

Protocols for Checking Compromised Credentials

Lucy Li*, Bijeeta Pal*, Junade Ali†, Nick Sullivan†, Rahul Chatterjee*, and Thomas Ristenpart*

* Cornell Tech

† Cloudflare

ABSTRACT

To prevent credential stuffing attacks, industry best practice now proactively checks if user credentials are present in known data breaches. Recently, some web services, such as HaveIBeenPwned (HIBP) and Google Password Checkup (GPC), have started providing APIs to check for breached passwords. We refer to such services as *compromised credential checking* (C3) services. We give the first formal description of C3 services, detailing different settings and operational requirements, and we give relevant threat models.

One key security requirement is the secrecy of a user’s passwords that are being checked. Current widely deployed C3 services have the user share a small prefix of a hash computed over the user’s password. We provide a framework for empirically analyzing the leakage of such protocols, showing that in some contexts knowing the hash prefixes leads to a 12x increase in the efficacy of remote guessing attacks. We propose two new protocols that provide stronger protection for users’ passwords, implement them, and show experimentally that they remain practical to deploy.

1 INTRODUCTION

Password database breaches have become routine [8]. Such breaches enable credential stuffing attacks, in which attackers try to compromise accounts by submitting one or more passwords that were leaked with that account from another website. To counter credential stuffing, companies and other organizations have begun checking if their users’ passwords appear in breaches, and, if so, they deploy further protections (e.g., resetting the user’s passwords or otherwise warning the user). Information on what usernames and passwords have appeared in breaches is gathered either from public sources or from a third-party service. The latter democratizes access to leaked credentials, making it easy for others to help their customers gain confidence that they are not using exposed passwords. We refer to such services as *compromised credential checking* services, or C3 services in short.

Two prominent C3 services already operate. HaveIBeenPwned (HIBP) [45] was deployed by CloudFlare in 2018 and is used by many web services, including Firefox [14], EVE Online [10], and 1Password [4]. Google released a Chrome extension called Password Checkup (GPC) [9, 44] in February 2019 that allows users to check if their username-password pair appears in a compromised dataset. Both services work by having the user share with the C3 server a prefix of the hash of their password or of the hash of their username-password pair. This leaks some information about user passwords, which is problematic should the C3 server be compromised or otherwise malicious. But until now there has been no thorough investigation into the damage from the leakage of current C3 services or suggestions for protocols that provide better privacy.

We provide the first formal treatment of C3 services for different settings, including exploration of their security requirements. A C3

service must provide secrecy of credentials provided by the client, and ideally, it should also preserve secrecy of the leaked datasets held by the C3 server. The computational and bandwidth overhead for the client and especially the server should also be low. The server might hold billions of leaked records, barring use of existing cryptographic protocols for private set intersection (PSI) [29, 36], which would use a prohibitive amount of bandwidth at this scale.

Current industry-deployed C3 services therefore reduce bandwidth requirements by dividing the leaked data into buckets before executing a PSI protocol. The client shares with the C3 server the identifier of the bucket where their credentials would be found, if present in the leak dataset. Then, the client and the server engage in a protocol between the bucket held by the server and the credential held by the client to determine if their credential is indeed in the leak. In current schemes, the prefix of the hash of the user credential is used as the bucket identifier. The client shares the hash prefix (bucket identifier) of their credentials with the C3 server.

Revealing hash prefixes of the credentials may be dangerous. We outline an attack scenario against such prefix-revealing C3 services. In particular, we consider a conservative setting where an attacker obtains the hash prefix shared with the C3 server (possibly by compromising the server) and also knows the username associated with the queried credential. We rigorously evaluate the security of HIBP and GPC under this threat model via a mixture of formal and empirical analysis.

We start by considering users with a password appearing in some leak and show how to adapt a recent state-of-the-art credential tweaking attack [40] to take advantage of the knowledge of hash prefixes. In a credential tweaking attack, one uses the leaked password to determine likely guesses (usually, small tweaks on the leaked password). Via simulation, we show that our variant of credential tweaking successfully compromises 80% of such accounts within 1,000 guesses, given the transcript of a query made to the HIBP server. This is 28% more than running the best known credential tweaking attack, without knowledge of the transcript.

We also consider user accounts not present in a leak. Here we found that the leakage from the hash prefix disproportionately affects security compared to the previous case. For these user accounts, obtaining the query to HIBP enables the attacker to guess 71% of passwords within 1,000 guesses, which is a 12x increase over the success with no hash prefix information. Similarly, for GPC, our simulation shows 34% of user passwords can be guessed in 10 or fewer attempts (and 61% in 1,000 attempts), should the attacker learn the hash prefix shared with the GPC server.

The attack scenarios described are conservative because they assume the attacker can infer which queries to the C3 server are associated to which usernames. This may not be always possible. Nevertheless, caution dictates that we would prefer schemes that leak less. We therefore present two new C3 protocols, one that

checks for leaked passwords (like HIBP) and one that checks for leaked username-password pairs (like GPC). Like GPC and HIBP, we *partition* the password space before performing PSI, but we do so in a way that reduces leakage significantly.

Our first scheme works when only passwords are queried. It utilizes a novel approach that we call frequency-smoothing bucketization (FSB). The key idea is to use an estimate of the distribution of human-chosen passwords to assign passwords to buckets in a way that flattens the distribution of accessed buckets. We show how to obtain good estimates (using leaked data), and, via simulation, that FSB reduces leakage significantly. In many cases the best attack given the information leaked by the C3 protocol works no better than having no information at all. While the benefits come with some added computational complexity and bandwidth, we show via experimentation that the operational overhead for the FSB C3 server or client is comparable with the overhead from GPC, while also leaking much less information than hash prefix based C3 protocols.

We also describe a more secure bucketizing scheme that provides better privacy/bandwidth tradeoff for C3 servers that store username-password pairs. In fact this scheme was also (independently) proposed in [44], and Google plans to transition to using it in their extension. It is a simple modification of their current protocol. We refer to it as IDB, ID-based bucketization, as it uses the hash prefix of only the user identifier for bucketization (instead of the hash prefix of the username-password pair as currently used by GPC). Not having password information in the bucket identifier hides the user’s password perfectly from an attacker who obtains the client queries (assuming that passwords are independent of usernames). We implement IDB and show that the average bucket size in this setting for a hash prefix of 16 bits is similar to that of GPC (around 9,166 entries per bucket).

Contributions. In summary, the main contributions of this paper are the following:

- We provide a formalization of C3 protocols and detail the security goals for such services.
- We discuss various threat models for C3 services, and analyze the security of two widely deployed C3 protocols. We show that an attacker that learns the queries from a client can severely damage the security of the client’s passwords, should they also know the client’s username.
- We give a new C3 protocol (FSB) for checking only leaked passwords that utilizes knowledge of the human-chosen password distribution to reduce the leakage.
- We give a new C3 protocol for checking leaked username-password pairs (IDB) that bucketizes using only usernames.
- We analyze the performance and security of both new C3 protocols to show feasibility in practice.

We will release as public, open source code our server and client implementations of FSB and IDB.

2 OVERVIEW

We investigate approaches to checking credentials present in previous breaches. Several third party services provide credential checking, enabling users and companies to mitigate credential stuffing




Figure 1: A C3S service allows a client to ascertain whether a username and password appear in public breaches known to the service.

and credential tweaking attacks [24, 40, 46], an increasingly daunting problem for account security.

To date, such C3 services have not received any analysis, and indeed their design rationale has only been discussed in blog posts [5, 43]. We start by describing the architecture of such services, and then we detail relevant threat models.

C3 settings. We provide a diagrammatic summary of the abstract architecture of C3 services in Figure 1. A C3 server has access to a breach database \tilde{S} . We can think of \tilde{S} as a set of size N , which consists of either a set of passwords $\{w_1, \dots, w_N\}$ or username-password pairs $\{(u, w_1), \dots, (u, w_N)\}$. This corresponds to two types of C3 services — *password-only C3 service* and *username-password C3 service*. For example, HIBP [5] is a password-only C3 service,¹ and Google’s service GPC [9] is an example of username-password C3 service.

A *client* has as input a credential $s = (u, w)$ and wants to determine if s is at risk due to exposure. The client and server therefore engage in a set membership protocol to determine if $s \in \tilde{S}$. Here, clients can be users themselves (query C3 service using, say, a browser extension), or other web services can query the C3 service on behalf of their users. Of course, clients may make multiple queries to the C3 service, though the number of queries might be rate limited.

The ubiquity of breaches means that, nowadays, the breach database \tilde{S} will be quite large. A recently leaked compilation of previous breached data contains 1.4 billion username password pairs [21]. The HIBP database has 501 million unique passwords [5]. Google’s blog specifies that there are 4 billion username-password pairs in their database of leaked credentials [43].

C3 protocols should be able to scale to handle set membership requests for these huge datasets for millions of requests a day. HIBP reported serving around 600,000 requests per day on average [6]. The design of C3S should therefore not be computationally expensive on the server-side. The number of network round trips required must be low, and we will restrict attention to protocols that can be completed with a single HTTPS request. Finally, we will want to minimize bandwidth usage.

Threat model. Both the C3 server’s database \tilde{S} and the client’s queried password should be considered confidential. While breaches

¹ Actually HIBP also allows checking if a user identifier (email) is leaked with a data breach. For the purpose of this study, however, we only focus on the above mentioned two C3 services.

Credentials checked	Name	Bucket identifier	B/w (KB)	RTL (ms)	Security loss
Password	HIBP	20-bits of SHA1(w)	15.9	208	12x
	FSB	Figure 6, $\bar{q} = 10^2$	261	527	2x
(Username, password)	GPC	16-bits of Argon2($u w$)	606	458	10x
	IDB	16-bits of Argon2(u)	606	487	1x

Figure 2: Comparison of different C3 protocols. HIBP [5] and GPC [9] are two C3 services used in practice. We introduce frequency-smoothing bucketization (FSB) and identifier-based bucketization (IDB). Security loss is computed assuming query budget $q = 10^3$ for users who has not been compromised before.

are often made public, we prefer to treat \tilde{S} as confidential even if it consists of only public information. Of course, by querying on \tilde{S} a malicious client will fundamentally be able to check if values are in the database. Ideally such a brute-force approach would be the best possible attack.

A malicious C3 server could deviate from its protocol, for example, by lying to the client about the contents of \tilde{S} in order to encourage them to pick a weak password. Monitoring techniques might be useful to catch such misdeeds. We do not consider active attacks further, as we focus instead on the more pressing issue of not leaking (u, w) to an honest-but-curious server that follows its protocol but wants to infer information about the user’s password.

In our threat model we consider targeted attacks, where the attacker has access to the username of the querying user. This is realistic, as an attacker can learn the username corresponding to a query by linking IP addresses to usernames. An attacker who compromises the C3 server might be able to find the IP address of the querying user. The attacker can send tracking emails to all the leaked usernames present in the breach dataset. If a client clicks on the link present in the email, the attacker would be able to retrieve the IP address of the user [27]. Thereby, the attacker can associate the email corresponding to the query to the C3 server.

For the rest of the paper, we will focus on this threat model where the attacker knows the querying user’s username, and refer to it as a known-username attack (KUA). The attacker can take advantage of the leaked data to find (any) leaked passwords associated to the target username and tailor its guesses based on them.

For this paper, we will focus on online attack settings, where the attacker tries to impersonate a user by guessing their password for other web services online. These are easy to launch and are one of the most prevalent forms of attacks [16, 28]. However, in an online setting, the web service can monitor the failed login attempts and lock an account out after too many incorrect password submissions. Therefore, the attacker gets only a small number of attempts, known as the guessing budget q of the attack.

Potential approaches. A C3 protocol requires, at core, a secure set membership query. Existing protocols for private set intersection (a generalization of set membership) [22, 31, 41, 42] cannot currently scale to the set sizes required in C3 settings, $N \approx 2^{30}$. For example, the basic PSI protocol that uses an oblivious pseudorandom function (OPRF) [31] computes $y_i = F_\kappa(u_i, w_i)$ for $(u_i, w_i) \in \tilde{S}$ where F_κ is the secure OPRF with secret key κ (held by the server). It sends all y_1, \dots, y_N to the client, and the client obtains $y = F_\kappa(u, w)$ for its input (u, w) by obviously computing it with the server. The

Symbol	Description
u / \mathcal{U}	user identifier, e.g. email / domain of users
w / \mathcal{W}	password / domain of passwords
\mathcal{S}	domain of credentials
\tilde{S}	set of leaked credentials
p	distribution of username-password pairs over $\mathcal{U} \times \mathcal{W}$
p_w	distribution of passwords over \mathcal{W}
\hat{p}_s	estimate of p_w used by C3 server
q	query budget of an attacker
\bar{q}	parameter to FSB, estimated query budget of an attack

Figure 3: Descriptions of the notation used in the paper.

client can then check if $y \in \{y_1, \dots, y_N\}$. But clearly for large N this is prohibitively expensive in terms of bandwidth. One can use Bloom filters to more compactly represent the set y_1, \dots, y_N , but the result is still too large. While more advanced PSI protocols exist that improve on these results asymptotically, they are unfortunately not yet practical for this C3 setting [30, 31].

Practical C3 schemes therefore relax the security requirements, allowing the protocol to leak some information about the client’s queried (u, w) but hopefully not too much. To date no one has investigated how damaging the leakage of currently proposed schemes is, which we turn to doing next. In Figure 2, we show all the different settings for C3 we discuss in the paper, and compare their security and performance.

3 BUCKETIZATION SCHEMES AND SECURITY MODELS

In this section we formalize the security models for a class of C3 schemes that bucketize the breach dataset into smaller sets (buckets). Intuitively, a straightforward approach for checking whether or not a client’s credentials are present in a large set of leaked credentials hosted by a server is to divide the leaked data into various buckets. The client and server can then perform a private set intersection between the user’s credentials and one of the buckets (potentially) containing that credential. The bucketization makes private set membership tractable, while only leaking to the server that the password may lie in the set associated to a certain bucket.

We give a general framework to understand the security loss and bandwidth overhead of different bucketization schemes, which we use to evaluate existing C3 services.

Notation. For ease of description of the constructions that follow, we fix some notation. Let \mathcal{W} be the set of all passwords, and p_w be the associated probability distribution; let \mathcal{U} be the set of all user identifiers, and p be the joint distribution over $\mathcal{U} \times \mathcal{W}$. We will use \mathcal{S} to denote the domain of credentials being checked, i.e., for password-only C3 service, $\mathcal{S} = \mathcal{W}$, and for username-password C3 service, $\mathcal{S} = \mathcal{U} \times \mathcal{W}$. Below we will use \mathcal{S} to give a generic scheme, and specify the setting only if necessary to distinguish. Similarly, $s \in \mathcal{S}$ denotes a password or a username-password pair, based on the setting. Let \tilde{S} be the set of leaked credentials, and $|\tilde{S}| = N$.

Let H be a cryptographic hash function from $\{0, 1\}^* \mapsto \{0, 1\}^\ell$, where ℓ is a parameter to the system. We use \mathcal{B} to denote the set of buckets, and we let $\beta: \mathcal{S} \mapsto \mathcal{B}^* \setminus \{\emptyset\}$ be a bucketizing function which maps a credential to a set of buckets. A credential can be mapped to multiple buckets, and every credential is assigned to at least one bucket. An inverse function to β is $\alpha: \mathcal{B} \mapsto \mathcal{S}^*$,

$\text{Guess}^{\mathcal{A}}(q)$ $(u, w) \leftarrow \rho \mathcal{U} \times \mathcal{W}$ $\{\tilde{w}_1, \dots, \tilde{w}_q\} \leftarrow \mathcal{A}(u, q)$ $\text{return } w \in \{\tilde{w}_1, \dots, \tilde{w}_q\}$	$\text{BucketGuess}_{\beta}^{\mathcal{A}'}(q)$ $(u, w) \leftarrow \rho \mathcal{U} \times \mathcal{W}; s \leftarrow (u, w)$ $h \leftarrow \beta(s)$ $\{\tilde{w}_1, \dots, \tilde{w}_q\} \leftarrow \mathcal{A}'(u, h, q)$ $\text{return } w \in \{\tilde{w}_1, \dots, \tilde{w}_q\}$
---	---

Figure 4: The guessing games to evaluate security of different C3 schemes.

which maps a bucket to the set of all credentials it contains; so, $\alpha(b) = \{s \in \mathcal{S} \mid b \in \beta(s)\}$. Note, $\alpha(b)$ can be very large given it considers all credentials in \mathcal{S} . We let $\tilde{\alpha}$ be the function that denotes the credentials in the buckets held by the C3 server, $\tilde{\alpha}(b) = \alpha(b) \cap \tilde{\mathcal{S}}$.

The client sends b to the server, and then the client and the server engage in a set intersection protocol between $\{s\}$ and $\tilde{\alpha}(b)$.

Bucketization schemes. Bucketization is dividing the credentials held by the server into smaller buckets. The client can use the bucketizing function β to find the set of buckets for a credential, and then pick one randomly to query the server. There are different ways to bucketize the credentials.

In the first method, which we call hash-prefix-based bucketization (HPB), the credentials are partitioned based on the first l bits of a cryptographic hash of the credentials. GPC [9] and HIBP [5] APIs use HPB. The distribution of the credentials is not considered in HPB, which causes it to incur higher security loss, as we show in Section 4.

We introduce a new bucketizing method, which we call frequency-smoothing bucketization (FSB), that takes into account the distribution of the credentials and replicates credentials into multiple buckets if necessary. The replication “flattens” the conditional distribution of passwords given a bucket identifier, and therefore vastly reduces the security loss. We discuss FSB in more details in Section 5.

In both HPB and FSB, the bucketization function depends on the user’s password. We give another bucketization approach — the most secure one — that bucketizes based only on the hash prefix of the user identifier. We call this identifier-based bucketizing (IDB). The approach is only applicable for username-password C3 services. We discuss IDB in Section 4.

Security measure. The goal of an attacker is to learn the user’s password. We will focus on online-guessing attacks, where an attacker tries to guess a user’s password over the login interfaces provided by a web service. An account might be locked for too many incorrect guesses (for example, 10), and the attack fails. Therefore, we will measure an attacker’s success given a certain guessing budget, say q . We will always assume the attacker has access to the username of the target user.

The security games are given in Figure 4. The game Guess models the situation in which no information besides the username is revealed to the adversary about the password. In the game BucketGuess , the adversary also gets access to a bucket that is chosen according to the credentials $s = (u, w)$ and the bucketization function β .

We define the advantage against a game as the maximum probability that the game outputs 1. Therefore,

$$\text{Adv}^{\text{GS}}(q) = \max_{\mathcal{A}} \Pr \left[\text{Guess}^{\mathcal{A}}(q) \Rightarrow 1 \right],$$

and

$$\text{Adv}_{\beta}^{\text{b-GS}}(q) = \max_{\mathcal{A}'} \Pr \left[\text{BucketGuess}_{\beta}^{\mathcal{A}'}(q) \Rightarrow 1 \right].$$

The probabilities are taken over the choices of username-password pairs and the selection of bucket from the bucketizing function β . Security loss, $\Delta_{\beta}(q)$, of a bucketizing protocol β is defined as the ratio of $\text{Adv}_{\beta}^{\text{b-GS}}(q)$ over $\text{Adv}^{\text{GS}}(q)$.

Note,

$$\Pr \left[\text{Guess}^{\mathcal{A}}(q) \Rightarrow 1 \right] = \sum_u \Pr [w \in \mathcal{A}(u, q) \wedge U = u].$$

To maximize this probability, the attacker must pick the q most probable passwords for each user. Therefore,

$$\text{Adv}^{\text{GS}}(q) = \sum_u \max_{w_1, \dots, w_q} \sum_{i=1}^q \Pr [W = w_i \wedge U = u]. \quad (1)$$

In $\text{BucketGuess}_{\beta}$, the attacker has access to the bucket identifier, and therefore the advantage is computed as

$$\begin{aligned} \text{Adv}_{\beta}^{\text{b-GS}}(q) &= \sum_u \sum_b \max_{w_1, \dots, w_q} \sum_{i=1}^q \Pr [W = w_i \wedge U = u \wedge B = b] \\ &= \sum_u \sum_b \max_{(u, w_1), \dots, (u, w_q) \in \alpha(b)} \sum_{i=1}^q \frac{\Pr [W = w_i \wedge U = u]}{|\beta((u, w_i))|} \end{aligned}$$

The second equation follows because for $b \in \beta((u, w))$, each bucket in $\beta(w)$ is equally likely to be chosen, so

$$\Pr [B = b \mid W = w \wedge U = u] = \frac{1}{|\beta((u, w))|}.$$

The joint distribution of usernames and passwords is hard to model. To simplify the equations, we divide the users targeted by the attacker into two groups: *compromised* (users whose previously compromised accounts are available to the attacker) and *uncompromised* (users for which the attacker has no information other than their usernames).

We assume there is no direct correlation between the username and password.² Therefore, an attacker cannot use the knowledge of only the username to tailor guesses. This means that in the uncompromised setting, we assume $\Pr [W = w \mid U = u] = \Pr [W = w]$. Assuming independence of usernames and passwords, we define in the uncompromised setting

$$\lambda_q = \text{Adv}^{\text{GS}}(q) = \max_{w_1, \dots, w_q} \sum_{i=1}^q \Pr [W = w_i]. \quad (2)$$

We give analytical and empirical analysis of security in this setting, and show that the security of uncompromised users is impacted by existing C3 schemes much more than that of compromised users.

In the compromised setting, the attacker can use the username to find other leaked passwords associated with that user, which then can be used to tailor guesses [40, 46]. Analytical bounds on

²Though prior work [33, 46] suggests knowledge of only username can improve efficacy of guessing user passwords, the improvement is minimal. See Appendix A for more on this analysis.

the compromised setting are less informative, so we evaluate this setting empirically in Section 6.

Bandwidth. The bandwidth required for a bucketization scheme is determined by the size of the buckets. The maximum size of the buckets can be determined using a balls-and-bins approach [20], assuming the client picks a bucket randomly from the possible set of buckets $\beta(s)$ for a credential s , and $\beta(s)$ also maps s to a random set of buckets. In total $m = \sum_{s \in \tilde{\mathcal{S}}} |\beta(s)|$ credentials (balls) are “thrown” into $n = |\mathcal{B}|$ buckets. If $m > |\mathcal{B}| \cdot \log |\mathcal{B}|$, then following the seminal results on balls-and-bins game [20], we can show the maximum number of passwords in a bucket with very high probability $1 - o(1)$ is less than $\frac{m}{n} \cdot \left(1 + \sqrt{\frac{n \log n}{m}}\right) \leq 2 \cdot \frac{m}{n}$. We will use this formula to compute an upper bound on the bandwidth requirement for specific bucketization schemes.

4 HASH-PREFIX-BASED BUCKETIZATION

Hash-prefix-based bucketization (HPB) schemes are a simple ways to divide the credentials stored by the C3 server. In this type of C3 scheme, a prefix of the hash of the credential is used as the criteria to group the credentials into buckets — all credentials that share the same hash-prefix are assigned to the same bucket. The total number of buckets depends on l , the length the hash-prefix. The number of credentials in the buckets depends on both l and $|\tilde{\mathcal{S}}|$. We will use $H^{(l)}(\cdot)$ to denote the function that outputs the l -bit prefix of the hash $H(\cdot)$. The client shares the hash prefix of the credential they wish to check with the server. While a smaller hash prefix reveals less information to the server about the user’s password, it also increases the size of each bucket held by the server, which in turn increases the communication overhead.

Hash-prefix-based bucketization is currently being used for credential checking in industry: HIBP [5] and GPC [9]. We introduce a new HPB protocol called IDB that achieves zero security loss for any query budget. Below we will discuss the design details of these three C3 protocols.

HIBP [5]. HIBP uses HPB bucketization to provide a password-only C3 service. They do not provide compromised username-password checking. HIBP maintains a database of leaked passwords, which contains more than 501 million passwords [5]. They use the SHA1 hash function, with prefix length $l = 20$; the leaked dataset is *partitioned* into 2^{20} buckets. The prefix length is chosen to ensure no bucket is too small or too big. With $l = 20$, the smallest bucket has 381 passwords, and the maximum bucket has 584 passwords [19]. This effectively makes the user’s password k -anonymous. However, k -anonymity provides limited protection, as shown by numerous prior works [35, 38, 49] and by our security evaluation.

The passwords are hashed using SHA1 and indexed by their hash prefix for fast retrieval. A client computes the SHA1 hash of their password w and queries HIBP with the 20-bit prefix of the hash; the server responds with all the hashes that shares the same 20-bit prefix. The client then checks if the full SHA1 hash of w is present among the set of hashes sent by the server. This is a weak form of PSI that does not hide the leaked passwords from the client — the client learns the SHA1 hash of the leaked passwords and can perform brute force cracking to recover those passwords.

HIBP justifies this design choice by observing that passwords in the server side leaked dataset are publicly available for download on the Internet. Therefore, HIBP lets anyone download the hashed passwords and usernames. This can be useful for parties who want to host their own leak checking service without relying on HIBP. However, keeping the leaked dataset up-to-date can be challenging, making a third-party C3 service preferable.

HIBP trades server side privacy for protocol simplicity. The protocol also allows utilization of heavy caching on content delivery networks (CDN), such as Cloudflare.³ The caching helps HIBP to be able to serve 8 million requests a day with 99% cache hit rate (as of August 2018) [18]. The human-chosen password distribution is “heavy-headed”, that is a small number of passwords are chosen by a large number of users. Therefore, a small number of passwords are queried a large number of times, which in turn makes CDN caching much more effective.

GPC [9]. Google provides a username-password C3S, called Password Checkup (GPC). The client — a browser extension — computes the hash of the username and password together using the Argon2 hash function with the first $l = 16$ bits to determine the bucket identifier. After determining the bucket, the client engages in a private set intersection (PSI) protocol with the server. The full algorithm is given in Figure 5. GPC uses an OPRF-based PSI protocol. Let $F_{(\cdot)}(\cdot)$ be a key-homomorphic pseudo-random function (PRF) such that $F_a(\cdot) \times F_b(\cdot) = F_{ab}(\cdot)$. Under the hood, F calls the hash function H on $u||w$, and then maps the hash output onto the elliptic curve point for further computation.

The server has a secret key κ which it uses to compute the $y_i = F_{\kappa}(u_i || w_i)$. The client shares with the server the bucket id b and the PRF output of username-password pair $x = F_r(u || w)$, for some randomly sampled r . The server returns the bucket $z_b = \{y_i \mid H(u_i || w_i) = b\}$ and $y = F_{\kappa}(x)$. Finally, the client completes the OPRF computation by computing $\tilde{x} = F_{\frac{1}{r}}(y) = F_{\kappa}(u || w)$, and checking if $\tilde{x} \in z_b$.

The GPC protocol is significantly more complex than HIBP, and it does not allow easy caching by CDNs. However, it provides secrecy of server side leaked data — the best case attack is to follow the protocol to brute-force check if a password is present in the leak database.

Bandwidth. HPB assigns every credential to only one bucket; therefore, $m = \sum_{w \in \tilde{\mathcal{S}}} |\beta(w)| = |\tilde{\mathcal{S}}| = N$. The total number of buckets $n = 2^l$. Following the discussion from Section 3, maximum bandwidth for a HPB C3S should be no more than $2 \cdot \frac{m}{n} = 2 \cdot \frac{N}{2^l}$.

We experimentally verified the bandwidth value, and the sizes of the buckets for HIBP, GPC, and IDB are given in Section 7.

Security. HPB schemes like HIBP and GPC expose a prefix of the user’s password (or username-password pair) to the server. As discussed earlier, we assume the attacker knows the username of the target user. In the uncompromised setting — where the user identifier does not appear in the leaked data available to the attacker, we show that giving the attacker the hash-prefix with a guessing budget of q queries is equivalent to giving as many as $q \cdot |\mathcal{B}|$ queries (with no hash-prefix) to the attacker.

³<https://www.cloudflare.com/>

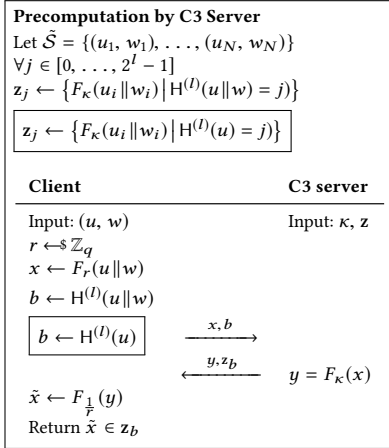


Figure 5: Algorithms for GPC, and the change in IDB given in the box. $F_{(\cdot)}$ is a PRF.

THEOREM 4.1. *Let $\beta_{\text{HPB}} : S \mapsto \mathcal{B}$ be the bucketization scheme that, for a credential $s \in S$, chooses a bucket that is a function of $H^{(l)}(s)$, where s contains the user’s password. The advantage of an attacker in this setting against previously uncompromised users is*

$$\text{Adv}_{\beta_{\text{HPB}}}^{\text{b-gs}}(q) \leq \text{Adv}^{\text{gs}}(q \cdot |\mathcal{B}|).$$

Proof: First, note that $|\beta_{\text{HPB}}(\cdot)| = 1$, as every password is assigned to exactly one of the buckets. Following the discussion from Section 3, assuming independence of usernames and passwords in the uncompromised setting, we can compute the advantage against game BucketGuess as,

$$\text{Adv}_{\beta_{\text{HPB}}}^{\text{b-gs}}(q) = \sum_{b \in \mathcal{B}} \max_{w_1, \dots, w_q} \sum_{i=1}^q \Pr[W = w_i] \leq \text{Adv}^{\text{gs}}(q \cdot |\mathcal{B}|)$$

We relax the $\alpha(b)$ notation to denote set of passwords (instead of username-password pairs) assigned to a bucket b . The inequality follows from the fact that each password is present in only one bucket. If we sum up the probabilities of the top q passwords in each bucket, the result will be at most the sum of the probabilities of the top $q \cdot |\mathcal{B}|$ passwords. Therefore, the maximum advantage achievable is $\text{Adv}^{\text{gs}}(q \cdot |\mathcal{B}|)$.

Theorem 4.1 only provides an upper bound on the security loss. Moreover, for the compromised setting, the analytical formula is less informative. So, we use empiricism to find the effective security loss against compromised and uncompromised users. We report all security simulation results in Section 6. Notably, with GPC with hash prefix length $l = 16$, an attacker can guess passwords of 60.5% of (previously uncompromised) user accounts in fewer than 1000 guesses, a 10x increase from the percent it can compromise without access to the hash prefix. (See Section 6 for more results.)

Identifier-based bucketization (IDB). As our security analysis and simulation show, the security degradation of HPB is dismal. The main issue with those protocols is that the bucket identifier is a deterministic function of the user password. We give a new C3 protocol that uses HPB style bucketing based on only username.

We call this identifier-based bucketization (IDB). IDB is defined for username-password C3 schemes.

IDB is a slight modification of the protocol used by GPC— we use the hash-prefix of the username, $H^{(l)}(u)$, instead of the hash-prefix of the username-password combination, $H^{(l)}(u \| w)$, as a bucket identifier. The scheme is described in Figure 5, using the changes in the boxed code. The bucket identifier is computed completely independent of the password (assuming username is independent of the password). Therefore, the attacker gets no additional advantage for knowing the bucket identifier.

Because IDB uses the hash-prefix of the username as the bucket identifier, two hash computations are required on the client side for each query (as opposed to one for GPC). With most modern devices, this is not a significant computing burden, but the protocol latency may be impacted, since we use a slow hash (Argon2). We show experimentally how the extra hash computation affects the latency of IDB in Section 7.

Since IDB does not use the user’s password to determine the bucket identifier, there is no security loss.

THEOREM 4.2. *With the IDB protocol, for all $q \geq 0$*

$$\text{Adv}_{\text{IDB}}^{\text{b-gs}}(q) = \text{Adv}^{\text{gs}}(q).$$

We provide the proof of this theorem in Appendix C. Because the bucket identifiers are chosen independent of the passwords, the conditional probability of the password given the bucket identifier remains the same as the probability without knowing the bucket identifier.

Overall, we can use a form of HPB to create a username-password C3S scheme with no security loss, but the password-only C3S schemes constructed using HPB lead to significant security loss. In the next section we solve this problem by introducing a more secure password-only C3S scheme.

5 FREQUENCY-SMOOTHING BUCKETIZATION

In the previous section we show how to build a username-password C3 service that does not degrade security. However, many services, such as HIBP, only provide a password-only C3 service. HIBP does not store username-password pairs so, should the HIBP server ever get compromised, an attacker cannot use their leak database to mount credential stuffing attacks. Moreover, IDB cannot be extended in any useful way to protect password-only C3 services.

Therefore, we introduce a new bucketization scheme to build secure password-only C3 services. We call this scheme frequency-smoothing bucketization (FSB). FSB assigns a password to multiple buckets based on its probability — frequent passwords are assigned to many buckets. Replicating a password into multiple buckets effectively reduces the conditional probabilities of that password given a bucket identifier. We do so in a way that makes the conditional probabilities of popular passwords similar to those of unpopular passwords to make it harder for the attacker to guess the correct password. FSB, however, is only effective for non-uniform credential distributions, such as password distributions.⁴ Therefore, FSB

⁴Usernames (e.g., emails) are unique for each users, so the distribution of usernames




Figure 6: Bucketizing function β_{FSB} for assigning passwords to buckets in FSB. Here \hat{p}_s is the distribution of passwords; $\mathcal{W}_{\bar{q}}$ is the set of top- \bar{q} passwords according to \hat{p}_s ; \mathcal{B} is the set of buckets; f is a universal hash function $f : W \mapsto \mathbb{Z}_{|\mathcal{B}|}$; $\tilde{\mathcal{S}}$ is the set of passwords hosted by the server.

cannot be used to build a username-password C3 service.

Implementing FSB requires knowledge of the distribution of human-chosen passwords. Of course, obtaining precise knowledge of the password distribution can be difficult; therefore, we will use an estimated password distribution, denoted by \hat{p}_s . Another parameter of FSB is \bar{q} , which is an estimate of the attacker’s query budget. We show that if the actual query budget $q \leq \bar{q}$, FSB has zero security loss. Larger \bar{q} will provide better security; however, it also means more replication of the passwords and larger bucket sizes. So, \bar{q} can be tuned to balance between security and bandwidth. Below we will give the two main algorithms of FSB scheme: β_{FSB} and $\tilde{\alpha}_{\text{FSB}}$, followed by bandwidth and security analysis for FSB.

Bucketizing function (β_{FSB}). To map passwords to buckets, we use a universal hash function $f : \mathcal{W} \mapsto \mathbb{Z}_{|\mathcal{B}|}$. The algorithm for bucketization $\beta_{\text{FSB}}(w)$ is given in Figure 6. The parameter \bar{q} is used in the following way: β replicates the most probable \bar{q} passwords, $\mathcal{W}_{\bar{q}}$, across all $|\mathcal{B}|$ buckets. Each of the remaining passwords are replicated proportional to their probability. A password w with probability $\hat{p}_s(w)$ is replicated exactly $\gamma = \left\lceil \frac{|\mathcal{B}| \cdot \hat{p}_s(w)}{\hat{p}_s(w_{\bar{q}})} \right\rceil$ times, where $w_{\bar{q}}$ is the \bar{q}^{th} most likely password. Exactly which buckets a password is assigned to are determined using the universal hash function f . Each bucket is assigned an identifier between $[0, |\mathcal{B}| - 1]$. A password w is assigned to the buckets whose identifiers fall in the range $[f(w), f(w) + \gamma - 1]$. The range can wrap around. For example, if $f(w) + \gamma > |\mathcal{B}|$, then the password is assigned to the buckets in the range $[0, f(w) + \gamma - 1 \bmod |\mathcal{B}|]$ and $[f(w), |\mathcal{B}| - 1]$.

Bucket retrieving function ($\tilde{\alpha}$). Retrieving passwords assigned to a bucket is challenging in FSB. An inefficient – linear in N – implementation of $\tilde{\alpha}$ is given in Figure 6. Storing the contents of each bucket separately is not feasible, since the number of buckets in FSB can be very large, $|\mathcal{B}| \approx N$. To solve the problem, we utilize the structure of the bucketizing procedure where passwords are assigned to buckets in continuous intervals. This allows us to use an interval tree [7] data structure to store the intervals for all of the passwords. Interval trees allow fast queries to retrieve the set of intervals that contain a queried point (or interval) – exactly what is needed to instantiate $\tilde{\alpha}$.

This efficiency comes with increased storage cost. To store N entries in an interval tree, we require $\mathcal{O}(N \log N)$ storage. The tree can be built in $\mathcal{O}(N \log N)$ time, and each query takes $\mathcal{O}(\log N + |\tilde{\alpha}(b)|)$

and username-password pairs are close to uniform.

time. The big-O notation only hides small constants.

Estimating password distributions. To construct the bucketization algorithm for FSB, the server needs an estimate of the password distribution (p_w). This estimate will be used by both the server and the client to assign passwords to buckets. One possible estimate is the histogram of the passwords in the leaked data $\tilde{\mathcal{S}}$. Histogram estimates are typically accurate for popular passwords, but such estimates are not complete – passwords that are not in the leaked dataset will have zero probability according to this estimate. Moreover, sending the histogram over to the client is expensive in terms of bandwidth and security critical. We also considered password strength meters, such as zxcvbn [47] as a proxy for a probability estimate. However, this estimate turned out to be too coarse for our purposes. For example, more than 10^5 passwords had a “probability” of greater than 10^{-3} .

We build a 3-gram password model \hat{p}_n using the leaked passwords present in $\tilde{\mathcal{S}}$. Markov models or n -gram models are shown to be effective at estimating human-chosen password distributions [34], and they are very fast to train and run (unlike neural network based password distribution estimators, such as [37]). However, we found the n -gram model assigns very low probabilities to popular passwords. The sum of the probabilities of the top 1,000 passwords as estimated by the 3-gram model is only 0.0012, whereas in practice the top 1000 passwords are chosen by 5.6% of users.

We therefore use a combined approach that uses a histogram model for the popular passwords and the 3-gram model for the rest of the distribution. Such combined techniques are also used in practice for password strength estimation [37, 47]. Let \hat{p}_s be the estimated password distribution used by FSB. Let \hat{p}_h be the distribution of passwords implied by the histogram of passwords present in $\tilde{\mathcal{S}}$. Let $\tilde{\mathcal{S}}_t$ be the set of the t most probable passwords according to \hat{p}_h . We used $t = 10^6$.

$$\hat{p}_s(w) = \begin{cases} \hat{p}_h(w) & \text{if } w \in \tilde{\mathcal{S}}_t, \\ \hat{p}_n(w) \cdot \frac{1 - \sum_{w \in \tilde{\mathcal{S}}_t} \hat{p}_h(w)}{1 - \sum_{w \in \tilde{\mathcal{S}}_t} \hat{p}_n(w)} & \text{otherwise.} \end{cases}$$

Bandwidth. We use the formulation provided in Section 3 to compute the bandwidth requirement for FSB. In this case, $m = |\mathcal{B}| \cdot \bar{q} + \frac{|\mathcal{B}|}{\hat{p}_s(w_{\bar{q}})} + N$, and $n = |\mathcal{B}|$. Therefore, the maximum size of a bucket is with high probability less than $2 \cdot \left(\bar{q} + \frac{1}{\hat{p}_s(w_{\bar{q}})} + \frac{N}{|\mathcal{B}|} \right)$. The details of this analysis are given in Appendix B.

In practice, we can choose the number of buckets to be such that $|\mathcal{B}| = N$. Then, the number of passwords in a bucket depends primarily on the parameter \bar{q} . Note, bucket size increases with \bar{q} .

Security analysis. We show that there is no security loss in the uncompromised setting for FSB when the actual number of guesses q is less than the parameter \bar{q} , and we give an upper bound for the security loss when q exceeds \bar{q} .

THEOREM 5.1. *If a frequency based bucketization scheme ensures $\forall w \in \mathcal{W}, |\beta_{\text{FSB}}(w)| = \min \left\{ |\mathcal{B}|, \left\lceil \frac{|\mathcal{B}| \cdot \hat{p}_s(w)}{\hat{p}_s(w_{\bar{q}})} \right\rceil \right\}$, then for the uncompromised users,*

$$(1) \quad \text{Adv}_{\beta_{\text{FSB}}}^{\text{b-gS}}(q) = \text{Adv}^{\text{gS}}(q) \text{ for } q \leq \bar{q}, \text{ and}$$

(2) for $q > \bar{q}$,

$$\frac{\lambda_q - \lambda_{\bar{q}}}{2} \leq \Delta_q \leq (q - \bar{q}) \cdot \hat{p}_s(w_{\bar{q}}) - (\lambda_q - \lambda_{\bar{q}})$$

The full proof is included in Appendix D. Intuitively, since the top q passwords are repeated across all buckets, having a bucket identifier does not allow an attacker to easily guess these q passwords. Moreover, the conditional probability of these q passwords given the bucket is greater than that of any other password in the bucket. Therefore, the attacker’s best choice is to guess the top q passwords, meaning that it does not get any additional advantage when $\bar{q} \leq q$, leading to part (1) of the theorem.

The proof of part (2) follows from the upper and lower bounds on the number of buckets each password beyond the top q is placed within. The bounds we prove show that the additional advantage in guessing the password in q queries is less than the number of additional queries times the probability of the \bar{q}^{th} password and at least half the difference in the guessing probabilities λ_q and $\lambda_{\bar{q}}$ (defined in Equation (2)).

Note that this analysis of security loss is based on the assumption that the FSB scheme has access to the precise password distribution, $\hat{p}_s = p_w$. We empirically analyze the security loss in Section 6 for $\hat{p}_s \neq p_w$, in both the compromised and uncompromised settings.

6 EMPIRICAL SECURITY EVALUATION

In this section we empirically evaluate and compare the security loss for different password-only C3 schemes we have discussed so far – hash-prefix-based bucketization (HPB) and frequency-smoothing bucketization (FSB).

We focus on known-username attacks (KUA), since in many deployment settings a curious (or compromised) C3 server can figure out the username of the querying user. We separate our analysis into two settings: previously *compromised* users, where the attacker has access to one or more existing passwords of the target user, and previously *uncompromised* users, where no password corresponding to the user is known to the attacker (or present in the breached data).

Recall, according to our threat model, we assume the adversary has knowledge of all the leak dataset C3 is using. This situation is realistic, since many password breaches are readily available for download online. For each setting the attacker also knows the bucketizing algorithm. The attacker obtains (possibly by compromising the C3 service) the bucket identifier that the client queried, as well as the user’s username or email. Our analysis will also show what an honest-but-curious C3 server would learn about the passwords of a user who participated in the protocol.

First we will look into the unrestricted setting where no password policy is enforced, and the attacker and the C3 server have the same amount of information about the password distribution. In the second experiment, we analyze the effect on security of giving the attacker more information compared to the C3 server (defender) by having a password policy that the attacker is aware of but the C3 server is not.

Password breach dataset. We used the breach dataset used in [40]. The dataset was derived from a previous breach compilation [21]

	\tilde{S}	T	$T \cap \tilde{S}$	T_{sp}	$T_{\text{sp}} \cap \tilde{S}$
# users	383.2	7.5	5.6 (76%)	4.8	3.7 (77%)
# passwords	255.2	5.4	3.6 (67%)	4.0	2.4 (60%)
# user-pw pairs	748.9	7.5	2.8 (37%)	4.9	1.8 (37%)

Figure 7: Number of entries (in millions) in the breach dataset \tilde{S} , test dataset T , and the site-policy test subset T_{sp} . Also reported are the intersections (of users, passwords, and user-password pairs, separately) between the test dataset entries and the whole breach dataset that the attacker has access to. The percentage values refer to the fraction of the values in each test set that also appear in the intersections.

dataset containing about 1.4 billion username-password pairs. The data was cleaned by removing non-ASCII characters and passwords longer than 30 characters. The authors of [40] also joined accounts with similar usernames and passwords using a method they called the *mixed method*. The usernames with only one email and password were removed, which in total removed 650 million username-password pairs. We obtained this joined and filtered dataset from the authors and performed our empirical analysis on that dataset. Removal of the 650 million pairs for users with only one password can only affect the experiment on the security for uncompromised users. Given the large size of the dataset, we expect our results on attack success are not impacted in any significant way by the removal of those accounts.

The final dataset consists of about 756 million username-password pairs.⁵ We remove 1% of username-password pairs to use as test data, denoted as T . The remaining 99% of the data is used to simulate the database of leaked credentials \tilde{S} . For the experiments with an enforced password policy, we took the username-password pairs in T that met the requirements of the password policy to create T_{sp} . We use T_{sp} to simulate queries from a website which only allows passwords that are at least 8 characters long and are not present in Twitter’s list of banned passwords [11]. For all attack simulations, the target user-password pairs are sampled from the test dataset T (or T_{sp}).

In Figure 7, we report some statistics about T , T_{sp} , and \tilde{S} . Notably, 5.6 million (76%) of the users in T are also present in \tilde{S} . This is likely because users in the joined breach compilation dataset have at least two passwords. If the 650 million singleton users had not been removed, we expect that this number would be smaller. Among the username-password pairs, 2.8 million (37%) pairs in T are also present in \tilde{S} . This means an attacker will be able to compromise 37% of the accounts (which is 50% of the previously compromised accounts) trivially with credential stuffing. In the site-policy enforced test data T_{sp} , a similar proportion of the users (77%) and username-password pairs (37%) are also present in \tilde{S} .

Experiment setup. We want to understand the impact of revealing a bucket identifier on the security of uncompromised and compromised users separately. As we can see from Figure 7, a large proportion of users in T are also present in \tilde{S} . We therefore split T into two parts: one with only username-password pairs from compromised users, T_{comp} (users with at least one password present in \tilde{S}), and another with only pairs from uncompromised users

⁵Note, there are duplicate username-password pairs in this dataset.

T_{uncomp} . We take two sets of random samples of 5000 username-password pairs⁶, one from T_{comp} , and another from T_{uncomp} . For each pair (u, w) , we run the games Guess and BucketGuess as specified in Figure 4. We record the results for guessing budgets of $q \in \{1, 10, 10^2, 10^3\}$. We repeat each of the experiments 5 times and report the averages in Figure 8.

For HPB, we compared implementations using hash prefixes of lengths $l \in \{12, 16, 20\}$. We use the SHA256 hash function with a salt, though the choice of hash function does not have a noticeable impact on the results.

For FSB, we used interval tree data structures to store the leaked passwords in $\tilde{\mathcal{S}}$ for fast retrieval of $\tilde{\alpha}(b)$. We used $|\mathcal{B}| = 2^{30}$ buckets and the hash function f is set to $f(x) = H^{(30)}(x)$, the 30-bit prefix of the (salted) SHA256 hash of the password.

Attack strategy. The attacker’s goal is to maximize its success in winning the games Guess and BucketGuess. In Equation (1) and Equation (2) we outline the advantage of attackers against Guess and BucketGuess, and thereby specify the best strategies for attacks. Guess denotes the baseline attack success rate in a scenario where the attacker does not have access to bucket identifiers corresponding to users’ passwords. Therefore the best strategy for the attacker \mathcal{A} is to output the q most probable passwords according to its best knowledge of the password distribution.

The optimal attack strategy for \mathcal{A}' in BucketGuess will be to find a list of passwords according to the following equation,

$$\operatorname{argmax}_{\substack{w_1, \dots, w_q \\ b \in \beta((u, w_i))}} \sum_{i=1}^q \frac{\Pr[W = w_i \mid U = u]}{|\beta((u, w_i))|},$$

where the bucket identifier b and user identifier u are provided to the attacker. This is equivalent to taking the top- q passwords in the set $\alpha(b)$ ordered by $\Pr[W = w \mid U = u] / |\beta((u, w))|$.

We compute the list of guesses outputted by the attacker for a user u and bucket b in the following way. For the compromised users, i.e., if $(u, \cdot) \in \tilde{\mathcal{S}}$, the attacker first considers the list of 10^4 targeted guesses generated based on the credential tweaking attack introduced in [40]. If any of these passwords belong to $\alpha(b)$ they are guessed first. This step is skipped for uncompromised users.

For the remaining guesses, we first construct a list of candidates L . L consists of the 10^6 most frequent passwords in $\tilde{\mathcal{S}}$ and 500×10^6 passwords generated from the 3-gram password distribution model \hat{p}_n . Each password w in L is assigned a weight $\hat{p}_s(w) / |\beta((u, w))|$ (See Section 5 for details on \hat{p}_s and \hat{p}_n). The list L is pruned to only contain unique guesses. Note L is constructed independent of the username or bucket identifier, and it is reordered based on the weight values. Therefore, it is constructed once for each bucketization strategy. Finally, based on the bucket identifier b , the remaining guesses are chosen from $\{\alpha(b) \cap (u, w) \mid w \in L\}$ in descending order of weight.

For the HPB implementation, each password is mapped to one bucket, so $|\beta(w)| = 1$ for all w . For FSB, $|\beta(\cdot)|$ can be calculated using the equation in Theorem 5.1.

Results. We report the success rates of the attack simulations

⁶There was a low standard deviation between results for different random samples of 5000 pairs.

in Figure 8. The baseline success rate (first row) is the advantage Adv^{GS} , computed using the same attack strategy stated above except with no information about the bucket identifier. The following rows record the success rate of the attack for HPB and FSB with different parameter choices. The estimated security loss (Δ_q) can be calculated by subtracting the baseline success rate from the HPB and FSB attack success rates.

The security loss from using HPB is devastating, especially for previously uncompromised users. Accessibility to the $l = 20$ -bit hash prefix, used by HIBP [5], allows an attacker to compromise 34% of previously uncompromised users in just one guess. In fewer than 10^3 guesses, that attacker can compromise more than 70% of the accounts (12x more than the baseline success rate with 10^3 guesses). Google Password Checker (GPC) uses $l = 16$ for its username-password C3 service. Against GPC, an attacker only needs 10 guesses per account to compromise 34% of accounts. Reducing the prefix length l can decrease the attacker’s advantage. However, that would also increase the bucket size. As we see for $l = 12$, the average bucket size is 62,309, so the bandwidth required to perform the credential check would be high.

FSB resists guessing attacks much better than HPB does. For $q \leq \bar{q}$ the attacker gets no additional advantage, even with the estimated password distribution \hat{p}_s . The security loss for FSB when $q > \bar{q}$ is much smaller than that of HPB, even with smaller bucket sizes. For example, the additional advantage over the baseline against FSB with $q = 100$ and $\bar{q} = 10$ is only 3%, despite FSB also having smaller bucket sizes than HPB with $l = 16$. Similarly for $\bar{q} = 100$, $\Delta_{10^3} = 3.6\%$. This is because the conditional distribution of passwords given an FSB bucket identifier is nearly uniform, making it harder for an attacker to guess the correct password in the bucket $\alpha(b)$ in q guesses.

For previously compromised users – users present in $\tilde{\mathcal{S}}$ – even the baseline success rate is very high: 37% of account passwords can be guessed in 1 guess and 53% can be guessed in fewer than 1,000 guesses. The advantage is supplemented even further with access to the hash prefix. As per the guessing strategy, the attacker first guesses the leaked passwords that are both associated to the user and in $\alpha(b)$. This turns out to be very effective. Due to the high baseline success rate the relative increase is low; nevertheless, in total, an attacker can guess the passwords of 80% of previously compromised users in fewer than 1,000 guesses. For FSB, the security loss for compromised users is comparable to the loss against uncompromised users for $q \leq \bar{q}$. Particularly for $\bar{q} = 10$ and $q = 100$, the attacker’s additional success is only 1.9%. Similarly, for $\bar{q} = 100$ an attacker gets at most 2.1% additional advantage for a guessing budget of $q=1,000$. Interestingly, FSB performs significantly worse for compromised users compared to uncompromised users for $q = 1$. This is because the FSB bucketing strategy does not take into account targeted password distributions, and the first guess in the compromised setting is based on the credential tweaking attack.

In our simulation, previously compromised users made up around 76% of the test set; it is unclear what is the actual proportion would be in the real world, so we do not combine results from the uncompromised and compromised settings.

As we can see, since the bucket sizes for FSB with $\bar{q} = 100$ and HPB with $l = 16$ are comparable, we will use $\bar{q} = 100$ as the

Protocol	Params	Bucket size		Uncompromised				Compromised			
		Avg	max	$q = 1$	$q = 10$	$q = 10^2$	$q = 10^3$	$q = 1$	$q = 10$	$q = 10^2$	$q = 10^3$
Baseline	N/A	N/A	N/A	0.6	1.3	2.5	5.9	37.4	50.4	51.4	52.6
HPB	$l = 20^{\ddagger}$	244	303	33.7	49.7	63.0	71.0	64.9	71.8	76.6	79.9
	$l = 16^{\dagger}$	3,896	4,138	18.3	34.3	47.6	60.5	58.2	65.0	71.1	75.7
	$l = 12$	62,309	63,173	8.4	18.0	31.6	45.1	53.5	58.2	63.9	70.0
FSB	$\bar{q} = 1$	76	112	0.6	5.7	69.9	71.0	51.3	53.3	79.4	79.9
	$\bar{q} = 10$	908	1,010	0.6	1.3	5.5	70.0	51.1	51.6	53.3	79.5
	$\bar{q} = 10^2$	5,635	5,876	0.6	1.3	2.5	9.4	50.7	51.5	52.1	54.7
	$\bar{q} = 10^3$	21,107	21,550	0.6	1.3	2.5	5.8	50.4	51.5	52.1	53.3

\ddagger HIBP uses $l = 20$ for its password-only C3 service. \dagger GPC uses $l = 16$ for username-password C3 service.

Figure 8: Comparison of attack success rate given q queries on different password-only C3 settings. All success rates are in percent (%) of the total number of samples (5,000). Bucket size, the number of passwords associated to a bucket, is measured on a random sample of 10,000 buckets.

parameter for FSB for further security and performance analysis. Note, GPC has a username-password C3 service and therefore, its bucket sizes will be larger. (See Figure 10.)

Password policy experiment. In the previous set of experiments, we assumed that the C3 server and the attacker use the same estimate of the password distribution. To simulate the effect when the attacker has a better estimate of the password distribution than the C3 server, we simulated a website which enforces a password policy. We assume that the policy is known to the attacker but not to the C3 server.

For our sample password policy, we required that passwords have at least 8 characters and that they must not be on Twitter’s banned password list [11]. The test samples are drawn from T_{sp} , username-password pairs from T where passwords follow this policy, and the attacker is also given the ability to tailor their guesses to this policy. The server still stores all passwords in \tilde{S} , without regard to this policy. Notably, the FSB scheme relies on a good estimate of the password distribution to be effective in distributing passwords evenly across buckets. Its estimate, when compared to the distribution of passwords in T_{sp} , should be less accurate than it was in the regular simulation, when compared to the password distribution from T .

We chose the parameters $k = 16$ for HPB and $\bar{q} = 100$ for FSB, because they were the most representative of how the HPB and FSB bucketization schemes compare to each other. These parameters also lead to similar bucket sizes, with around 5,000 passwords per bucket. Overall, we see that the success rate of an attacker decreases in this simulation compared to the general experiment (without a password policy). This is likely due to the fact that after removing popular passwords, the remaining set of passwords that we can choose from has higher entropy, and each password is harder to guess. FSB still defends much better against the attack than HPB does, even though the password distribution estimate used by the FSB implementation is quite inaccurate, especially at the head of the distribution. FSB assigns larger probability estimates to passwords that are banned according to our password policy.

We also see that due to the inaccurate estimate by the C3 server for FSB, we start to see some security loss for an adversary with guessing budget $q = 100$. In the general simulation, the password estimate \hat{p}_s used by the server was closer to p , so we didn’t have any noticeable security loss where $q \leq \bar{q}$.

Protocol	Uncompromised				Compromised			
	$q = 1$	10	10^2	10^3	$q = 1$	10	10^2	10^3
Baseline	0.1	0.4	1.2	3.1	37.9	46.7	47.0	47.8
HPB ($l = 16$)	9.6	17.8	27.0	41.2	51.0	55.7	59.5	63.8
FSB ($\bar{q} = 10^2$)	0.1	0.4	1.4	9.6	46.8	47.1	47.3	51.2

Figure 9: Attack success rate (in %) comparison for HPB with $l = 16$ (effectively GPC) and FSB with $\bar{q} = 10^2$ for password policy simulation. The first row records the baseline success rate $\text{Adv}^{\text{GS}}(q)$. There were 5,000 samples each from the uncompromised and compromised settings.

7 PERFORMANCE EVALUATION

In this section, we implement different approaches to checking compromised credentials and evaluate their computational overheads. For fair comparison, in addition to the algorithms we propose, FSB and IDB, we also implement HIBP and GPC with our breach dataset.

Setup. We build C3 services as serverless web applications that provide REST APIs. We used AWS Lambda [1] for the server-side computation and Amazon DynamoDB [3] to store the data. The benefit of using AWS Lambda is it can be easily deployed as Lambda@Edge and integrated with Amazon’s content delivery network (CDN), called CloudFront [2]. (HIBP uses Cloudflare as CDN to serve more than 600,000 requests per day [6].) We used Javascript to implement the server and the client side functionalities. The server is implemented as a Node-JS app. We provisioned the Lambda workers to have maximum 3GB of memory. For cryptographic operations, we used a Node-JS library called Crypto [12].

For pre-processing and pre-computation of the data we used a desktop with an Intel Core i9 processor and 128 GB RAM. Though some of the computation (e.g., hash computations) can be expedited using GPUs, we did not use any for our experiment. We used the same machine to act as the client. The round trip network latency of the Lambda API from the client machine takes about 130 milliseconds. Recall that the breach dataset we use contains 255 million unique passwords and 749 million unique username-password pairs. (See Figure 7.)

To measure the performance of each scheme, we pick 20 random passwords from the test set T and run the full C3 protocol with each one. We report the average time taken for each run in Figure 10. In the figure, we also give the break down of the time taken by the server and the client for different operations. The network latency had very high standard deviation (25%), though all other

measurements had low ($< 1\%$) standard deviation compared to the mean.

HIBP. The implementation of HIBP is the simplest among the four schemes. The set of passwords in \tilde{S} is hashed using SHA256 and split into 2^{20} buckets based on the first 20 bits of the hash value (we picked SHA256 because we also used the same for FSB). Because the bucket sizes in HIBP are so small (< 500), each bucket is stored as a single value in a DynamoDB cell, where the key is the hash prefix. For larger leaked datasets, each bucket can be split into multiple cells. The client sends the 20 bit prefix of the SHA256 hash of their password, and the server responds with the corresponding bucket.

Among all the protocols HIBP is the fastest (but also weakest in terms of security). It takes only 208 ms on average to complete a query over WAN. Most of the time is spent in round-trip network latency and the query to DynamoDB. The only cryptographic operation on the client side is a SHA256 hash of the password, which takes less than 1 ms.

FSB. The implementation of FSB is more complicated than that of HIBP. Because we have more than 1 billion buckets for FSB and each password is replicated in potentially many buckets, storing all the buckets explicitly would require too much storage overhead. We use interval trees [7] to quickly recover the passwords in a bucket without explicitly storing each bucket. Each password w in the breach database is represented as an interval specified by $\beta_{\text{FSB}}(w)$. We stored each node of the tree as a separate cell in DynamoDB. We retrieved the intervals (passwords) intersecting a particular value (bucket identifier) by querying the nodes stored in DynamoDB. FSB also needs an estimate of the password distribution to get the interval range for a tree. We use \hat{p}_s as described in Section 4. The description of \hat{p}_s takes 8.9 MB of space that needs to be included as part of the client side code. This is only a one-time cost during client installation.

The depth of the interval tree is $\log N$, where N is the number of intervals (passwords) in the tree. Since each node in the tree is stored as a separate key-value pair in the database, one client query requires $\log N$ queries to DynamoDB. To reduce this cost, we split the interval tree into r trees over different ranges of intervals, such that the i -th tree is over the interval $[(i-1) \cdot \lfloor |\mathcal{B}|/r \rfloor, i \cdot \lfloor |\mathcal{B}|/r \rfloor - 1]$. The passwords whose bucket intervals span across multiple ranges are present in all corresponding trees. We used $r = 64$, as it ensures each tree has around 4 million passwords, and the total storage overhead is less than 1% more than if we stored one large tree.

Each interval tree of 4 million passwords was generated in parallel and took 3 hours in our server. Each interval tree takes 400 MB of storage in DynamoDB, and in total 25 GB of space. FSB is the slowest among all the protocols, mainly due to multiple DynamoDB calls, which cumulatively take 273 ms (half of the total time, including network latency). This can be sped up by using a better implementation of interval trees on top of DynamoDB, such as storing a whole subtree in a DynamoDB cell instead of storing each tree node separately. We can also split the range of the range tree into more granular intervals to reduce each tree size. Nevertheless, as the round trip time for FSB is small (527 ms), we leave such optimization for future work. The maximum amount of memory used by the server is less than 81 MB during an API call.

On the client side, the computational overhead is minimal. The

Protocol	Client			Server		Total time	Bucket size
	Crypto	Server call	Comp	DB call	Crypto		
HIBP	1	205	2	40	–	208	244
FSB	1	524	2	273	–	527	3,086
GPC	47	402	9	71	6	458	9,164
IDB	72	405	10	73	6	487	9,166

Figure 10: Time taken in milliseconds to make a C3 API call. The client and server columns contain the time taken to perform client side and server side operations respectively.

client performs one SHA256 hash computation. The network bandwidth consumed for sending the bucket of hash values from the server takes on average 261 KB.

IDB and GPC. Implementations of IDB and GPC are very similar. We used the same platform – AWS Lambda and DynamoDB – to implement these two schemes. All the hash computations used here are Argon2id with default parameters, since GPC in [9] uses Argon2. During precomputation, the server computes the Argon2 hash of each username-password pair and raises it to the power of the server’s key κ . These values can be further (fast) hashed to reduce their representation size, which saves disk space and bandwidth. However, hashing would make it difficult to rotate server key. We therefore store the exponentiated Argon2 hash values in the database, and hash them further during the online phase of the protocol. The hash values are indexed and bucketized based on either $H^{(l)}(u||w)$ (for GPC) or $H^{(l)}(u)$ (for IDB). We used $l = 16$ for both GPC and IDB, as proposed in [9].

The server (for both IDB and GPC) only performs one elliptic curve exponentiation, which on average takes 6 ms. The remaining time incurred is from network latency and calling Amazon DynamoDB.

On the client side, one Argon2 hash has to be computed for GPC and two for IDB. Computing the Argon2 hash of the username-password pairs takes on an average 20 ms on the desktop machine. We also tried the same Argon2 hash computation on a personal laptop (Macbook Pro), and it took 8 ms. In total, hashing and exponentiation takes 47 ms for GPC, and 72 ms (an additional 25 ms) for IDB. The cost of checking the bucket is also higher (compared to HIBP and FSB) due to larger bucket sizes.

IDB takes only 31 ms more time on average than GPC (due to one extra Argon2 hashing), while also leaking no additional information about the user’s password. It is the most secure among all the protocols we discussed (should username-password pairs be available in the leak dataset), and runs in a reasonable time.

8 DEPLOYMENT DISCUSSION

Here we discuss different ways C3 services can be used and associated threats that need to be considered. A C3 service can be queried while creating a password – during registration or password change – to ensure the new password is not present in a leak. In this setting C3 is queried from a web server, and the client IP is potentially not revealed to the server. This, we believe, is a safer setting to use than the one we will discuss below.

In another scenario, a user can directly query a C3 service. A user can look for leaked passwords themselves by visiting a web site or using a browser plugin, such as 1Password [4] or Password

Checkup [9]. This is the most prevalent use case of C3. For example, the client can regularly check with a C3 service to proactively safeguard user accounts from potential credential stuffing attacks.

However, there are several security concerns with this setting. Primarily, the client’s IP is revealed to the C3 server in this setting, making it easier for the attacker to deanonymize the user. Moreover, multiple queries from the same user can lead to a more devastating attack. Below we give two new threat models that need to be considered for secure deployment of C3 services (where bucket identifiers depend on the password).

Regular password checks. A user or webservice might want to regularly check their passwords with C3 services. Therefore, a compromised C3 server may learn multiple queries from the same user, which can enable potentially powerful attacks. For FSB the bucket identifier is chosen randomly, so knowing multiple bucket identifiers for the same password will help an attacker narrow down the password search space and significantly improve attack success.

We can mitigate this problem for FSB by derandomizing the client side bucket selection using a client side state (e.g. browser cookie) so the client always selects the same bucket for the same password. We let the c be the client side cookie. To check a password w with the C3 server, the client picks the j^{th} bucket from the range $\beta(w)$, where $j \leftarrow f(w||c) \bmod |\beta(w)|$.

This derandomization ensures queries from the same device are deterministic (after the cookie is set). However, if the attacker can link queries of the same user from two different devices, the mitigation is ineffective. If the cookie is stolen from the client device, then the security of FSB is effectively reduced to that of HPB with similar bucket sizes.

Similarly, if an attacker can track the interaction history between a user and a C3 service, it can obtain better insight about the user’s passwords. For example, if a user who regularly checks with a C3 service stops checking a particular bucket identifier, that could mean the associated password may appear in the most up-to-date leaked dataset, and the attacker can use that information to guess the user’s password(s).

Checking similar passwords. Another important issue is querying the C3 service with multiple correlated passwords. Some web services, like 1Password, use HIBP to check multiple passwords for a user. As shown by prior work, passwords chosen by the same user are often correlated [24, 40, 46]. An attacker who can see bucket identifiers of multiple correlated passwords can mount a stronger attack. Such an attack would require estimating the joint distribution over passwords. We leave analysis of this threat model for future work.

9 RELATED WORK

Private set intersection. The protocol task facing C3 services is private set membership, a special case of private set intersection (PSI) [29, 36]. The latter allows two parties to find the intersection between their private sets without revealing any additional information.

Even state-of-the-art PSI protocols do not scale to the sizes needed for our application. For example, Kiss et al. [30] proposed an efficient PSI protocol for unequal set sizes based on oblivious

pseudo-random functions (OPRF). It performs well for sets with millions of elements, but the bandwidth usage scales proportionally to the size of the set and so performance is prohibitive in our setting. Other efficient solutions to PSI [22, 31, 41, 42] have similar prohibitive bandwidth usage.

Private information retrieval (PIR) [23] is another cryptographic primitive used to retrieve information from a server. Assuming the server’s dataset is public, the client can use PIR to privately retrieve the entry corresponding to their password from the server. But in our setting we also want to protect the privacy of the dataset leak. Even if we relaxed that security requirement, the most advanced PIR schemes [17, 39] require exchanging large amounts of information over the network, so they are not useful for checking leaked passwords. PIR with two non-colluding servers can provide better security [26] than the bucketization-based C3 schemes, with communication complexity sub-polynomial in the size of the leaked dataset. However, it requires building a C3 service with two servers guaranteed to not collude, which may be practically difficult.

Compromised credential checking. To the best of our knowledge, HIBP was the first publicly available C3 service. Junade Ali designed the current HIBP protocol which uses bucketization via prefix hashing to limit leakage. Google’s Password Checker extends this idea to use PSI, which minimizes the information about the leak revealed to clients. They also moved to checking username, password pairs.

Google’s system was described in a paper by Thomas et al. [44], which became available to us after we began work on this paper. They introduced the design and implementation of their Google Password Checker and report on measurements of its initial deployment. They recognized that their first generation protocol leaks some bits of information about passwords, but did not analyze the potential impact on password guessability. They also propose (what we call) the ID-based protocol as a way to avoid this leakage. Our paper provides further motivation for their planned transition to it.

Thomas et al. point out that password-only C3 services are likely to have high false positive rates. Our new protocol FSB, being in the password-only setting, inherits this limitation. That said, should one want to do password-only C3 (e.g., because storing username, password pairs is considered too high a liability given their utility for credential stuffing), FSB represents the best known approach.

Other C3 services include, for example, Vericlouds [15] and GhostProject [13]. They allow users to register with an email address, and regularly keep the user aware of any leaked (sensitive) information associated with that email. Such services send information to the email address, and the user implicitly authenticates (proves ownership of the email) by having access to the email address. These services are not anonymous and must be used by the primary user. Moreover, these services cannot be used for password-only C3.

Distribution-sensitive cryptography. Our FSB protocol uses an estimate of the distribution of human chosen passwords, making it an example of distribution-sensitive cryptography, in which constructions use contextual information about distributions in order to improve security. Previous distribution-sensitive approaches include Woodage et al. [48], who introduced a new type of secure

sketch [25] for password typos, and Lacharite et al.'s [32] frequency-smoothing encryption. While similar in that they use distributional knowledge, their constructions do not apply in our setting.

10 CONCLUSION

We explore different settings and threat models associated with checking compromised credentials (C3). The main concern is the secrecy of the user passwords that is being checked. We show, via simulations, that the existing industry deployed C3 services (such as HIBP and GPC) do not provide adequate security. Indeed an attacker who obtains the query to such a C3 service and the username of the querying user can severely damage the secrecy of the password. We give more secure C3 protocols for checking leaked passwords and username-password pairs. We implemented and deployed different C3 protocols on AWS Lambda and evaluated their computational and bandwidth overhead. We finish with several nuanced threat models and deployment discussions that should be considered when deploying C3 services.

REFERENCES

- [1] 2018. Argon2. <https://www.npmjs.com/package/argon2/>.
- [2] 2018. CloudFront. <https://aws.amazon.com/cloudfront/>.
- [3] 2018. DynamoDb. <https://aws.amazon.com/dynamodb/>.
- [4] 2018. Finding Pwned Passwords with IPassword. <https://blog.agilebits.com/2018/02/22/finding-pwned-passwords-with-1password/>.
- [5] 2018. Have I Been Pwned: API v2. <https://haveibeenpwned.com/API/v2>.
- [6] 2018. I Wanna Go Fast: Why Searching Through 500M Pwned Passwords Is So Quick. <https://www.troyhunt.com/i-wanna-go-fast-why-searching-through-500m-pwned-passwords-is-so-quick/>.
- [7] 2018. Interval Tree. https://en.wikipedia.org/wiki/Interval_tree.
- [8] 2018. List of data breaches. https://en.wikipedia.org/wiki/List_of_data_breaches.
- [9] 2018. Password Check. <https://security.googleblog.com/2019/02/protect-your-accounts-from-data.html>.
- [10] 2018. SECURITY UPDATE - Q2 2018. <https://www.eveonline.com/article/pc29kq/an-update-on-security-the-fight-against-bots-and-rmt>.
- [11] 2018. Twitter's List Of 370 Banned Passwords. <http://www.businessinsider.com/twitters-list-of-370-banned-passwords-2009-12>. Accessed: 2015-11-06.
- [12] 2019. Crypto Nodejs. <https://nodejs.org/api/crypto.html>.
- [13] 2019. GhostProject. <https://ghostproject.fr/>.
- [14] 2019. Testing Firefox Monitor, a New Security Tool. <https://blog.mozilla.org/futurereleases/2018/06/25/testing-firefox-monitor-a-new-security-tool/>.
- [15] 2019. Vericlouds. <https://my.vericlouds.com/>.
- [16] 4iQ. 2018. Identities in the Wild: The Tsunami of Breached Identities Continues. https://4iq.com/wp-content/uploads/2018/05/2018_IdentityBreachReport_4iQ.pdf/.
- [17] Carlos Aguilar-Melchor, Joris Barrier, Laurent Fousse, and Marc-Olivier Killijian. 2016. XPIR: Private information retrieval for everyone. *Proceedings on Privacy Enhancing Technologies* 2016, 2 (2016), 155–174.
- [18] Junade Ali. 2018. Optimising Caching on Pwned Passwords (with Workers). <https://blog.cloudflare.com/optimising-caching-on-pwnedpasswords>.
- [19] Junade Ali. 2018. Validating Leaked Passwords with k-Anonymity. <https://blog.cloudflare.com/validating-leaked-passwords-with-k-anonymity/>.
- [20] P. Berenbrink, T. Friedetzky, Z. Hu, and R. Martin. 2008. On weighted balls-into-bins games. *Theoretical Computer Science* 409, 3 (2008), 511–520.
- [21] Julio Casal. Dec, 2017. 1.4 Billion Clear Text Credentials Discovered in a Single Database. <https://medium.com/4iqdelvedeep/1-4-billion-clear-text-credentials-discovered-in-a-single-database-3131d0a1ae14>.
- [22] Hao Chen, Kim Laine, and Peter Rindal. 2017. Fast private set intersection from homomorphic encryption. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1243–1255.
- [23] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. 1995. Private information retrieval. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*. IEEE, 41–50.
- [24] Anupam Das, Joseph Bonneau, Matthew Caesar, Nikita Borisov, and XiaoFeng Wang. 2014. The Tangled Web of Password Reuse.. In *NDSS*, Vol. 14. 23–26.
- [25] Y. Dodis, L. Reyzin, and A. Smith. 2004. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. In *Eurocrypt 2004*, C. Cachin and J. Camenisch (Eds.). Springer-Verlag, 523–540. LNCS no. 3027.
- [26] Zeev Dvir and Sivakanth Gopi. 2015. 2-server PIR with sub-polynomial communication. In *Proceedings of the forty-seventh Annual ACM Symposium on the Theory of Computing*. ACM, 577–584.
- [27] Steven Englehardt, Jeffrey Han, and Arvind Narayanan. 2018. I never signed up for this! Privacy implications of email tracking. *Proceedings on Privacy Enhancing Technologies* 2018, 1 (2018), 109–126.
- [28] Verizon Enterprise. 2017. 2017 Data breach investigations report.
- [29] Michael J Freedman, Kobbi Nissim, and Benny Pinkas. 2004. Efficient private matching and set intersection. In *Advances in Cryptography—EUROCRYPT*. Springer, 1–19.
- [30] Ágnes Kiss, Jian Liu, Thomas Schneider, N Asokan, and Benny Pinkas. 2017. Private set intersection for unequal set sizes with mobile applications. *Proceedings on Privacy Enhancing Technologies* 2017, 4 (2017), 177–197.
- [31] Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. 2016. Efficient batched oblivious PRF with applications to private set intersection. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 818–829.
- [32] Marie-Sarah Lacharité and Kenneth G Paterson. 2018. Frequency-smoothing encryption: preventing snapshot attacks on deterministically encrypted data. *IACR Transactions on Symmetric Cryptology* 2018, 1 (2018), 277–313.
- [33] Yue Li, Haining Wang, and Kun Sun. 2016. A study of personal information in human-chosen passwords and its security implications. In *IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications*. IEEE, 1–9.
- [34] Jerry Ma, Weining Yang, Min Luo, and Ninghui Li. 2014. A Study of Probabilistic Password Models. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 689–704.
- [35] Ashwin Machanavajjhala, Johannes Gehrke, Daniel Kifer, and Muthuramakrishnan Venkatasubramanian. 2006. 1-diversity: Privacy beyond k-anonymity. In *22nd International Conference on Data Engineering (ICDE'06)*. IEEE, 24–24.
- [36] Catherine Meadows. 1986. A more efficient cryptographic matchmaking protocol for use in the absence of a continuously available third party. In *1986 IEEE Symposium on Security and Privacy*. IEEE, 134–134.
- [37] William Melicher, Blase Ur, Sean M Segreti, Saranga Komanduri, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. [n. d.]. Fast, lean and accurate: Modeling password guessability using neural networks.
- [38] A Narayanan and V Shmatikov. 2008. Robust de-anonymization of large datasets. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy, May 2008*.
- [39] Femi Olumofin and Ian Goldberg. 2011. Revisiting the computational practicality of private information retrieval. In *International Conference on Financial Cryptography and Data Security*. Springer, 158–172.
- [40] Bijeta Pal, Tal Daniel, Rahul Chatterjee, and Thomas Ristenpart. 2019. Beyond Credential Stuffing: Password Similarity using Neural Networks. *IEEE Symposium on Security and Privacy* (may 2019).
- [41] Benny Pinkas, Thomas Schneider, Gil Segev, and Michael Zohner. 2015. Phasing: Private set intersection using permutation-based hashing. In *24th {USENIX} Security Symposium ({USENIX} Security 15)*. 515–530.
- [42] Benny Pinkas, Thomas Schneider, Christian Weinert, and Udi Wieder. 2018. Efficient circuit-based PSI via cuckoo hashing. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 125–157.
- [43] Jennifer Pullman, Kurt Thomas, and Elie Bursztein. 2019. Protect your accounts from data breaches with Password Checkup. <https://security.googleblog.com/2019/02/protect-your-accounts-from-data.html>.
- [44] Kurt Thomas, Jennifer Pullman, Kevin Yeo, Ananth Raghunathan, Patrick Gage Kelley, Luca Invernizzi, Borbala Benko, Tadek Pietraszek, Sarvar Patel, Dan Boneh, and Elie Bursztein. 2019. Protecting Accounts from Credential Stuffing with Password Breach Alerting. In *USENIX Security Symposium*. USENIX.
- [45] Troy Hunt. 2018. Have I Been Pwned? <https://haveibeenpwned.com/Passwords/>.
- [46] Ding Wang, Zijian Zhang, Ping Wang, Jeff Yan, and Xinyi Huang. 2016. Targeted online password guessing: An underestimated threat. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. ACM, 1242–1254.
- [47] Dan Lowe Wheeler. 2016. zxcvbn: Low-budget password strength estimation. In *Proc. USENIX Security*.
- [48] Joanne Woodage, Rahul Chatterjee, Yevgeniy Dodis, Ari Juels, and Thomas Ristenpart. 2017. A new distribution-sensitive secure sketch and popularity-proportional hashing. In *Annual International Cryptology Conference*. Springer, 682–710.
- [49] Lei Zhang, Shushil Jajodia, and Alexander Brodsky. 2007. Information disclosure under realistic assumptions: Privacy versus optimality. In *Proceedings of the 14th ACM conference on Computer and communications security*. ACM, 573–583.

Distance	%
0	1.2
≤ 1	1.7
≤ 2	2.3
≤ 3	3.1
≤ 4	4.6

Figure 11: Statistics on samples with low edit distance between username and password, as a percentage of a random sample of 10^5 username-password pairs.

A CORRELATION BETWEEN USERNAME AND PASSWORDS

In Section 3 the username and password choices of previously uncompromised users can be modeled independently.

To check whether this assumption would be valid or not, we randomly sampled 10^5 username-password pairs from the dataset used in Section 6 and calculated the Levenshtein edit distance between each username and password in a pair. We have recorded the result of this experiment in Figure 11.

We found that the mean edit distance between a username and password was 9.4, while the mean password length was 8.4 characters and the mean username length was 10.0 characters. This supports that while there are some pairs where the password is almost identical to the username, a large majority are not related to the username at all.

B BANDWIDTH OF FSB

To calculate the maximum bandwidth used by FSB, we use the balls-and-bins formula as described in Section 3. Each password w is stored in $|\beta(w)|$ buckets, so the total number of balls, or passwords being stored, can be calculated as

$$\begin{aligned}
m &= \sum_{w \in \tilde{\mathcal{S}}} |\beta(w)| \\
&= \sum_{w \in \mathcal{W}_{\bar{q}} \cap \tilde{\mathcal{S}}} |\mathcal{B}| + \sum_{w \in \tilde{\mathcal{S}} \setminus \mathcal{W}_{\bar{q}}} \left\lceil \frac{|\mathcal{B}| \cdot \hat{p}_s(w)}{\hat{p}_s(w_{\bar{q}})} \right\rceil \\
&\leq |\mathcal{W}_{\bar{q}} \cap \tilde{\mathcal{S}}| \cdot |\mathcal{B}| + \sum_{w \in \tilde{\mathcal{S}} \setminus \mathcal{W}_{\bar{q}}} \left(\frac{|\mathcal{B}| \cdot \hat{p}_s(w)}{\hat{p}_s(w_{\bar{q}})} + 1 \right) \\
&\leq |\mathcal{B}| \cdot \bar{q} + |\mathcal{B}| \cdot \frac{1}{\hat{p}_s(w_{\bar{q}})} + N
\end{aligned}$$

The first equality is obtained by replacing the definition of $\beta(w)$; the second inequality holds because $\lceil x \rceil \leq x + 1$; the third inequality holds because $\tilde{\mathcal{S}} \subseteq W$.

The number of bins $n = |\mathcal{B}|$, and $m > n \log n$, if $\bar{q} > \log n$. Therefore, the maximum bucket size for FSB would with high probability be no more than $2 \cdot \left(\bar{q} + \frac{1}{\hat{p}_s(w_{\bar{q}})} + \frac{N}{|\mathcal{B}|} \right)$.

C PROOF OF THEOREM 4.2

Because the IDB bucketization scheme does not depend on the password, $\Pr[B = b \mid W = w \wedge U = u] = \Pr[B = b \mid U = u]$

$$\begin{aligned}
&\text{Adv}_{\text{IDB}}^{\text{b-gs}}(q) \\
&= \sum_u \sum_b \max_{w_1, \dots, w_q} \sum_{i=1}^q \Pr[W = w_i \wedge U = u] \cdot \Pr[B = b \mid U = u] \\
&= \sum_u \left(\sum_b \Pr[B = b \mid U = u] \right) \max_{w_1, \dots, w_q} \sum_{i=1}^q \Pr[W = w_i \wedge U = u] \\
&= \text{Adv}^{\text{gs}}(q)
\end{aligned}$$

The first step follows from independence of password and bucket choice, and the third step is true because there is only one bucket for each username.

D PROOF OF THEOREM 5.1

First we calculate the general form of the BucketGuess $_{\beta_{\text{FSB}}}$ advantage. Then, we show that for $q \leq \bar{q}$, $\text{Adv}_{\beta_{\text{FSB}}}^{\text{b-gs}}(q) = \text{Adv}^{\text{gs}}(q)$, and we bound the difference in the advantages for the games when $q > \bar{q}$.

$$\begin{aligned}
\text{Adv}_{\beta_{\text{FSB}}}^{\text{b-gs}}(q) &= \sum_u \sum_b \max_{\substack{w_1, \dots, w_q \\ \in \alpha(b)}} \sum_{i=1}^q \frac{\Pr[W = w_i \wedge U = u]}{|\beta_{\text{FSB}}(w_i)|} \\
&= \sum_b \max_{\substack{w_1, \dots, w_q \\ \in \alpha(b)}} \sum_{i=1}^q \frac{\hat{p}_s(w_i)}{|\beta_{\text{FSB}}(w_i)|}
\end{aligned}$$

The second step follows from the independence of usernames and passwords in the uncompromised setting.

We will use $\mathcal{W}_{\bar{q}}$ to refer to the top \bar{q} passwords according to password distribution $\hat{p}_s = p_w$, and $w_{\bar{q}}$ to refer to the \bar{q} th most popular password according to \hat{p}_s .

For $w \in \mathcal{W}_{\bar{q}}$, we can calculate the fraction in the summation exactly as $\frac{\hat{p}_s(w)}{|\beta_{\text{FSB}}(w)|} = \frac{\hat{p}_s(w)}{|\mathcal{B}|}$.

For any other $w \in \mathcal{W} \setminus \mathcal{W}_{\bar{q}}$, we can bound the fraction using the bound on the number of buckets a password is placed in.

$$\frac{|\mathcal{B}| \cdot \hat{p}_s(w)}{\hat{p}_s(w_{\bar{q}})} \leq |\beta_{\text{FSB}}(w)| < \frac{|\mathcal{B}| \cdot \hat{p}_s(w)}{\hat{p}_s(w_{\bar{q}})} + 1.$$

We can use the lower bound on $|\beta_{\text{FSB}}(w)|$ to find that

$$\frac{\hat{p}_s(w)}{|\beta_{\text{FSB}}(w)|} \leq \frac{\hat{p}_s(w_{\bar{q}})}{|\mathcal{B}|}.$$

Using the upper bound on $|\beta_{\text{FSB}}(w)|$,

$$\frac{\hat{p}_s(w)}{|\beta_{\text{FSB}}(w)|} > \frac{\hat{p}_s(w)}{\frac{|\mathcal{B}| \cdot \hat{p}_s(w)}{\hat{p}_s(w_{\bar{q}})} + 1} = \frac{\hat{p}_s(w) \cdot \hat{p}_s(w_{\bar{q}})}{|\mathcal{B}| \cdot \hat{p}_s(w) + \hat{p}_s(w_{\bar{q}})} = \frac{\hat{p}_s(w_{\bar{q}})}{|\mathcal{B}| + \frac{\hat{p}_s(w_{\bar{q}})}{\hat{p}_s(w)}}$$

Since the values of $\frac{\hat{p}_s(w)}{|\beta_{\text{FSB}}(w)|}$ are always larger for $w \in \mathcal{W}_{\bar{q}}$, the values of w_1, \dots, w_q chosen for each bucket will be the top \bar{q} passwords overall, along with the top $q - \bar{q}$ of the remaining passwords in the bucket, ordered by $\frac{\hat{p}_s(\cdot)}{|\beta_{\text{FSB}}(\cdot)|}$.

To find an upper bound on $\text{Adv}_{\beta_{\text{FSB}}}^{\text{b-gs}}(q)$,

$$\begin{aligned} & \sum_b \max_{\substack{w_1, \dots, w_q \\ \in \alpha(b)}} \sum_{i=1}^q \frac{\hat{p}_s(w_i)}{|\beta_{\text{FSB}}(w_i)|} \\ & \leq \sum_b \left(\sum_{w \in \mathcal{W}_{\bar{q}}} \frac{\hat{p}_s(w)}{|\mathcal{B}|} + (q - \bar{q}) \frac{\hat{p}_s(w_{\bar{q}})}{|\mathcal{B}|} \right) \\ & = \lambda_{\bar{q}} + (q - \bar{q}) \cdot p_{\bar{q}} \end{aligned}$$

For $q \leq \bar{q}$, we have $\text{Adv}_{\beta_{\text{FSB}}}^{\text{b-gs}}(q) \leq \lambda_{\bar{q}}$.

To find a lower bound on $\text{Adv}_{\beta_{\text{FSB}}}^{\text{b-gs}}(q)$, let $w_{\bar{q}+1}^*, \dots, w_q^*$ be the $q - \bar{q}$ passwords in $\alpha(b) \setminus \mathcal{W}_{\bar{q}}$ with the highest probability of occurring, according to $\hat{p}_s(\cdot)$.

$$\begin{aligned} & \sum_b \max_{\substack{w_1, \dots, w_q \\ \in \alpha(b)}} \sum_{i=1}^q \frac{\hat{p}_s(w_i)}{|\beta_{\text{FSB}}(w_i)|} \\ & > \sum_b \left(\sum_{w \in \mathcal{W}_{\bar{q}}} \frac{\hat{p}_s(w)}{|\mathcal{B}|} + \sum_{i=\bar{q}+1}^q \frac{\hat{p}_s(w_{\bar{q}})}{|\mathcal{B}| + \frac{\hat{p}_s(w_{\bar{q}})}{\hat{p}_s(w_i^*)}} \right) \\ & \geq \lambda_{\bar{q}} + \sum_{i=\bar{q}+1}^q \left[\frac{|\mathcal{B}| \cdot \hat{p}_s(w_i^*)}{\hat{p}_s(w_{\bar{q}})} \right] \cdot \frac{\hat{p}_s(w_{\bar{q}})}{|\mathcal{B}| + \frac{\hat{p}_s(w_{\bar{q}})}{\hat{p}_s(w_i^*)}} \\ & \geq \lambda_{\bar{q}} + \sum_{i=\bar{q}+1}^q \frac{|\mathcal{B}| \cdot \hat{p}_s(w_i^*)}{|\mathcal{B}| + \frac{\hat{p}_s(w_{\bar{q}})}{\hat{p}_s(w_i^*)}} \geq \lambda_{\bar{q}} + \sum_{i=\bar{q}+1}^q \frac{\hat{p}_s(w_i^*)}{1 + \frac{\hat{p}_s(w_{\bar{q}})}{\hat{p}_s(w_i^*) \cdot |\mathcal{B}|}} \\ & \geq \lambda_{\bar{q}} + \sum_{i=\bar{q}+1}^q \hat{p}_s(w_i^*)/2 \geq \lambda_{\bar{q}} + (\lambda_q - \lambda_{\bar{q}})/2 = \frac{\lambda_q + \lambda_{\bar{q}}}{2} \end{aligned}$$

Therefore, $\Delta_q \geq \frac{\lambda_q - \lambda_{\bar{q}}}{2}$.

Note, for every password to be assigned to a bucket, $|\mathcal{B}| \geq \hat{p}_s(w_{\bar{q}})/\hat{p}_s(w)$, or for all $w \in \mathcal{W}$, $\frac{\hat{p}_s(w_{\bar{q}})}{\hat{p}_s(w) \cdot |\mathcal{B}|} \leq 1$.