTCHAs? A Large Scale Evaluation. 399–413. https://doi.org/10.1109/SP.2010.31
[9] Elie Bursztein, Matthieu Martin, and John C. Mitchell. 2011. Text-based CAPTCHA strengths and weaknesses. 125–138. https://doi.org/10.1145/2046707.2046724
[10] Moses Charikar, Kevin C. Chen, and Martin Farach-Colton. 2002. Finding Frequent Items in Data Streams. 693–703. https://doi.org/10.1007/3-540-45465-9_59
[11] Rahul Chatterjee, Anish Athayle, Devdatta Akhawe, Ari Juels, and Thomas Ristenpart. 2016. pASSWORD tYPOS and How to Correct Them Securely. 799–818. https://doi.org/10.1109/SP.2016.53
[12] Rahul Chatterjee, Joanne Woodage, Yuval Pnueli, Anusha Chowdhury, and Thomas Ristenpart. 2017. The TypTop System: Personalized Typo-Tolerant Password Checking. 329–346. https://doi.org/10.1145/3133956.3134000
[13] Graham Cormode and Shan Muthukrishnan. 2005. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms* 55, 1 (2005), 58–75.
[14] George B Dantzig. 1957. Discrete-variable extremum problems. *Operations research* 5, 2 (1957), 266–288.
[15] Cynthia Dwork. 2011. Differential privacy. *Encyclopedia of Cryptography and Security* (2011), 338–340.
[16] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. 1054–1067. https://doi.org/10.1145/2660267.2660348
[17] Dinei Florencio and Cormac Herley. 2007. A large-scale study of web password habits. In *Proceedings of the 16th international conference on World Wide Web*. ACM, 657–666.
[18] David Freeman, Sakshi Jain, Markus Dürmuth, Battista Biggio, and Giorgio Giacinto. 2016. Who Are You? A Statistical Approach to Measuring User Authenticity.
[19] Haichang Gao, Jeff Yan, Fang Cao, Zhengya Zhang, Lei Lei, Mengyun Tang, Ping Zhang, Xin Zhou, Xuqin Wang, and Jiawei Li. 2016. A Simple Generic Attack on Text Captchas.
[20] ghacks 2011. Amazon Login May Accept Password Variants. (2011). https://www.ghacks.net/2011/01/31/amazon-login-may-accept-password-variants/
[21] Maximilian Golla, Daniel V Bailey, and Markus Dürmuth. 2017. " I want my money back!" Limiting Online Password-Guessing Financially.. In *SOUPS*.
[22] Maximilian Golla and Markus Dürmuth. 2018. On the Accuracy of Password Strength Meters. 1567–1582. https://doi.org/10.1145/3243734.3243769
[23] Ariel Gordon and Richard Allen Lundeen. 2014. Efficiently throttling user authentication. (Nov. 25 2014). US Patent 8,898,752.
[24] Benjamin Harsha, Robert Morton, Jeremiah Blocki, John Springer, and Melissa Dark. 2020. Bicycle Attacks Considered Harmful: Quantifying the Damage of Widespread Password Length Leakage. *arXiv preprint arXiv:2002.01513* (2020).
[25] Have I Been Pwned 2019. Have I Been Pwned. (2019). https://haveibeenpwned.com
[26] C. Herley and P. Van Oorschot. 2012. A Research Agenda Acknowledging the Persistence of Passwords. *IEEE Security Privacy* 10, 1 (Jan 2012), 28–36. https://doi.org/10.1109/MSP.2011.150
[27] Apple Inc. Apple Differential Privacy Technical Overview. (????). https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf Retrieved 25, Apr. 2019.
[28] Dmitry Kogan, Nathan Manohar, and Dan Boneh. 2017. T/Key: Second-Factor Authentication From Secure Hash Chains. 983–999. https://doi.org/10.1145/3133956.3133989
[29] Saranga Komanduri, Richard Shay, Patrick Gage Kelley, Michelle L Mazurek, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, and Serge Egelman. 2011. Of passwords and people: measuring the effect of password-composition policies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2595–2604.

---

[9]https://wayworkshop.org

[30] Ariel Kulik and Hadas Shachnai. 2010. There is no EPTAS for two-dimensional knapsack. *Inform. Process. Lett.* 110, 16 (2010), 707–710.

[31] LinkedIn n.d.. LinkedIn Password Corpus. (n.d.). https://hashes.org/public.php

[32] David Malone and Kevin Maher. 2012. Investigating the distribution of password choices. In *Proceedings of the 21st international conference on World Wide Web*. ACM, 301–310.

[33] Moni Naor, Benny Pinkas, and Eyal Ronen. 2019. How to (not) Share a Password: Privacy Preserving Protocols for Finding Heavy Hitters with Adversarial Behavior. 1369–1386. https://doi.org/10.1145/3319535.3363204

[34] Arvind Narayanan and Vitaly Shmatikov. 2006. How to break anonymity of the netflix prize dataset. *arXiv preprint cs/0610105* (2006).

[35] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust de-anonymization of large datasets (how to break anonymity of the Netflix prize dataset). *University of Texas at Austin* (2008).

[36] Thu Pham. 2019. STOP THE PWNAGE: 81% OF HACKING INCIDENTS USED STOLEN OR WEAK PASSWORDS. (2019). https://duo.com/decipher/stop-the-pwnage-81-of-hacking-incidents-used-stolen-or-weak-passwords Retrieved May 20, 2020.

[37] Benny Pinkas and Tomas Sander. 2002. Securing Passwords Against Dictionary Attacks. 161–170. https://doi.org/10.1145/586110.586133

[38] prowebscraper. 2019. Top 10 Captcha Solving Services Compared. (2019). https://prowebscraper.com/blog/top-10-captcha-solving-services-compared/

[39] RockYou 2010. RockYou Password Corpus. (2010). http://downloads.skullsecurity.org/passwords/rockyou.txt.bz2.

[40] Ravi Sandhu, Colin Desa, and Karuna Ganesan. 2005. System and method for password throttling. (April 19 2005). US Patent 6,883,095.

[41] Stuart Schechter, Cormac Herley, and Michael Mitzenmacher. 2010. Popularity is everything: A new approach to protecting passwords from statistical-guessing attacks. In *Proceedings of the 5th USENIX conference on Hot topics in security*. USENIX Association, 1–8.

[42] S Schecter and C Herley. 2016. The Binomial Ladder Frequency Filter and its Applications to Shared Secrets. *MSR-TR-2018-18* (2016).

[43] Apple Differential Privacy Team. Learning with Privacy at Scale. (????). https://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html Retrieved 25, Apr. 2019.

[44] TechNewsWorld 2019. Microsoft Exposes Russian Cyberattacks on Phones, Printers, Video Decoders. (2019). https://www.technewsworld.com/story/86171.html

[45] Andreas Tuerk. 2019. To stay secure online, Password Checkup has your back. (2019). https://www.blog.google/technology/safety-security/password-checkup/

[46] Liam Tung. 2019. Ransomware crooks hit Synology NAS devices with brute-force password attacks. (2019). https://www.zdnet.com/article/ransomware-crooks-hit-synology-nas-devices-with-brute-force-password-attacks/ Retrieved May 20, 2020.

[47] Blase Ur, Sean M. Segreti, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Saranga Komanduri, Darya Kurilova, Michelle L. Mazurek, William Melicher, and Richard Shay. 2015. Measuring Real-World Accuracies and Biases in Modeling Password Guessability. 463–481.

[48] Luis von Ahn, Manuel Blum, Nicholas J. Hopper, and John Langford. 2003. CAPTCHA: Using Hard AI Problems for Security. 294–311. https://doi.org/10.1007/3-540-39200-9_18

[49] Luis Von Ahn, Benjamin Maurer, Colin McMillen, David Abraham, and Manuel Blum. 2008. recaptcha: Human-based character recognition via web security measures. *Science* 321, 5895 (2008), 1465–1468.

[50] Ding Wang, Haibo Cheng, Ping Wang, Xinyi Huang, and Gaopeng Jian. 2017. ZipfâĂŹs law in passwords. *IEEE Transactions on Information Forensics and Security* 12, 11 (2017), 2776–2791.

[51] Ding Wang, Gaopeng Jian, Xinyi Huang, and Ping Wang. 2014. Zipf's Law in Passwords. Cryptology ePrint Archive, Report 2014/631. (2014). http://eprint.iacr.org/2014/631.

[52] Ding Wang and Ping Wang. 2016. On the Implications of Zipf's Law in Passwords. 111–131. https://doi.org/10.1007/978-3-319-45744-4_6

[53] Ding Wang, Zijian Zhang, Ping Wang, Jeff Yan, and Xinyi Huang. 2016. Targeted Online Password Guessing: An Underestimated Threat. 1242–1254. https://doi.org/10.1145/2976749.2978339

[54] Daniel Lowe Wheeler. 2016. zxcvbn: Low-Budget Password Strength Estimation. 157–173.

[55] Guixin Ye, Zhanyong Tang, Dingyi Fang, Zhanxing Zhu, Yansong Feng, Pengfei Xu, Xiaojiang Chen, and Zheng Wang. 2018. Yet Another Text Captcha Solver: A Generative Adversarial Network Based Approach. 332–348. https://doi.org/10.1145/3243734.3243754

[56] Stuart Schechter Yuan Tian, Cormac Herley. 2019. StopGuessing: Using Guessed Passwords to Thwart Online Guessing. In *4th IEEE European Symposium on Security and Privacy*. IEEE.

[57] ZDNet 2019. Facebook passwords are not case sensitive. (2019). https://www.zdnet.com/article/facebook-passwords-are-not-case-sensitive-update/